StorageWorks by Compaq

Heterogeneous Open SAN Design Reference Guide

Part Number: AA-RMPND-TE

Third Edition (February 2002)

This document is a guide to designing and building large Storage Area Networks (SANs). It describes how Compaq storage systems, storage management tools, and Fibre Channel products can be used in open heterogeneous SANs.

COMPAQ

© 2002 Compaq Information Technologies Group, L.P.

Compaq, the Compaq logo, and StorageWorks Registered in U. S. Patent and Trademark Office.

SANworks, Tru64, and OpenVMS are trademarks of Compaq Information Technologies Group, L.P. in the United States and other countries.

Microsoft, Windows, and Windows NT are trademarks of Microsoft Corporation in the United States and other countries.

Intel, is a trademark of Intel Corporation in the United States and other countries.

UNIX is a trademark of The Open Group in the United States and other countries.

All other product names mentioned herein may be trademarks of their respective companies.

Compaq shall not be liable for technical or editorial errors or omissions contained herein. The information in this document is provided "as is" without warranty of any kind and is subject to change without notice. The warranties for Compaq products are set forth in the express limited warranty statements accompanying such products. Nothing herein should be construed as constituting an additional warranty.

Compaq service tool software, including associated documentation, is the property of and contains confidential technology of Compaq Computer Corporation. Service customer is hereby licensed to use the software only for activities directly relating to the delivery of, and only during the term of, the applicable services delivered by Compaq or its authorized service provider. Customer may not modify or reverse engineer, remove, or transfer the software or make the software or any resultant diagnosis or system management data available to other parties without Compaq's or its authorized service provider's consent. Upon termination of the services, customer will, at Compaq's or its service provider's option, destroy or return the software and associated documentation in its possession.

Printed in the U.S.A.

Heterogeneous Open SAN Design Reference Guide
Third Edition (February 2002)
Part Number: AA-RMPND-TE

Contents

About This Guide

1 Understanding SAN Topologies

2 SAN Topologies

3 SAN Fabric Design Rules

4 Heterogeneous SAN Platform and Storage System Rules

5 Enterprise Backup Solution

6 SAN Management

7 SAN Security

A Supported Products

Glossary

Index

Figures

Tables

About This Guide

The following sections are covered:

- Text Conventions
- Symbols in Text
- Symbols on Equipment
- Getting Help
- Compaq Authorized Reseller

Text Conventions

This document uses the conventions in Table 1 to distinguish elements of text.

Table 1: Text Conventions

Element	Convention	Examples
• **Named Keys** • **Key Sequences**	**Bold**	**Home, Print Screen, Num Lock, Esc, PgUp** A plus sign (**+**) between two keys means that you should press them simultaneously: **Ctrl+A**, **Ctrl+Home**, **Alt+Ctrl+Del**
• Menu Items • Directory Names • Button Names • Dialog Box Names	Initial Caps (for UNIX directory names, the exact case of every character is displayed).	On the File menu, choose Save. Save the file in the C:\StorageSets\Default directory. UNIX: Save the file in the /home/newuser/practice directory. To back up files, click the Backup Now button. In the Save As dialog box, choose the drive then the folder.
filenames	Unless case sensitive, use *lowercase italics*. If filenames are case-sensitive (UNIX) or are easier to understand with some upper case letters, the exact case of each character is displayed.	To configure storage, edit *storageset.ini*. Changes are stored in *NewSystemConfigurationFile.ini*. (UNIX, AIX, Solaris): Errors are logged to *MixedCaseFile.txt.*
Menu Command Sequences	Initial Caps, with a right angle bracket (>) between items. Menu items are displayed as shown on screen.	To compare documents, choose: Tools > Documents > Compare.

Table 1: Text Conventions (Continued)

Element	Convention	Examples
• User Input and System Responses (Output and Error Messages) • COMMAND NAMES • Drive Names	`Monospace font`. `COMMAND NAMES` appear in upper case, unless they are case sensitive (UNIX command names are case sensitive and will not appear in uppercase). Entered `<variables>` are displayed in angle brackets (< >) and all lower case.	User Input and System Responses: • To exit from the program, type `Exit`. • At the prompt, type this command: `SHOW THIS_CONTROLLER` (no variable) • To see your settings, give the command: `SHOW <storagesets> FULL` (with variable) • You will see the `Continue?` message. Command Names • Use `SET THIS_CONTROLLER` to change parameters. • To manage storage, enter `RUN` *`sysmgr.exe`* • UNIX: To list files, give the `ls` command. • Drive Names: Navigate to your CD-ROM drive (usually `D:` or `E:`).
URLs	Sans serif font.	For update notices, visit: http://www.compaq.com/products/updates

Symbols in Text

These symbols may be found in the text of this guide. They have the following meanings.

WARNING: Text set off in this manner indicates that failure to follow directions in the warning could result in bodily harm or loss of life or damage to equipment.

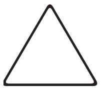

CAUTION: Text set off in this manner indicates that failure to follow directions could result in damage to equipment or loss of information.

IMPORTANT: Text set off in this manner presents clarifying information or specific instructions.

NOTE: Text set off in this manner presents commentary, sidelights, or interesting points of information.

Symbols on Equipment

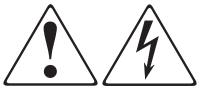

Any enclosed surface or area of the equipment marked with these symbols indicates the presence of electrical shock hazards. Enclosed area contains no operator serviceable parts.

WARNING: To reduce the risk of injury from electrical shock hazards, do not open this enclosure.

Any RJ-45 receptacle marked with these symbols indicates a network interface connection.

WARNING: To reduce the risk of electrical shock, fire, or damage to the equipment, do not plug telephone or telecommunications connectors into this receptacle.

Any surface or area of the equipment marked with these symbols indicates the presence of a hot surface or hot component. Contact with this surface could result in injury.

WARNING: To reduce the risk of injury from a hot component, allow the surface to cool before touching.

Power supplies or systems marked with these symbols indicate the presence of multiple sources of power.

WARNING: To reduce the risk of injury from electrical shock, remove all power cords to completely disconnect power from the supplies and systems.

Any product or assembly marked with these symbols indicates that the component exceeds the recommended weight for one individual to handle safely.

WARNING: To reduce the risk of personal injury or damage to the equipment, observe local occupational health and safety requirements and guidelines for manually handling material.

Related Documents

Topic	Document Title
Tru64 UNIX	HSG80 ACS Solution Software Version 8.6 for Compaq Tru64 UNIX Installation and Configuration Guide
Tru64 UNIX	Tru64 UNIX Kit V1.0 for Enterprise Virtual Array Installation and Configuration Guide
Compaq OpenVMS	OpenVMS Kit V1.0 for Enterprise Virtual Array Installation and Configuration Guide
Compaq OpenVMS	HSG80 ACS Solution Software Version 8.6 for Compaq OpenVMS Installation and Configuration Guide
Windows NT (Intel), Windows 2000	Windows NT and 2000 Kit V1.0 for Enterprise Virtual Array Installation and Configuration Guide
Windows NT (Intel), Windows 2000	HSG80 ACS Solution Software Version 8.6 for Windows NT and Windows 2000 Installation and Configuration Guide

Topic	Document Title
Windows NT (Intel), Windows 2000	Enterprise/Modular Storage RAID Array Fibre Channel Cluster for Windows NT/Windows 2000 Installation Guide
Windows NT (Intel), Windows 2000	Fibre Channel External Boot Support for Windows NT and Windows 2000 Application Note
Windows NT (Intel), Windows 2000	RAID Array 4100 User Guide
Windows NT (Intel), Windows 2000	RA4100 SAN Solution User Guide
Windows NT (Intel), Windows 2000	MSA1000 User Guide
Sun Solaris	Sun Solaris Kit V1.0 for Enterprise Virtual Array Installation and Configuration Guide
Sun Solaris	HSG80 ACS Solution Software Version 8.6 for Sun Solaris Installation and Configuration Guide
HP-UX	HSG80 ACS Solution Software Version 8.6 for HP-UX Installation and Configuration Guide
HP-UX	HP-UX Fibre Channel Hub Application Note
IBM AIX	HSG80 ACS Solution Software Version 8.6 for IBM AIX Installation and Configuration Guide
Linux	HSG80 ACS Solution Software Version 8.6 for Linux X86 and Alpha Installation and Configuration Guide
Linux	RA4100 SAN Solution User Guide
Novell NetWare	HSG80 ACS Solution Software Version 8.6 for Novell NetWare Installation and Configuration Guide
Novell NetWare	RAID Array 4100 User Guide
Novell NetWare	RA4100 SAN Solution User Guide
SGI IRIX	HSG80 ACS Solution Software Version 8.6 for SGI IRIX Installation and Configuration Guide
HSV110 Controller	Enterprise Virtual Array HSV Controller User Guide
HSG80 Controller	HSG80 Array Controller ACS Version 8.6 CLI Reference Guide
StorageWorks Command Console	Command Console V2.4 User Guide
Enterprise Backup Solution	Enterprise Backup Solution Reference Guide
Fibre Channel SAN Switch	Fibre Channel SAN Switch Management Guide
Fibre Channel Storage Switch, 8 & 16 Port	StorageWorks Fibre Channel SAN Switch Installation and Hardware Guide
StorageWorks Command Console	Command Console for the SAN Switch Installation Guide
Fibre Channel SAN Switch Configurations	StorageWorks Combining 16-Port Switches to Construct Higher Port Count Switches Application Note

Getting Help

If you still have a question after reading this guide, contact service representatives or visit our website.

Compaq Technical Support

In North America, call the Compaq technical support at 1-800-OK-COMPAQ. This service is available 24 hours a day, 7 days a week.

NOTE: For continuous quality improvement, calls may be recorded or monitored.

Outside North America, call Compaq technical support at the nearest location. Telephone numbers for worldwide technical support are listed on the Compaq website: http://www.compaq.com.

Be sure to have the following information available before you call Compaq:

- Technical support registration number (if applicable)
- Product serial numbers
- Product model names and numbers
- Applicable error messages
- Operating system type and revision level
- Detailed, specific questions

Compaq Website

The Compaq website has the latest information on this product as well as the latest drivers. Access the Compaq website at: http://www.compaq.com/storage. From this website, select SANworks.

Compaq Authorized Reseller

For the name of your nearest Compaq Authorized Reseller:

- In the United States, call 1-800-345-1518.
- In Canada, call 1-800-263-5868.
- Elsewhere, see the Compaq website for locations and telephone numbers.

1

Understanding SAN Topologies

Guide Overview

This guide describes scalable heterogeneous Open StorageWorks SANs by Compaq. A Storage Area Network (SAN) is a network of shared storage resources that can be allocated across a heterogeneous environment. SANs connect multiple heterogeneous servers to single or multiple shared storage systems using Fibre Channel switches.

SANs redefine previous concepts of computer storage. Recent advances in technology make it possible to organize and manage storage as a resource independent of the local system. SANs enable new ways to configure and use storage. A SAN keeps data preserved in a common pool or multiple pools, where it can be easily accessed and managed. This method of storage reduces operational costs, supports business growth, and strengthens the corporate infrastructure.

Heterogeneous Open StorageWorks SANs by Compaq can be configured to provide:

- Data protection
- High availability
- Long distance accessibility
- High connectivity with high bandwidth
- Common storage access for multi-vendor platforms
- Storage consolidation
- Local, centralized, or distributed storage access
- Centralized backup and restore
- Centralized storage management
- Disaster tolerance

A heterogeneous, scalable SAN is the answer for the increasingly complex and exponentially growing storage needs of today’s businesses. SANs provide effective centralized management of primary storage (disks) and secondary storage (tapes) at a lower total cost. The topologies and design rules described in this guide will help you design a SAN that will fill your immediate business needs and enable your SAN to be scaled in incremental steps as your business needs demand expanded capacity.

This Chapter

This chapter introduces the concepts of SAN design and SAN topologies. It first provides context for SAN design by describing the design philosophy for StorageWorks SANs by Compaq. Next, it provides general information about several of the basic SAN topologies. Finally, it explores the SAN design considerations that result when you apply the design philosophy to the SAN topologies.

Understanding Heterogeneous Open SANs

About the Design Philosophy

The SAN design philosophy maximizes the features and functionality provided by StorageWorks by Compaq SAN products. These products support an Open SAN design through cooperating combinations of moderately sized components. The use of varying numbers of smaller subsystem components:

- Allows greater flexibility in SAN design
- Provides for incremental scaling over time
- Accommodates diverse geographic and data locality requirements

Open SAN design is facilitated through the use of multiple port functionality and simplified design approaches.

Multiple Port Functionality

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

  For example, Compaq Fibre Channel switches have multiple ports and can be interconnected across long distances to achieve large network configurations. This connectivity is also known as a Fibre Channel "fabric." Each fabric has ports into which several computer servers, storage systems, and related components can be integrated.

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

  Each Compaq controller has multiple ports for connection to the fabric. In addition, each controller can support a heterogeneous mixture of servers, which means that controllers and servers can be added to the existing infrastructure 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

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

1. **Compaq standard topology design**.

   Compaq standard topologies 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 a Compaq design**.

   Each of the Compaq 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 implement it to meet the data locality and connectivity requirements. By varying a standard design, you use the SAN experience of Compaq to leverage your own design efforts. This approach is recommended for anyone with an intermediate level of SAN experience.

3. **Custom design using the StorageWorks by Compaq 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 StorageWorks SANs by Compaq. 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.

Types of SAN Topologies

There are four basic Compaq standard SAN topologies divided into two types of multi-switch SAN Fabric implementations.

The first type of implementation has all switches connected to servers or storage. This type includes:

- Cascaded Fabric
- Meshed Fabric
- Ring Fabric

The other type of implementation has some switches that connect mainly to other switches. This type of implementation includes:

- Backbone SAN Fabrics

Cascaded, Meshed, and Ring SAN Fabrics

The first three fabric topologies (Refer to Figure 1–1) are SAN topologies where all switches in the fabric are utilized for connecting servers and storage. For more details on these topologies refer to Chapter 2, "SAN Fabric Topologies". Typically, in these types of fabric arrangements, a small percentage of the total number of switch ports is used for inter-switch connectivity. Refer to Connectivity in this chapter. Trade-offs can be made, allowing for an increase in the number of Inter-Switch Links (ISLs), by reducing the total number of ports used for servers and storage.

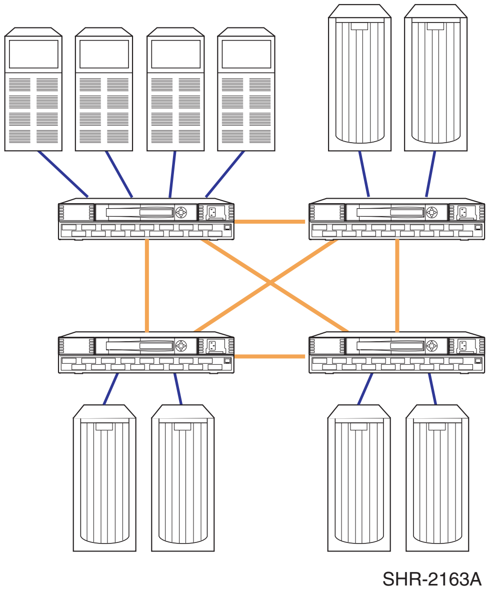

Figure 1–1: Meshed Fabric

Backbone SAN Fabric

Backbone SAN Fabrics (See Figure 1–2) are implementations where one or more Fibre Channel switches within the fabric are used primarily for connection to other switches. The server and storage connections in a Backbone SAN Fabric design are concentrated on edge switches. These switches are attached to a central backbone switch or switches. Refer to Chapter 2, "Tree Backbone Fabrics".

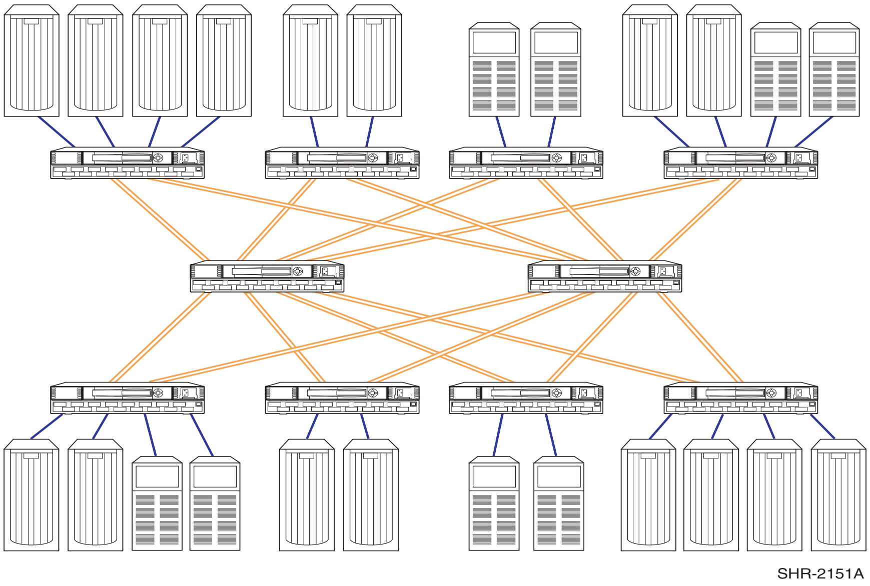

Figure 1–2: Backbone SAN - Tree Backbone Fabric

Topology Variations

In addition to these base topology designs, other variations are possible. For example, in a Backbone SAN Fabric design it may be desirable to locate centralized primary storage or centralized secondary storage, such as tape backup, directly on a backbone switch. Conversely, in a SAN Fabric design it could be desirable to utilize a high number of inter-switch links between two adjacent switches, based on specific performance and access requirements.

Other variations of Compaq standard topologies are possible, provided that the SAN design rules are followed.

Custom-Designed SAN Topologies

Compaq standard SAN topologies or subsets of these topologies can meet most SAN implementation requirements. There may be specific cases where Compaq standard topologies (or variants) do not meet your specific needs or requirements. In these cases, a custom SAN design can be created provided that 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 9, "Best Practices".

Design Considerations

A SAN design starts with requirements, a design philosophy, and the standard topologies and design rules. You create a SAN by making choices and reconciling these three elements. The process of creating a SAN presents many design considerations. 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

Whether choosing a Compaq standard SAN topology or creating a custom design, you should thoroughly review each of the design considerations described in this guide. Some items may be more or less important to individual designs, but each should be carefully evaluated with regard to expected future needs and growth, in addition to current needs.

Geographic Layout

The physical layout of campuses, facilities, and the location of servers and storage within individual buildings can be a major factor in determining the appropriate SAN design. Support for long distances throughout the SAN can easily accommodate specific and varying geographic needs. A StorageWorks by Compaq SAN can be implemented with multiple inter-switch cable segment distances of up to 10 km each for a total distance of 70 km, or when combined with a single 100 km very long distance segment, up to 160 km total distance across the SAN.

Support for these long distances also provides for interconnection of existing independent SAN islands, that may be geographically separated, into a single long distance SAN. Refer to Chapter 9, "Merging SAN Fabrics" for further information.

Data Locality

A major factor in determining the optimal SAN topology design is the desired or required location or method of storage deployment. Storage 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.

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 servers and storage connected to the same Fibre Channel switch. Centralized or "many-to-one" is where the primary data access is between many servers and centrally located storage. Distributed or "many-to-many" is where data access varies between many different servers and many different storage systems. Many-to-many data access is conducive to applications or deployments where you wish to implement 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 SAN fabric.

The need for connectivity directly affects the total number of Fibre Channel switches required in a SAN. Including data locality and geographic requirements, the total number of ports required for connectivity of servers and storage is a major consideration when evaluating different SAN topology designs. You should consider future connectivity requirements and choose a design that can be scaled or that can migrate to a topology design with more capacity. Your design should also provide adequate fabric performance by implementing the required number of ISLs.

If the total number of ports required exceeds a given topology design, you should consider higher capacity topologies or perhaps deploy multiple independent SANs.

Storage Capacity

The total storage capacity requirements, both present and future, should be calculated to ensure that the design is adequate to meet your needs. There are two aspects to storage capacity. The first is adequate size. Storage size can be expanded by adding larger capacity disk devices, adding additional disk devices, 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 may be 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

Compaq 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. Compaq 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. Refer to Appendix A, "Heterogeneous Open SAN Products".

Scalability and Migration

A major benefit provided by Compaq 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.

Compaq-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 9, "Best Practices", for information about scaling and migrating different SAN topology designs, as some transitions are easier to perform then 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" and Appendix A, "Supported Products".

Data Availability

Data availability is a measure of how reliably data can be accessed despite the normal failure of SAN components.

SAN implementations can provide different levels of data availability. There are varying levels of availability designs ranging from the simplest single component and single path designs up to the highest availability, or No Single Point Of Failure (NSPOF) designs. In many cases the different levels can be implemented within the same SAN allowing mixed data protection levels, depending on the level of protection required for specific applications or data. For further information, reference Chapter 2, "Data Availability in a SAN".

Disaster Tolerance

Disaster tolerance is a measure of how reliably data can be accessed and restored despite multiple simultaneous failures of SAN components, often localized to a single site or locality in the SAN.

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.

NOTE: Disaster tolerance requires a SAN design that implements two or more separate fabrics, which provides the highest availability no single point of failure protection.

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 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.

All Compaq Fibre Channel switches implement a "non-blocking" design. That is, any pair of ports can be active and transfer data without impacting data transfer between another pair of ports on the switch. This feature should be carried throughout the design of the fabric itself. Oversubscription must be avoided in order to have a “non-blocking” fabric.

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 will 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 your fabric design provides for an adequate number of ISLs between all switches, and also minimizes scenarios where many devices or ports are attempting to access a single destination device or 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. Consider the types of applications that will be utilized on the SAN relative to data locality classification. Applications can usually be classified as high bandwidth or high throughput. What is important is that any design chosen provides an adequate level of performance based on the data locality classification of the major applications being utilized.

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 SAN Appliance, regardless of the arrangement or location of the storage components, and the preferred data access method. The SAN Appliance connects directly to the SAN through a Fibre Channel switch, providing it with high bandwidth and automatic connection to all devices in the SAN. The SAN hardware devices can be monitored and managed utilizing SAN Web-based tools either resident or invoked through the SAN Appliance. For implementations not utilizing the SAN Appliance, tools can be run on a local server or client configured for access in the SAN. Refer to Chapter 6, "SAN Management".

Fabric Zoning

Zoning is a fabric management service used to define logical device subsets within a SAN. Zoning enables resource partitioning for management and access control. The Compaq Fibre Channel switch zoning feature provides a way to control SAN access at the node device or port level. Zoning can be used to separate one physical fabric into many logical fabrics consisting of selected server and storage devices or ports. 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 and the number of zones and zone membership are effectively unlimited. Nodes can be in multiple zones to allow overlapping, depending on the desired access control.

Selective Storage Presentation

Compaq storage systems employ an exclusive LUN selection or masking feature called Selective Storage Presentation (SSP). This feature allows you to assign or selectively present storage sets (LUNs) from a Compaq storage system to multiple servers and host bus adapters of differing types in a 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

Compaq is addressing the issue of security over a SAN. StorageWorks by Compaq SAN hardware and SAN management tools provide reliable access to data, robust data storage, and enforcement of data access restrictions. For more information, refer to Chapter 7, "SAN Security".

Summary

SAN design requires the consideration of many factors. You provide the requirements for your SAN. Compaq provides a design philosophy and standard SAN topologies that come together in implementations to meet your requirements. In designing your SAN you must consider physical factors like:

- Geographic layout
- Connectivity
- Storage capacity

You must also consider abstract factors like:

- Operating systems
- Data availability
- Manageability

With its design philosophy, design rules, and standard topologies, Compaq has provided many reusable elements to aid you in the SAN design process. This allows you to focus on the design considerations that are unique and most relevant to your particular requirements.

2

SAN Topologies

Overview

This chapter describes the Compaq standard SAN topologies. You should review each of the SAN design considerations previously described in the first chapter before you select an appropriate topology design. The design considerations enable you to generate an accurate 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 a Compaq standard SAN topology design, a subset or variation of a Compaq design, or you can design a custom SAN topology. With all methods, 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".

Regardless of the approach taken, you should review the Compaq standard SAN topologies section in this chapter. This provides a greater understanding of the different aspects of SAN implementation. Compaq recommends that you first consider implementing one of the Compaq standard topologies or a variation of one of these designs. If your requirements cannot be met by one of these topology designs, then you can implement a customized SAN topology design, provided you follow the design rules.

Compaq Standard SAN Topologies

The Compaq topology designs reflect the proper application of the SAN design rules. Each of the topology designs is tailored for particular data access and connectivity needs. Collectively these designs provide a wide range of options for selecting the appropriate SAN design for your specific requirements. Variations of these designs can be validated by adhering to the appropriate rule set for each topology type.

The different types of Compaq standard SAN topologies are described in detail in the following sections. Consider the advantages of each design category before choosing a specific topology design.

SAN Fabric Topologies

SAN Fabric topology designs include:

- Cascaded Fabrics: See the section Cascaded Fabrics.
- Meshed Fabrics: See the section Modified Meshed Fabrics.
- Ring Fabrics: See the section Ring Fabric.
- Tree Backbone Fabrics: See the section Tree Backbone Fabrics.

The smallest SAN consists of a single Fibre Channel switch, server, and storage system. You can scale a SAN Fabric up to 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 topologies as a way to connect existing smaller SANs or SAN islands. If you have already deployed multiple small SANs, you can connect them into larger fabrics up to the maximum fabric size shown for each of the topology types. For example, 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 2–2, "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 as in the Tree Backbone fabric topology.

If more capacity is required, you can deploy multiple independent SANs. As larger fabric configurations are supported over time, the independent SANs may be interconnected to form even larger SANs, provided the maximum fabric support limit rules are followed.

Advantages of Using Cascaded, Meshed, or Ring SAN Fabric Designs

Each of the design types can be:

- Implemented as a separate SAN for specific departments or applications within a large 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 10 km, or even 100 km.
- Used to begin deploying 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 all Availability levels. See “Levels of Availability”.
- Upgraded to higher capacity topologies or topologies optimized for different data access types if needs change.

Cascaded Fabrics

A cascaded fabric SAN (See Figure 2–1) is a string of switches or levels of switches connected together by one or more ISLs. The switches are arranged in a linear array, each one connected to the switch that is next in line, or arranged in a vertical cascade with multiple levels off a single top switch.

Cascaded fabric designs are well suited to applications where data access is local relative to the same switch. Access requirements for a server (or groups of servers) are typically to the same storage system s or storage sets that are attached to the same switch. Groups of servers and the storage being accessed can be connected to the same switches providing the highest level of performance. Cascading provides a means to scale the SAN for additional connectivity of servers and storage, and allows for centralized management and backup.

Cascading 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

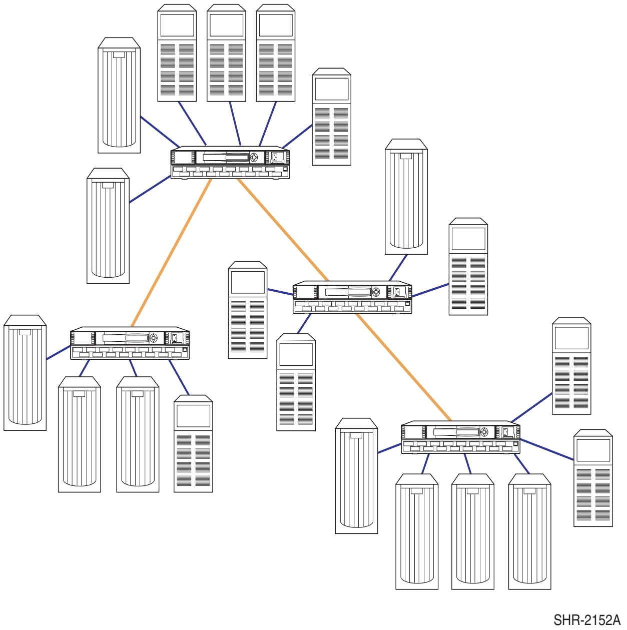

Figure 2–1: Cascaded Fabric SAN

in the fabric. Compaq 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..

Advantages of a Cascaded Fabric

- 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

Modified Meshed Fabrics

A meshed fabric is similar to a cascaded fabric, however, in a meshed fabric design, all switches are interconnected so there are at least two paths or routes from any one switch to any other switch in the fabric. This is true even if it is implemented with single ISLs between switches. This type of connectivity provides a level of fabric resiliency. If a single ISL or ISL switch 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.

In a full mesh design with all switches connected to all other switches, the efficiency of fabric relative to available ports can decrease as more switches are added. The efficiency of this fabric design can be improved by implementing a slightly modified mesh design, (See Figure 2–2).

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 efficiency.

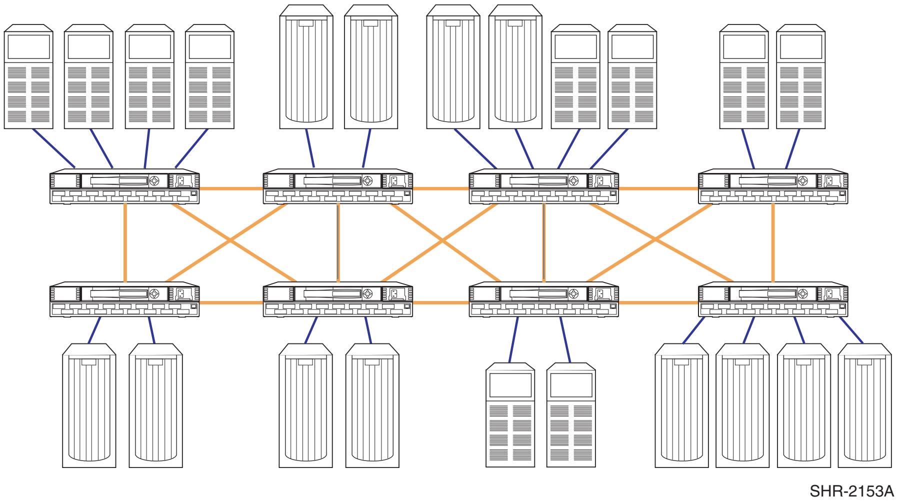

Figure 2–2: Modified Meshed Fabric SAN

Meshed fabrics are well suited to applications where data access is a mix of local and distributed. The full connectivity ensures many-to-many access, while at the same time allowing localized access to individual switches, servers and storage.

Advantages of a Meshed Fabric

- 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.
- 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 SANworks 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
  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.

Ring Fabric

A ring fabric (See Figure 2–3, "Ring Fabric SAN") is a continuous ring of switches connected together into a single fabric. Each switch is connected to adjacent switches on either side 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. The maximum number of switches supported in the ring of a Ring fabric is 14, which results in a maximum switch hop count of 7 under normal circumstances. (Because the ring provides connectivity in two directions, any two devices are never more than 7 hops apart in a 14-switch ring). If using less than 14-switches in a ring, additional switches can be added outside of the ring. For example, in a 10-switch ring, 10 additional switches can be added provided that each additional switch outside of the ring is connected directly to a switch in the ring. This ensures that the 7 hop limit is maintained.

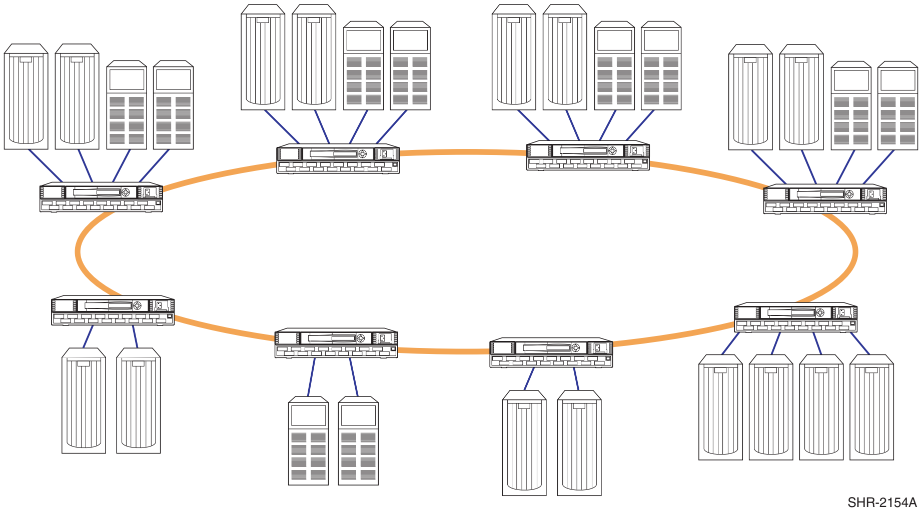

Figure 2–3: Ring Fabric SAN

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 vertical through the same 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.

The 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 within the ring provides a way to centrally manage the SAN and allows for centralized backup. This design allows data access through the ring, if required, provided an adequate number of ISLs are specified in the design. The ring fabric is not recommended for applications that require many-to-many connectivity.

Advantages of a Ring Fabric

- 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 SANworks tools can navigate and manage the Storage Area Network in the ring fabric, 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 SAN Fabrics

A Backbone SAN Fabric design has one or more Fibre Channel switches acting as fabric backbones dedicated to connecting to other switches within the fabric. The backbone switches act as routers with full bandwidth and redundant connectivity to all other switches. This type of implementation offers the best "many-to-many" connectivity and evenly distributed bandwidth throughout the fabric.

The design is well-suited for implementations where data traffic patterns may vary and even be random at times, but the underlying need is for full network “many-to-many” connectivity with high performance. This topology is also ideal if you require or plan to implement in the future, SAN-wide storage pooling and sharing.

Tree Backbone Fabrics

See Figure 2–4. A Tree Backbone fabric is a three-level Backbone SAN Fabric implementation where the center level switch or switches are dedicated as backbone switches. The backbone switches connect to all edge switches on the upper and lower tiers in the fabric. Servers and storage can be connected to any of the edge switches on either tier to maximize the flexibility of connections.

If required, design trade-offs can be made that allow for the connection of centralized primary or secondary storage directly on the backbone switches (this may reduce access or available bandwidth to specific portions of the fabric). Future capabilities will allow interconnecting Backbone SANs using ports on the backbone switches, or additional dedicated backbone switches to build larger fabrics.

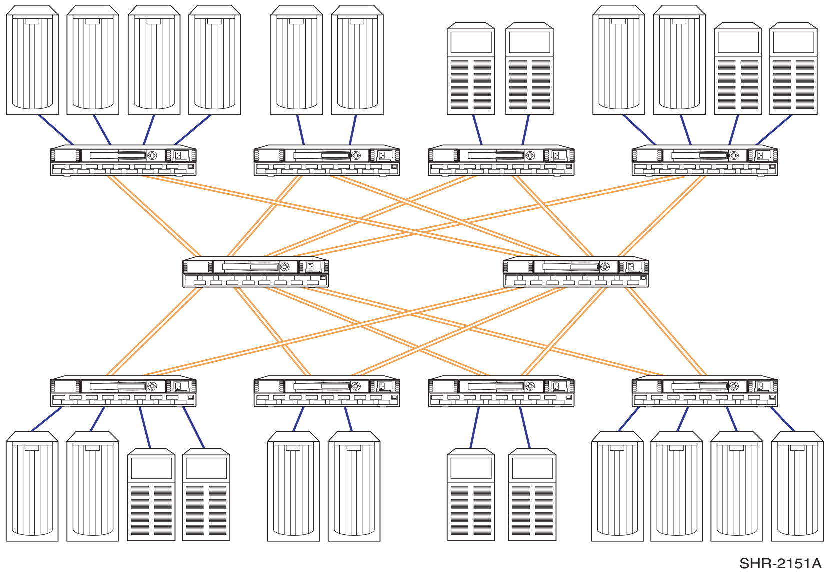

Figure 2–4: Tree Backbone SAN

StorageWorks by Compaq SAN fabrics currently support a large number of switches in a fabric. See Figure 2–5 for an example diagram of a Tree Backbone SAN with 20 switches.

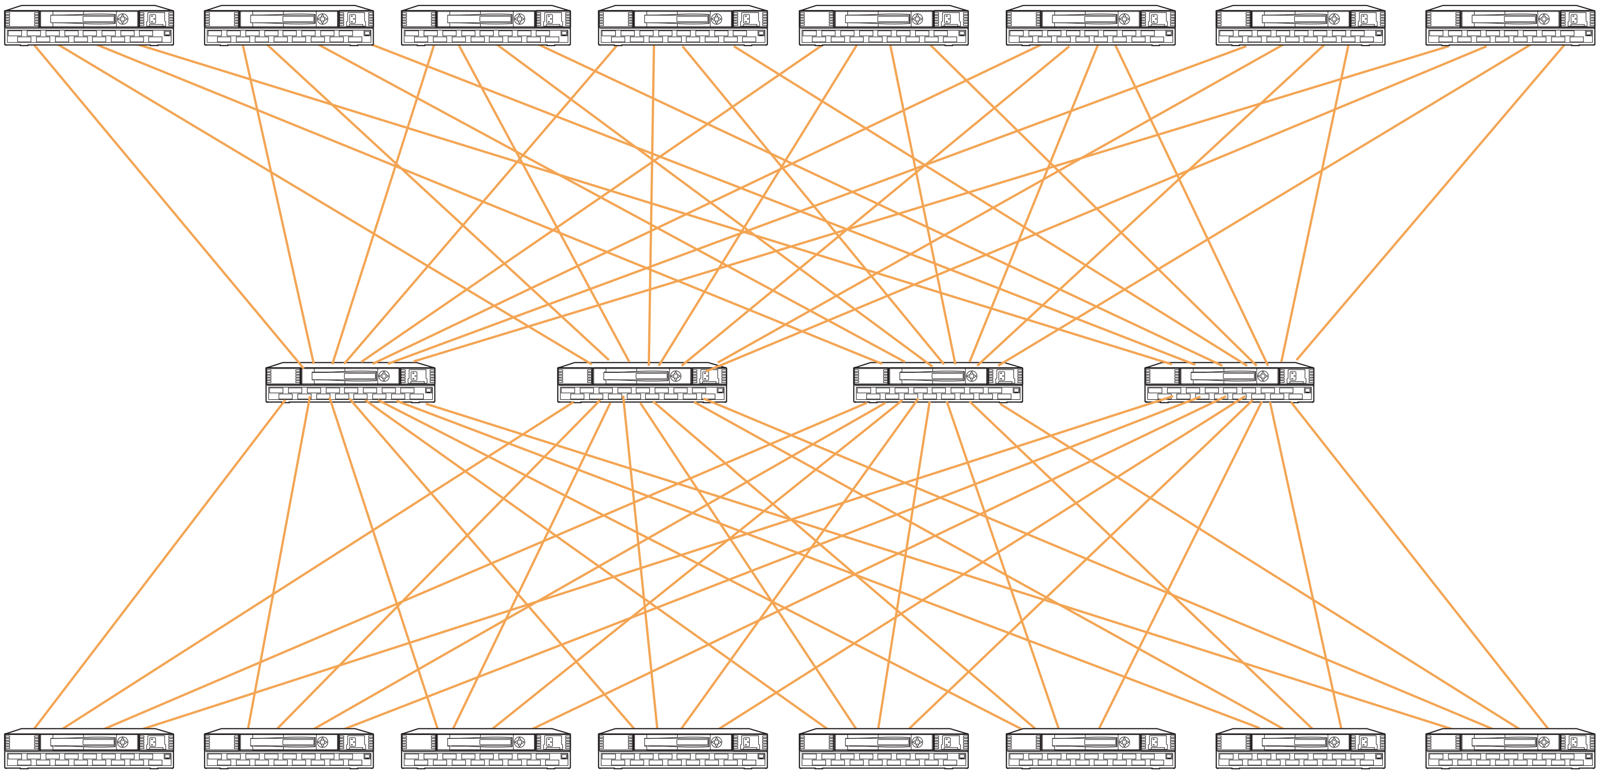

Figure 2–5: Tree Backbone SAN with 20 Switches

Fat Tree and Skinny Tree Designs

Fat trees and skinny trees are two types of backbone SAN topologies. The main 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 uses 50% of the edge switch ports as ISL connections while skinny trees use less than 50%. 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 end ports. In Figure 2–6, six 16-port switches in a skinny tree configuration yield 64 ports while the same switches in a fat tree yield only 32 ports (See Figure 2–7).
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.

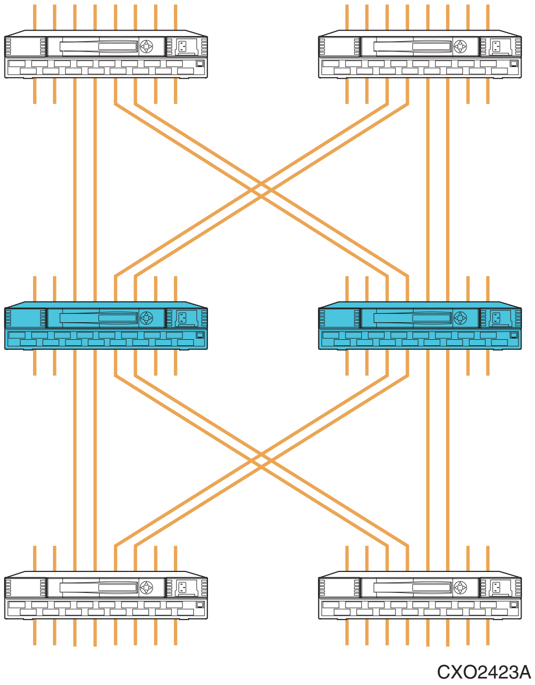

Figure 2–6: 64-Port Skinny Tree

The configuration for Figure 2–6 has 6 switches, 2 of which are backbone switches (shaded) and 4 edge switches. The 4 edge switches each have 12 available end ports 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 end ports is 64.

There are 8 ISLs on either side of the backbone switches. The cross sectional bandwidth is approximately 8 times 100 MB/s, or 800 MB/s.

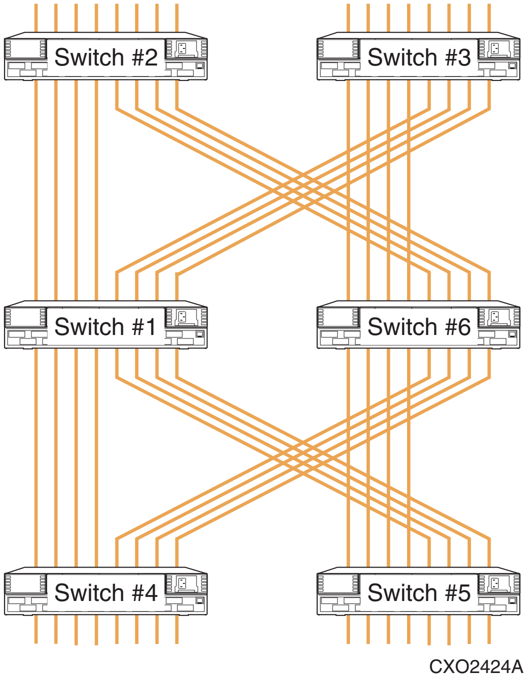

Figure 2–7: 32-Port Fat Tree

The configuration for Figure 2–7 also has 6 switches, with 2 backbone switches and 4 edge switches. The 4 edge switches each have 8 available end ports to connect to servers or storage. No ports are available on the backbone switches for end ports, so the total number of end ports is 32.

The cross sectional bandwidth for this fat tree is approximately 1600 MB/s.

For the SAN fabric, additional ISL connections must be added as requirements increase for cross-sectional SAN bandwidth. The SAN fabric must have enough ISL connections to limit contention or oversubscription for these ISL connections and also to provide the expected cross-sectional bandwidth. Contention or oversubscription for ISL connections describes the funnel effect of many ports trying to traverse a limited number of ISL connections. If you have 12 end ports on an edge switch trying to transfer data to other switches and only 4 ISLs then your ISL connections have a 3 to 1 oversubscription. There are 3 end ports contending for each ISL connection.

Skinny tree configurations which have more input ports (end ports) than output ports (ISLs) are more susceptible to oversubscription than fat trees, which have double the number of ISLs.

Compaq provides pre-configured fat tree and skinny tree SAN topologies as ready to deploy products. The SAN Switch Integrated 32 is a pre-configured 6-switch fat tree SAN topology, the SAN Switch Integrated 64 is a pre-configured 6-switch skinny tree SAN topology. Both products house the 6 switches in a rack mountable enclosure with all inter-switch cabling pre-connected. Available edge switch device 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.

Advantages of Tree 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 SANworks 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 deployed or distributed across a wide area with multiple inter-switch distances up to 10 km each or a single long distance switch separation of 100 km.

  Accommodates diverse geographic needs while supporting a wide variety of servers and storage needs and traffic 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

Each of the SAN topologies can be characterized by how well they provide specific data access. Refer to Chapter 1, "Data Locality". Table 2–1 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 the 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 2–1: Topology Usage Rating

	Data Locality		
SAN Topology	Local “one-to-one”	Centralized “Many-to-One”	Distributed “Many-to-Many”
Cascaded	Highest	Not Recommended	Not Recommended
Mesh	Medium	Medium	High
Ring	Highest	Medium	Not Recommended
Skinny Tree Backbone	Medium	High	High
Fat Tree Backbone	High	Highest	Highest

Topology Maximums

Table 2–2 indicates the maximum number of switches and ports supported for each of the Compaq standard SAN topologies.

NOTE: The maximums shown presume 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 SANworks Management Appliance also reduces the total number of ports available for servers and storage. See Chapter 6, "SANworks Management Appliance".

Table 2–2: Topology Maximums

SAN Topology Min/Max Number of Switches	Total Number of Ports		Maximum Number of Device Ports[1]	
	Single Fabric	Two Fabrics	Single Fabric	Two Fabrics
Cascaded, Meshed, Ring[2] Single Fabric: 2 to 20 Two Fabrics: 4 to 40	320	640	~250	~500
Tree Backbone Single Fabric: 5 to 20 Two Fabrics: 10 to 40			192	384

[1] Assumes 16-port switches. Indicates the number of ports available for server and storage connectivity. For Cascaded, Meshed, and Ring fabrics the number of device ports is approximate since this is dependent on the specific switch arrangement and the number of ISLs utilized.

[2] The maximum number of switches supported in a ring is 14; additional switches can be added if you reduce the number of switches in the ring. A maximum of 20 switches can be implemented in a single fabric if there are no more than 10 switches in the ring within the fabric, and no more than 7 hops between any two devices in the fabric.

Data Availability in a SAN

Data availability in a SAN can be influenced by many factors. The fabric architectural design, the number of Fibre Channel switches, and the number of ISLs between switches can all have a direct effect on the fabric 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 affects data availability and accessibility, as well as performance.

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 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 Path

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, but does offer the highest connectivity level relative to port count.

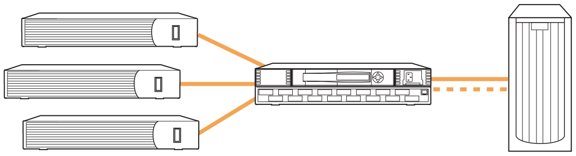

Figure 2–8: Level 1: Maximum Connectivity

Level 2: Single Meshed or Cascaded Fabric/Single Server and Storage Path

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.

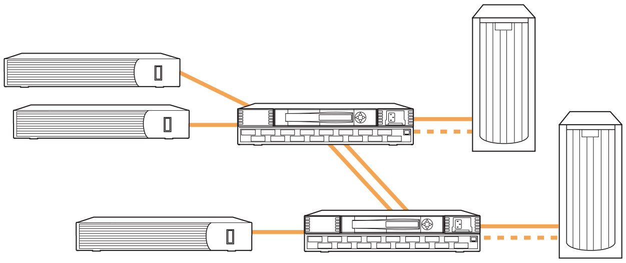

Figure 2–9: Level 2: Fabric Resiliency

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 controller when set in multiple-bus fail over mode. Each controller must be cabled to a different switch. Refer to Figure 2–10.

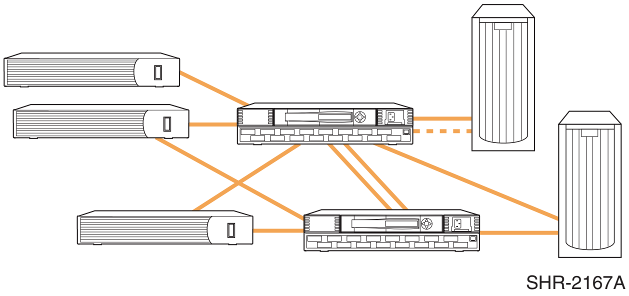

Figure 2–10: Level 3: Single Fabric High Availability Multi-Pathing

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 activity 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 vulnerabilities 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 level of protection is available for all Compaq standard SAN topologies by replicating the chosen design in two or more separate fabrics (Compaq highly recommends that the fabric design be symmetrical). Although this increases 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 fabric 4 switch-meshed design with multiple server and storage paths 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.

Utilizing two fabrics allows for non-disruptive software and firmware code updates. For example, given the two fabrics shown in Figure 2–11, you can fail over 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..

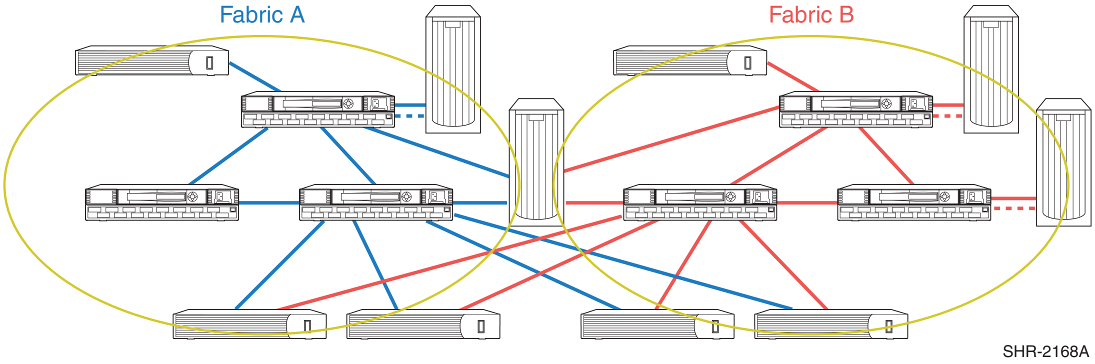

Figure 2–11: Level 4: Dual Fabric High Availability Multi-Pathing Fault Tolerant

Table 2–3 characterizes data availability and indicates the supported topologies for each level.

Table 2–3: 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**[1]	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

For more information refer to the section "Cabling Scheme Options" in Chapter 4.

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. Table 2–4 characterizes data availability levels relative to cost and total number of available ports.

Table 2–4: Availability Factors

Fabric Design	Level	Hardware Cost Factor[1]	Available Ports[2]
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** [3]	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 nonmeshed 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 Compaq standard SAN topologies can be scaled incrementally to increase connectivity and overall capacity. You should always plan for expected future growth when doing 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 9, "Best Practices" for information about migrating topologies.

Table 2–5 lists the migration paths and the options for scalability for all topologies.

Table 2–5: Topology Migration & Scaling

SAN Topology	Migration	Scalability (For All Topologies)
Cascaded	Convert to Meshed, Ring or Tree	• Increase the number of switches
Meshed	Convert to Ring, or Tree	• Use higher port count switches
Ring	Convert to Meshed or Tree	• Transition to a different topology
Tree	Add additional backbone switches	• Deploy multiple fabrics

3

SAN Fabric Design Rules

Heterogeneous SAN Fabric Design Configuration Rules

The sections in this chapter contain SAN Fabric design configuration rules for Heterogeneous SANs. The exact configuration rules for a SAN begin with the base SAN Fabric rules. These are modified depending on the specific topology implementation rules, platform or operating system and storage system rules described in Chapter 4, "Heterogeneous SAN Platform and Storage System Rules"**,** and the requirements of applications being run on the SAN. Read the documentation and release notes for all hardware and software products that are being utilized on the SAN for additional configuration information details. See the Preface, "Related Documents" for a list of related documentation.

General SAN Fabric Rules

1. Up to 20 switches total in a fabric (SAN fabric using Compaq SAN Switch 8, SAN Switch 16, SAN Switch 8-EL, SAN Switch 16-EL, and SAN Switch Integrated 32/64 model switches intermixed (refer to the SAN switch documentation). Each SAN Switch Integrated 32/64 model switch adds 6 switches to the fabric switch count.
2. Up to 4 switches total in a SAN fabric using Compaq Fibre Channel Switch 8 or Switch 16 model switches only or when intermixed with StorageWorks SAN Switch 8, SAN Switch 16, SAN Switch 8-EL, or SAN Switch 16-EL model switches.

   Intermixing of Compaq Fibre Channel switches and Compaq SAN switches requires compatibility mode be set in the SAN switches (refer to the SAN switch documentation).
3. The Compaq FC-AL Switch 8 is supported for cascaded attachment to the SAN through a single FL-port on a Compaq SAN Switch 8, SAN Switch 16, SAN Switch 8-EL, or SAN Switch 16-EL. In this configuration, RA4000/4100 storage systems are accessible only from servers attached directly to the FC-AL switch.
4. Up to 7 switch hops (8 switches) maximum between any two devices in the SAN. Each SAN Switch Integrated 32/64 model switch adds up to 2 hops between devices depending on the specific device-to-switch connections and device-to-device access (refer to Chapter 2, "Tree Backbone Fabrics", Figure 2–6 and Figure 2–7).
5. Within a single fabric where switches are interconnected, each switch must have a unique domain number (Domain ID) and a unique World Wide Name (WWN). All switch configuration parameters in each switch must be the same.
6. Up to 16 inter-switch links (ISLs) on a switch, with up to 8 active ISLs to the same destination.
7. Any mix of servers and storage systems is allowed in a SAN provided the specific platform, operating system, and storage system rules are followed. Refer to the appropriate sections in this guide and the documentation listed in the section "Related Documents" in the preface.

8. Compaq recommends all switches in a fabric use the same switch firmware revision. Two successive Fabric firmware versions can be temporarily used in one Fabric during rolling upgrades.

Specific Fabric Topology Rules

1. Up to 20 switches in a Cascaded, Meshed, or Backbone SAN Fabric topology.
2. Up to 14 switches configured in a ring with a Ring SAN Fabric topology. Up to 20 switches in a Ring SAN Fabric provided that no more than 10 switches are in a ring and no more then 10 switches are outside of the ring, one switch cascaded from each of the 10 ring switches.
3. Up to 7 switch hops between any 2 devices in a Cascaded, Meshed, or Ring SAN Fabric topology.

SAN Fabric Zoning Rules

The fabric zoning feature is supported with all Compaq Fibre Channel switch models. Zoning can be used to logically separate devices and different hardware platforms and operating systems in the same physical SAN. Use of zoning is required under these specific conditions:

- When mixing different hardware platforms, operating systems or storage systems that are currently only supported in homogenous SANs, and it is unknown whether there are interaction problems. Refer to Table 4–6 and Table 4–3 for specific information about zoning in heterogeneous SANs.
- When there are known interaction problems between different hardware platforms or operating systems and specific storage system types.
- When the number of nodes or ports in the SAN exceeds a storage system connection support limit. There is a connection limit for storage systems using the HSG60/80 or HSV110 controllers. The version of ACS or VCS controller code determines the specific limit.

SANworks Management Appliance Rules and Recommendations

Whenever a management appliance is placed in a fabric with heterogeneous servers, a dedicated storage management zone must be created. This zone is specifically for the management appliance and the elements it is to monitor and manage.

For example, create a zone called SANAPP_#_ZONE that would contain the appliance host bus adapter WWID and the WWIDs of all the HSG or HSV controllers managed by this SWMA. Because fabric devices can be in multiple zones, this will have no effect on other zones containing the same HSG and HSV controller WWIDs.

Currently, the management appliance communicates with HSG or HSV controllers in-band, that is, within the Fibre Channel fabric itself. It is not necessary or recommended to include either the switch WWIDs or server HBA WWIDs in this zone. Management communication to these devices from the SANworks Management Appliance is done out-of-band or outside the fabric via TCP/IP.

1. It is recommended that an SWMA be used to manage the SAN when a fabric contains more than four Fibre Channel switches.

2. When using an SWMA, one SWMA is required for each fabric in multiple fabric SANs. In DRM and other multi-fabric environments, it is recommended to have one SWMA per fabric per site, where only one appliance is active at any given time.
3. Multiple appliances per fabric are allowed as long as only one appliance is accessing a controller pair at a time. Zoning is required to actively isolate an appliance and its controllers from other appliances and their controllers.
4. Element Manager for HSG may manage up to 25 HSG controller pairs per appliance. The Element Manager for HSV may manage up to 16 HSV controller pairs per appliance.
5. An SWMA is required to manage Enterprise Virtual Arrays.
6. A license key is required for each Enterprise Virtual Array managed by the SANworks Management Appliance. Additional licenses are required to take advantage of value added software functionality.
7. For SANs with more than 1024 HBAs, an HSV controller must be zoned so that it can see no more than 1024 HBAs. It may be necessary to add a zone to a SAN to satisfy the 1024 HBA limit.

Refer to Chapter 4, "Heterogeneous SAN Platform and Storage System Rules", for rules about mixing specific platforms in a Heterogeneous SAN without the need for fabric zoning.

SAN Component Interconnect Descriptions and Rules

The following sections describe rules for SAN component interconnects–switch port interfaces and physical cabling.

Fibre Channel Switch Interface Usage Descriptions

- E-Port interface for switch to switch connectivity
- F-Port interface for fabric attached device–initiators (host bus adapters) and targets (storage ports)
- FL-Port interface for public loop fabric-aware, 24-bit Fibre Channel addressable devices–initiators or Compaq FC-AL Switch 8 SAN attachment
- FL-Port interface for private loop 8-bit Fibre Channel addressable devices–private FC-AL initiators and targets. This requires use of the Compaq SAN Switch 8 and SAN Switch 16 models' QuickLoop Feature. Typically this feature is only required when a specific platform can only be configured with a private FC-AL host bus adapter driver.
- G-Port, default interface when nothing appears to be attached to the port.

Access with QuickLoop

The QuickLoop switch feature allows private FC-AL initiators and targets configured in a QuickLoop to communicate with each other through the switch. Since all initiators configured *inside* a QuickLoop are private they cannot communicate with targets outside of the QuickLoop. QuickLoop is only supported on Compaq SAN Switch 8 and SAN Switch 16 models when used with MA6000, MA8000, RA8000, EMA12000/16000, and ESA12000 RAID Array systems.

Fiber Optic Interconnects/Distance Rules

1. The minimum[1] allowable bend radius of fiber optic cable specified by cable size is 25 mm for 50, 62.5, and 9 micron fiber optic cable.

2. There is a minimum fiber optic cable segment length between Fibre Channel devices (a transmitter and a receiver). This length is 0.5 meters for 50 and 62.5 micron cable and 2.0 meters for 9 micron cable. The minimum length does not apply to patch cords through a passive patch panel; it only applies to the total distance between the transmitter and receiver of the devices being connected through the patch panel. This distance should be greater than the minimum length for the type of cable being used.

3. Up to 200 meters maximum distance per cable segment between devices and switches or switches and switches using 62.5/125 micron, multi-mode fiber optic cable and short wavelength GBICs or GLMs[2].

NOTE: Information on the use of 62.5 micron fiber optic cable is provided to facilitate use of previously installed cable. Compaq recommends 50 micron fiber optic cable for any new installation requiring multi-mode fiber.

4. Up to 500 meters maximum distance per cable segment between devices and switches or switches and switches using 50/125 micron multi-mode fiber optic cable and short wavelength GBICs or GLMs.

5. Up to 10 km (6.3 miles) maximum distance between any two switches using 9/125 micron single-mode fiber optic cable and long wavelength GBICs.

6. Maximum of 70 km (44 miles) distance between devices when configured through seven 10 km switch-to-switch segments. This rule is based on a maximum of 10 km between any two switches and a maximum of seven hops between devices.

7. From greater than 10 km up to 100 km (63 miles) distance between two switches using 9/125 micron single-mode fiber optic cable and very long distance GBICs. A maximum of one 100 km very long distance segment per SAN.

NOTE: Refer to the Data Replication Manager (DRM) solution documentation for specific interconnect and distance rules related to DRM configurations.

8. Maximum of 160 km total distance across the SAN using multiple segments. This can be implemented using a single 100 km segment and six 10 km segments (7 hops) or other combinations such as two 50 km segments and five 10 km segments (7 hops). In all cases the individual segments must be configured with the proper GBICs and cable type and the maximum of 7 hops must not be exceeded across the SAN.

9. For longer 50-micron short wave multi-mode optical fiber cables (up to 500 meters), a third party vendor must be contacted. The cables must be duplex, tight buffered multi-mode 50/125 µm (Belcore GR-409 compliant) and the connectors must be SC duplex low metal (NTT-SC Belcore 326, IEC-874-19 SC compliant).

[1] Specification is based on representative cables. Consult manufacturer for parameters of specific cables.

[2] GLMs are used in the HSG60 (MA6000) and HSG80 (MA/RA8000, EMA/ESA12000, EMA16000) storage controllers

10. For 9-micron long wave single-mode optical fiber cables (up to 100 kilometers), a third party vendor must be contacted. The cables must be duplex, tight buffered, single-mode 9/125 µm (Belcore GR-409 compliant) and the connectors must be SC duplex low metal (NTT-SC Belcore 326, IEC-874-19 SC compliant).
11. The mixing of 9-micron, 50-micron, and 62.5-micron fiber cables in the same cable segment is not supported.

Fiber Optic Cable Loss Budgets

The information in this section is based on the Fibre Channel Physical Interface Specification. Refer to the specification document for more information.

NOTE: Media losses are not specified due to variances between different fiber optical cable manufacturers. In all cases the specification that must be followed is the total channel insertion loss, which includes media losses.

1. For 62.5/125 micron fiber optic cable up to 200 meters:
   Maximum of 3 dB total channel insertion loss[1], 0.75 dB loss per mated connector pair.

NOTE: Information on the use of 62.5 micron fiber optic cable is provided to facilitate use of previously installed cable. Compaq recommends 50 micron fiber optic cable for any new installation requiring multi-mode fiber.

2. For 50/125 micron fiber optic cable up to 500 meters:
   Maximum of 3.85 dB total channel insertion loss, 0.75 dB loss per mated connector pair.
3. For 9/125 micron fiber optic cable up to 10 km:
   Maximum of 7.8 dB total channel insertion loss, 0.75 dB loss per mated connector pair.
4. For 9/125 micron fiber optic cable up to 100 km:
   Maximum of 21.5 dB total channel insertion loss, 0.75 dB loss per mated connector pair.
5. Use of optical fiber patch panels is supported provided the total channel insertion loss between the transmitter and receiver for the cable segment routed through the patch panel does not exceed the maximum listed for the connector and cable type in use.

[1] Channel insertion loss is the combined passive loss from connectors, splices, and media between the transmitter and receiver.

Table 3–1: Storage Product Interconnect/Transport Support

	Storage Product		
Interface/Transport	**Heterogeneous SAN**	**DRM**	**Enterprise Backup Solutions (EBS)**
Fibre Channel via 50 micron multi-mode fiber optic cable and short-wave GBICs	Up to 500 meters per cable segment	Up to 500 meters per cable segment	Up to 500 meters per cable segment
Fibre Channel via 62.5 micron multi-mode fiber optic cable and short-wave GBICs	Up to 200 meters per cable segment	Up to 200 meters per cable segment	Up to 200 meters per cable segment
Fibre Channel via 9 micron single-mode fiber optic cable and long-wave GBICs	Up to 10 km per cable segment	Up to 10 km per cable segment	Up to 10 km per cable segment
Fibre Channel via 9 micron single-mode fiber optic cable and very long distance GBICs	Up to 100 km per cable segment and 160 km total distance	Up to 100 km per cable segment and 160 km total distance	Up to 10 km per cable segment
Fibre Channel via Wave Division Multiplexing (WDM) and Dense Wave Division Multiplexing (DWDM)	Not Supported	Up to 100 km per segment and 160 km total distance	Not Supported
ATM over single T1/E1 Wide Area Network (WAN)	Not Supported	Supported No Distance Limit	Not Supported
ATM over single T1/E1 WAN (Inverse Multiplexing)	Not Supported	Supported No Distance Limit	Not Supported
ATM over T3/E3 WAN	Not Supported	Supported No Distance Limit	Not Supported
ATM over fractional and/or shared T3/E3 and OC3 WAN	Not Supported	Supported No Distance Limit	Not Supported

General Fabric Performance Recommendations

The performance of an application on a heterogeneous SAN is usually seen from the perspective of "storage performance". The intervening SAN and competing workloads are not usually considered. The fact is that the performance of the storage will be dependent on the interaction of all the components and applications in the SAN. Some of the possible component limiting factors include the host CPU(s), FC HBA, SAN topology, SAN traffic, RAID controllers, or the specific configuration of disks used behind the controllers. This is a dynamic workload environment and at any given moment any part of the SAN can dominate the performance. This complexity can be simplified if the environment is divided into the categories of servers, SAN infrastructure, and data storage devices. This document is primarily interested in the SAN infrastructure. There are issues with storage and servers that are unique to SAN implementations and we will address some of those as well.

SAN Infrastructure Performance

For the purposes of discussion we will assume a multi-switch fabric, since single switches always offer the highest performance with minimum latency. A multi-switch fabric has two factors that decrease overall fabric-wide infrastructure performance:

- latency through multiple switches (hops)
- oversubscription or congestion of ISLs

Performance testing and measurement by Compaq has shown switch latency to be less than 5% of the time lost due to congestion of a full frame from another path. This implies that the number of switches and hops between devices is not a major factor for performance. However, as devices send frames through more switches and hops, the chances are increased that other traffic in the SAN may contend for the same ISL or path. This may result in lower performance due to oversubscription of a particular ISL or path that is serving multiple devices.

Oversubscription has been determined to be the largest contributing factor to reduced Fibre Channel performance. When devices must contend for the same ISL or path, the result is that each competing device will receive 1/n of the available bandwidth on the path (where n is the number of contending devices).

While the topology and size of the SAN have been seen to affect performance, staying within the rules and recommendations outlined in this guide minimizes these factors. The topology designs have been defined to accommodate a particular data access or data locality type. Recommendations on the number of ISLs based on device-to-device access ratios serve to ensure that adequate bandwidth is available across the SAN, minimizing oversubscription.

Compaq recommends following these guideline in configuring your SAN.

- Whenever possible, devices that exchange the highest amount of data should be connected to the same Fibre Channel switch.
- For high bandwidth, the number of application servers should be balanced with storage by using as much one-to-one access as possible
- When devices exchanging data are on different switches:
  - Minimize the number of hops between devices
  - For high bandwidth (large transfer size) applications, configure a maximum of two active storage controller ports per ISL
  - For high throughput (small transfer size) applications, configure a maximum of 20 active storage controller ports per ISL
  - For mixed applications, configure a maximum of four active storage controller ports per ISL

To plan the SAN you need to know the capabilities of all the devices connected. The maximum performance characteristics of the some of the StorageWorks devices are listed in Table 3–2. User applications may not necessarily reach these levels of performance as applications may perform additional levels of processing before each I/O. The controller specifications listed show both cache (no disk access) and media (with disk access) limitations.

NOTE: Unlike parallel SCSI, Fibre Channel SANs provide full duplex operations (simultaneous read and write traffic). The performance numbers listed reflect half-duplex traffic flow. In some cases multiple application exchanges may well be able to see nearly twice the MBs per second figures stated herein.

Table 3–2: Storage Component Performance Specifications

Components	IO/sec (small transfer sizes) (maximum)	MB/sec (large transfer sizes) (maximum)
HP A3740-60101 (3.0 Tachyon)	22,000	97
SUN 64 Bit S-Bus JNI HBA SW-SWSA4-SC	9,452	80
W2K/WNT4.0 DS-KGPSA-CB	28,100	102
FC Switch (see note) FC ISL	1,800,000 110,000	1,600 100
One HSV110 of a controller pair (2 active ports)	80,000 Cache Access 36,000 Media Access	200 Cache Access 200 Media Access
Two HSV110 controllers (4 active ports)	161,000 Cache Access 55,000 Media Access	400 Cache Access 400 Media Access
Single HSG80 controller (1 active, 1 standby port)	10,700 Cache Access 5,200 Media Access	77 Cache Access 54 Media Access
Single HSG80 controller (2 active ports)	10,600 Cache Access 5,200 Media Access	98 Cache Access 54 Media Access
Dual HSG80 controller (2 active, 2 standby ports)	21,100 Cache Access 10,500 Media Access	154 Cache Access 102 Media Access
Dual HSG80 controller (4 active ports)	28,800 Cache Access 10,800 Media Access	195 Cache Access 102 Media Access
Single DLT Tape drive	N/A	5 Media Access
Single SDLT Tape drive	N/A	11 Media Access

All numbers are for planning purposes and represent each device’s best-known performance.

The limits are based on I/O performance (I/Os per second) - typical of small transfer applications such as databases and mail, and bandwidth performance (MBs per second) - typical of large transfer applications such as video and graphics.

The HSG80 performance numbers are based on V8.5F of the ACS. Users should expect differences when using the S or P variants of the ACS firmware.

NOTE: Fibre channel switch performance limits are theoretical.

4

Heterogeneous SAN Platform and Storage System Rules

This chapter describes rules related to specific platforms, operating systems, and storage products. For additional information refer to the relevant platform specific application notes and individual product documentation. Refer to the preface for a list of related documentation.

General Platform/Operating System Rules

1. Each platform listed is supported in all SAN Fabric topology configurations unless otherwise noted in this guide or the applicable platform application notes.
2. Any mix of servers (clustered and standalone) and storage systems is allowed in a SAN, provided that you follow all individual platform rules, storage system rules, and maximums listed in this guide and in platform specific documentation.
3. Refer to the section "Combined Shared Access Interoperability Table" for information related to mixing platforms on a single shared storage system. In certain situations multiple storage systems may be required to accommodate the requirements of different platforms or operating systems.
4. All Compaq and multi-vendor hardware platforms and operating systems that are supported in a homogeneous SAN are supported in a heterogeneous SAN. Refer to Table 4–3 to determine if zoning is required for specific combinations of supported heterogeneous platforms.
5. Servers and storage systems can attach to multiple fabrics. The number of separate fabrics per server is based on the specific server model capabilities and the maximum number of Fibre Channel host bus adapters supported.
6. Refer to the section "High Availability Configuration Considerations" for cabling scheme options for platforms that support high availability multi-pathing.

Specific Platform/Operating System Rules – Enterprise Virtual Array, MA6000, MA/RA8000, EMA/ESA12000, and EMA16000 Storage Systems

Table 4–1 summarizes the rules for configuring each hardware platform or operating system with the supported storage system types. The table lists the platforms, HBAs, storage system types, and the preferred method of storage attachment for each platform type. The sections following the table describe additional rules not listed in the table. Refer to Table 4–6 for information about configuring storage systems for shared access to multiple heterogeneous platforms in a SAN.

NOTE: This document specifies information for VCS 1.02 and ACS 8.6.

For the most current information on HBAs, firmware, and drivers see this document:

Supplemental Tables for the Heterogeneous Open SAN Design Reference Guide

http://www.compaq.com/products/storageworks/san/documentation.html

Table 4–1: Platform/Storage System SAN Attachment Summary – Enterprise Virtual Array, MA6000, MA/RA8000, EMA/ESA12000, and EMA16000 Storage Systems

Platform or Operating System VCS/ACS versions in ( )	Platform HBA SAN Attachment	Multi-Path Support	Enterprise VirtualArray SAN Attachment	MA6000, MA8000, RA8000, EMA12000/16000, ESA12000 Storage System SAN Attachment
Compaq OpenVMS 7.2-1H1 w/TIMA v300 (VCS 1.02) 7.2-2 TIMA DEC-AXPVMS-VMS722_FIBRE_SCSI-V0100--4 (VCS 1.02, ACS 8.6) 7.3 TIMA DEC-AXPVMS-VMS73_FIBRE_SCSI-V0200--4 (VCS 1.02, ACS 8.6) Clusters	F-Port KGPSA-BC 380574-001 KGPSA-CA 168794-B21	Native	F-Port Operating System Parameter: OPENVMS Unit Identifier Required	F-Port using FABRIC topology Multiple-Bus Failover SCSI-3 Mode Connection name O/S parameter: VMS Unit Identifier required
	KGPSA-DA 261329-B21 (ACS 8.6)		N/A	
Tru64 UNIX 4.0F (ACS 8.6) TruCluster Software Products Version 1.6 (ACS 8.6): 4.0F Patch Kit 7	F-Port KGPSA-BC 380574-001 KGPSA-CA 168794-B21	N/A	N/A	F-Port using FABRIC topology Transparent Failover SCSI-2 Mode Command Console LUN enabled/disabled or SCSI-3 Mode Connection name O/S parameter: TRU64_UNIX
Tru64 UNIX 5.1 (ACS 8.6) 5.1 Patch Kit 4, 5.1A Patch Kit 1 (VCS 1.02, ACS 8.6) TruCluster Server Version 5.1, 5.1A (VCS 1.02, ACS 8.6)	F-Port KGPSA-BC 380574-001 KGPSA-CA 168794-B21	Native	F-Port Operating System Parameter: TRU64	F-Port using FABRIC topology Transparent or Multiple-Bus Failover SCSI-2 Mode Command Console LUN enabled/disabled or SCSI-3 Mode Connection name O/S parameter: TRU64_UNIX
	KGPSA-DA 261329-B21 (ACS 8.6)		N/A	

Table 4–1: Platform/Storage System SAN Attachment Summary – Enterprise Virtual Array, MA6000, MA/RA8000, EMA/ESA12000, and EMA16000 Storage Systems

Platform or Operating System VCS/ACS versions in ( )	Platform HBA SAN Attachment	Multi-Path Support	Enterprise VirtualArray SAN Attachment	MA6000, MA8000, RA8000, EMA12000/16000, ESA12000 Storage System SAN Attachment
HP-UX 10.20, 11.0, 11i (ACS 8.6) MC/ServiceGuard Clusters	L-Port (QuickLoop) **10.20, 11.0, 11i:** HPA3404A A3591A A3636A A3740A A5158A A6685A **11.0/11i:** F-Port HPA5158A A6685A 218409-B21	HP PV Links (10.20, 11.0) Secure Path (11.0) HPA5158A, A6685A	N/A	L-Port (QuickLoop) using LOOP_HARD topology (10.20 11.0, 11i) F-Port using FABRIC topology (11.0, 11i) Transparent or Multiple-Bus Failover (10.20, 11.0) SCSI-2 Mode Command Console LUN enabled/disabled or SCSI-3 Mode Connection name O/S parameter: HP
IBM AIX 4.3.3, 5.1 (ACS 8.6)	F-Port 197819-B21 SWIA1-PD	Secure Path		F-Port using FABRIC topology Transparent Failover or Multiple-Bus Failover SCSI-2 Mode Command Console LUN enabled/disabled SCSI-3 Mode Connection name O/S parameter: WINNT
Linux Caldera 2.3.1 (ACS 8.6) Intel	F-Port 167433-B21	N/A	N/A	F-Port using FABRIC topology Transparent Failover SCSI-3 Mode Connection name O/S parameter:SUN
Linux Redhat 7.0, 7.1 (ACS 8.6) Alpha/Intel	F-Port 167433-B21	N/A	N/A	F-Port using FABRIC topology Transparent Failover SCSI-3 Mode Connection name O/S parameter:SUN
Linux SuSE 6.3, 7.0 (ACS 8.6) Alpha/Intel	F-Port 167433-B21	N/A	N/A	F-Port using FABRIC topology Transparent Failover SCSI-3 Mode Connection name O/S parameter:SUN
Microsoft Windows 2000 SP2 (VCS 1.02, ACS 8.6) Windows NT 4.0 SP6a (VCS1.02, ACS 8.6)	F-Port KGPSA-BC 380574-001 KGPSA-CB	Secure Path	F-Port Operating System Parameter: WINDOWS	F-Port using FABRIC topology Transparent or Multiple-Bus Failover SCSI-2 Mode Command Console LUN disabled or SCSI-3 Mode Connection name O/S parameter: WINNT

Table 4–1: Platform/Storage System SAN Attachment Summary – Enterprise Virtual Array, MA6000, MA/RA8000, EMA/ESA12000, and EMA16000 Storage Systems

Platform or Operating System VCS/ACS versions in ( )	Platform HBA SAN Attachment	Multi-Path Support	Enterprise VirtualArray SAN Attachment	MA6000, MA8000, RA8000, EMA12000/16000, ESA12000 Storage System SAN Attachment
Microsoft Windows 2000 Datacenter 1.1a (ACS 8.6) MSCS	F-Port KGPSA-BC 380574-001 KGPSA-CB	Secure Path	N/A	F-Port using FABRIC topology Transparent or Multiple-Bus Failover SCSI-2 Mode Command Console LUN disabled or SCSI-3 Mode Connection name O/S parameter: WINNT
Novell NetWare 4.2, 5.1, 6.0 (ACS 8.6) 5.1 with Clusters 1.01 (ACS 8.6) 6.0 with Clusters 1.06 (ACS 8.6)	FL-Port 120186-B21 223180-B21	Secure Path (5.1, 6.0)	N/A	F-Port using FABRIC topology Transparent (4.2, 5.1, 6.0) or Multiple-Bus Failover (5.1, 6.0) SCSI-2 Mode Command Console LUN enabled/disabled or SCSI-3 Mode Connection name O/S parameter: NETWARE
SGI IRIX 6.5.11, 6.5.12 (ACS 8.6)	F-Port SGI PCI-FC-1POPT SGI XT-FC-1POPT	N/A	N/A	F-Port using FABRIC topology Transparent Failover SCSI-2 Mode Command Console LUN disabled or SCSI-3 Mode Connection name O/S parameter:SGI
Sun Solaris 2.6 (32-bit) (ACS 8.6) Sun Solaris 7, 8 (32/64-bit) (VCS 1.02, ACS 8.6) Veritas Clusters 1.3 (VCS 1.02) Sun Solaris 8 (64-bit cPCI) (ACS 8.6) SUN Clusters v2.2 w/ 2.6, 7 (ACS 8.6) Veritas Clusters 1.2.1 (ACS 8.6)	F-Port 380575-001 (32-bit Sbus) 123503-001 (64-bit Sbus) 380576-001 (32-bit PCI)	Secure Path (Sbus, PCI)	F-Port Operating System Parameter: SOLARIS	F-Port using FABRIC topology Transparent or Multiple-Bus Failover SCSI-2 Mode Command Console LUN enabled/disabled or SCSI-3 Mode Connection name O/S parameter:SUN
	254457-B21 (64-bit cPCI) (ACS 8.6)	N/A	N/A	N/A

Compaq OpenVMS

- Supports Multiple-Bus failover mode. Multi-path driver is embedded in the operating system.
- Supports multi-path high availability configuration implemented in separate fabrics or a single fabric.
- Zoning required when used in a Heterogeneous SAN with HP-UX or IBM AIX.

7.2-1H1, 7.2-2, 7.3 – VCS 1.02

Enterprise Virtual Array:

- A Unit Identifier is required for each Virtual Disk
- Version 7.2-1H1 requires TIMA v300 Kit
- Version 7.2-2 requires TIMA DEC-AXPVMS-VMS722_FIBRE_SCSI-V0100—4 Kit
- Version 7.3 requires TIMA DEC-AXPVMS-VMS73_FIBRE_SCSI-V0200—4 Kit

7.2-2, 7.3 – ACS 8.6

MA6000, MA/RA8000, EMA/ESA12000, and EMA16000 storage systems:

- A Unit Identifier is required for each controller.
- Version 7.2-2 requires TIMA DEC-AXPVMS-VMS722_FIBRE_SCSI-V0100—4 Kit
- Version 7.3 requires TIMA DEC-AXPVMS-VMS73_FIBRE_SCSI-V0200—4 Kit
- Servers and storage systems configured for DRM must be zoned. One DRM zone per fabric for each DRM configuration.

Tru64 UNIX

- Tru64 UNIX versions V5.1/V5.1A support TruCluster Server Version 5.1/5.1A
- Tru64 UNIX versions 4.0F support TruCluster Software Products Version 1.6
- Supports multi-path high availability configuration implemented in separate fabrics or a single fabric
- Zoning is required when used in a Heterogeneous SAN with HP-UX, IBM AIX or Linux

5.1, 5.1A – VCS 1.02

Enterprise Virtual Array:

- Version 5.1 requires Patch Kit 4, version 5.1A requires Patch Kit 1
- Supports Multiple-Bus Failover mode. Multi-path driver is embedded in the V5.1/V5.1A operating systems.

4.0F, 5.1, 5.1A – ACS 8.6

- MA6000, MA/RA8000, EMA/ESA12000, and EMA16000 storage systems:
- Version 4.0F requires Patch Kit 7
- Version 5.1A requires Patch Kit 1
- Tru64 UNIX version 4.0F supports Transparent failover mode.
- Tru64 UNIX version 5.1and 5.1A supports Transparent and Multiple-Bus failover mode. Multi-path driver is embedded in the V5.1/V5.1A operating systems.
- Zoning is required when a SAN is configured for multiple TruCluster products with Tru64 UNIX 4.0F. Each TruCluster configured with Tru64 UNIX 4.0F must be in its own zone.
- Servers and storage systems configured for DRM must be zoned. One DRM zone per fabric for each DRM configuration.

HP-UX

10.20, 11.0, 11i – ACS 8.6

MA6000, MA/RA8000, EMA/ESA12000, and EMA16000 storage systems:

- Version 11.0 and 11i use the Compaq FC-Fabric host bus adapter or HP native HBA for Fabric attachment. Version 10.20, 11.0, and 11i can use the HP FC-AL host bus adapter configured for LOOP topology in a QuickLoop for FC-AL attachment.
- Supports Transparent failover mode and Multiple-Bus failover mode. Multiple-Bus failover mode is supported for HP-UX version 10.20 and 11.0 (FC-AL) using the HP PV Links multi-path driver and supported for HP-UX version 11.0 using the SANworks Secure Path multi-path driver. Supports multi-path high availability configuration implemented in separate fabrics or in a single fabric with zoned paths.
- Zoning is required when used in a Heterogeneous SAN with other operating systems.
- Supports MC/ServiceGuard Version A.11.13

IBM AIX

4.3.3, 5.1 – ACS 8.6

MA6000, MA/RA8000, EMA/ESA12000, and EMA16000 storage systems:

- Supports Transparent failover mode and Multiple-Bus failover mode. SANworks Secure Path multi-path driver is required for Multiple-Bus failover.
- Supports HACMP/ES Clusters 4.4.1 ES (5.1 only).
- If you configure greater than 4 (up to 6) servers (assuming one Fibre Channel HBA per server) for access to a single controller host port on an MA6000, MA/RA8000, EMA/ESA12000, and EMA 16000 storage system, and 1 or more of those servers is IBM AIX based, set the disk queue depth to 20 for all LUNs.
- Zoning is required when used in a Heterogeneous SAN with other operating systems.

Linux

Caldera (Intel) 2.3.1, Redhat (Alpha/Intel) 7.0, 7.1
SuSE (Alpha/Intel) 6.3, 7.0 – ACS 8.6

MA6000, MA/RA8000, EMA/ESA12000, and EMA16000 storage systems:

- Supports Transparent failover mode.
- Zoning is required when used in a Heterogeneous SAN with other operating systems.

Microsoft Windows

- Supports multi-path high availability configuration implemented in separate fabrics or in a single fabric with zoned paths.
- Zoning is required when used in a heterogeneous SAN with Linux, IBM AIX or HP-UX.
- Supports MSCS.

Windows 2000 w/SP2, Windows NT 4.0 w/SP6a – VCS 1.02, ACS 8.6

Enterprise Virtual Array:

- Supports Multiple-Bus Failover mode. SANworks Secure Path multi-path driver is required for Multiple-Bus failover.

MA6000, MA/RA8000, EMA/ESA12000, and EMA16000 storage systems:

- Supports Transparent failover mode and Multiple-Bus failover mode. SANworks Secure Path multi-path driver is required for Multiple-Bus failover.
- Servers and storage systems configured for DRM must be zoned. One DRM zone per fabric for each DRM configuration.
- Supports booting over the SAN Fabric. Refer to the section “Booting from the SAN”
- Extended Configurations with Microsoft Windows NT 4.0/Windows 2000
- If you configure greater than 4 (up to 8) servers (assuming one Fibre Channel HBA per server) for access to a single controller host port on an MA6000, MA/RA8000, EMA/ESA12000, and EMA16000 storage system, and 1 or more of those servers is Windows based, select the “Extended Configuration” check box in the StorageWorks Windows NT/Windows 2000 Platform Kit Fibre Channel Software Setup utility custom installation setup for each Windows server. Select this option to adjust registry settings for your KGPSA host bus adapter to operate in an "Extended Configuration" environment.

NOTE: The default for this option is checked, so be sure to uncheck this option when you have 4 or fewer servers configured for access to a single controller host port.

Windows 2000 Datacenter – ACS 8.6

MA6000, MA/RA8000, EMA/ESA12000, and EMA16000 storage systems:

- Supports Transparent failover mode and Multiple-Bus failover mode. SANworks Secure Path multi-path driver is required for Multiple-Bus failover.
- Not supported for heterogeneous shared access on a single storage system. The storage system can only be accessed by Datacenter servers.

SANworks Secure Path for Windows

For MA6000, MA/RA8000, EMA/ESA12000, and EMA16000 storage systems:

- Supports using single host bus adapter with Secure Path and Multiple-Bus failover. Refer to the white paper, *Using Secure Path for Servers with Single Host Bus Adapter (HBA)*, 14JK-0301A-WWEN.
- One instance of the Secure Path Manager can support multiple managed entities called profiles. For 3.x versions of Secure Path, a single profile can consist of up to 8 standalone servers connected to and sharing up to 8 storage systems, or up to 8 clustered servers connected to and sharing up to 8 storage systems. You cannot manage both standalone and clustered servers in the same profile.
- Secure Path configurations utilizing 4 active controller ports connected to the same server or servers offer the flexibility to use the 4 active ports for either increased total LUN count, or increased PATH accessibility to a lesser number of LUNs. Refer to the section “High Availability Configuration Considerations” for more information.
- Provides for dynamic port I/O load distribution in non-clustered servers when configured for maximum paths.

- Distribute units equally across both controllers for proper static load balancing using the Unit Preferred Path parameter to assign units to a specific controller at initial boot.
- SSP/LUN level masking - Storagesets (LUNs) must be enabled for access from both paths using the storage LUN presentation or Unit Connection Name parameter feature.
- For Windows NT or Windows 2000, when using Compaq Secure Path in single or dual fabric configurations with both Multiple-bus Failover and Transparent Failover storage systems, the Transparent Failover storage systems must be in a different fabric zone.

Novell NetWare

- Zoning required when used in a Heterogeneous SAN with Sun, HP-UX or IBM AIX

4.2 – ACS 8.6

MA6000, MA/RA8000, EMA/ESA12000, and EMA16000 storage systems:

- Supports Transparent failover mode.

5.1, 6.0 – ACS 8.6

MA6000, MA/RA8000, EMA/ESA12000, and EMA16000 storage systems:

- Supports NetWare Clusters.
- Supports Transparent failover mode and Multiple-Bus failover mode. SANworks Secure Path multi-path driver is required for Multiple-Bus failover.
- Supports multi-path high availability configuration implemented in separate fabrics or in a single fabric with zoned paths.

SGI IRIX

- Zoning required when used in a Heterogeneous SAN with HP-UX or IBM AIX

6.5.11, 6.5.12 – ACS 8.6

MA6000, MA/RA8000, EMA/ESA12000, and EMA16000 storage systems:

- Supports Transparent failover mode.

Sun Solaris

- Zoning required when used in a Heterogeneous SAN with Netware, HP-UX or IBM AIX
- Supports multi-path high availability configuration implemented in separate fabrics or in a single fabric with zoned paths.

7, 8 – VCS 1.02

Enterprise Virtual Array:

- Supports Multiple-Bus Failover mode. SANworks Secure Path multi-path driver is required for Multiple-Bus failover.
- Supports VERITAS Clusters 1.3. A cluster must be in its own zone.

2.6 and 7, 8 (32-bit or 64-bit) – ACS 8.6

MA6000, MA/RA8000, EMA/ESA12000, and EMA16000 storage systems:

- Supports Transparent failover mode and Multiple-Bus failover mode. SANworks Secure Path multi-path driver is required for Multiple-Bus failover.
- Supports Sun Clusters version 2.2 (Sun 2.6, and 7) and VERITAS Clusters 1.2.1. A cluster must be in its own zone.
- Supports cPCI on Sun Solaris 8
- Servers and storage systems configured for DRM must be zoned. One DRM zone per fabric for each DRM configuration.

Heterogeneous SAN Platform Interoperability for Enterprise Virtual Array

For the Enterprise Virtual Array, heterogeneous SAN platform interoperability is defined in Table 4–2. A “Yes” in the table indicates that the listed platforms can be configured for shared access to the same Enterprise Virtual Array. “Zoning Required” indicates the platforms listed must be configured in different fabric zones in order to co-exist in the same physical SAN or share the same Enterprise Virtual Array.

Table 4–2: SAN/Platform Interoperability for Single Shared Enterprise Virtual Array

Platform or Operating System	**Compaq OpenVMS** 7.2-1H1, 7.3	**Tru64 UNIX** 5.1, 5.1A	**Microsoft Windows NT** 4.0 SP6a **Windows 2000** SP2	**Sun Solaris** 7, 8
Compaq OpenVMS 7.2-1H1, 7.3	Yes	Yes	Yes	Yes
Tru64 UNIX 5.1, 5.1A	Yes	Yes	Yes	Yes
Microsoft Windows NT 4.0 SP6a **Windows 2000** SP2	Yes	Yes	Yes	Yes
Sun Solaris 7, 8	Yes	Yes	Yes	Yes

Heterogeneous SAN Platform Interoperability for MA6000, MA/RA8000, EMA/ESA12000, and EMA16000 Storage Systems

For MA6000, MA/RA8000, EMA/ESA12000, and EMA16000 storage systems there are three levels of heterogeneous interoperability rules: platform zoning rules, controller SCSI-modes, and controller failover modes.

The platform zoning rules define which platforms or operating systems must be in different fabric zones in order to coexist in the same physical SAN. Refer to Table 4–3, Compatible SCSI Modes for MA6000, MA/RA8000, EMA/ESA12000, and EMA16000 Storage Systems Using ACS 8.6.

The controller SCSI-mode and controller failover rules define which platforms or operating systems can be configured for shared access to a single shared storage system based on controller SCSI-mode and failover mode compatibility. Refer to Table 4–4 “Compatible SCSI Modes for MA6000, MA/RA8000, EMA/ESA12000, and EMA16000 Storage Systems Using ACS 8.6” and Table 4–5 “Compatible Failover Modes for MA6000, MA/RA8000, EMA/ESA12000, and EMA16000 Storage Systems Using ACS 8.6”.

An additional table, Table 4–6, combines (and to some extent repeats) the information from the other tables into single tables that can be quickly referenced to determine the settings and rules for mixing all possible combinations of any two platforms.

Platform Zoning Rules

This table summarizes the zone compatibility for different platforms in a SAN. Platforms in the same columns can coexist in the same zone.

Table 4–3: SAN/Platform Zoning Requirements - MA6000, MA/RA8000, EMA/ESA12000, and EMA16000 Storage Systems

Zone 1	Zone 2	Zone 3	Zone 4	Zone 5
SGI OpenVMS NetWare Tru64 UNIX	Sun OpenVMS Windows NT Windows 2000 Tru64 UNIX	SGI OpenVMS SUN Linux	HP-UX	IBM AIX

NOTE: The above table is summarized as:

- NetWare and Sun platforms are incompatible in the same zone.
- Linux is incompatible in the same zone with Tru64 UNIX or Microsoft Windows.
- HP-UX and IBM AIX platforms are incompatible in zones with all platforms.

Compatible Controller SCSI-Modes and Controller Failover Modes

The following tables summarize information about supported controller SCSI modes and failover modes for all platforms. Table 4–4 summarizes information about compatible SCSI modes and Table 4–5 summarizes information about supported storage system failover modes for all platforms for MA6000, MA/RA8000, EMA/ESA12000, and EMA16000 storage systems using ACS 8.6.

Table 4–4: Compatible SCSI Modes for MA6000, MA/RA8000, EMA/ESA12000, and EMA16000 Storage Systems Using ACS 8.6

SCSI-2 CCL	SCSI-2 No CCL	SCSI-3
		OpenVMS 7.2-2, 7.3
Tru64 4.0F, 5.1, 5.1A	Tru64 4.0F, 5.1, 5.1A	Tru64 4.0F, 5.1, 5.1A
HP-UX 10.20, 11.0, 11i	HP-UX 10.20, 11.0, 11i	HP-UX 10.20, 11.0, 11i
IBM AIX 4.3.3, 5.1	IBM AIX 4.3.3, 5.1	IBM AIX 4.3.3, 5.1
		Linux Caldera 2.3.1
		Linux Redhat 7.0, 7.1
		Linux SuSE 6.3, 7.0
	Microsoft Windows NT 4.0 Windows 2000	Microsoft Windows NT 4.0 Windows 2000
Novell NetWare 4.2, 5.1, 6.0	Novell NetWare 4.2, 5.1, 6.0	Novell NetWare 4.2, 5.1, 6.0
SGI 6.5.11, 6.5.12		SGI 6.5.11, 6.5.12
Sun 2.6, 7, 8	Sun 2.6, 7, 8	Sun 2.6, 7, 8

Table 4–5: Compatible Failover Modes for MA6000, MA/RA8000, EMA/ESA12000, and EMA16000 Storage Systems Using ACS 8.6

Transparent	Multiple-Bus
	OpenVMS 7.2-2, 7.3
Tru64 4.0F, 5.1, 5.1A	Tru64 5.1, 5.1A
HP-UX 10.20, 11.0, 11i	HP-UX 10.20, 11.0
IBM AIX 4.3.3, 5.1	IBM AIX 4.3.3, 5.1
Linux Caldera 2.3.1	
Linux Redhat 7.0, 7.1	
Linux SuSE 6.3, 7.0	
Microsoft Windows NT 4.0 Windows 2000	Microsoft Windows NT 4.0 Windows 2000
Novell NetWare 4.2, 5.1, 6.0	Novell NetWare 5.1, 6.0
SGI 6.5.11, 6.5.12	
Sun 2.6, 7, 8	Sun 2.6, 7, 8

Combined Shared Access Interoperability Table

Table 4–6 combines the information from the previous tables into a single table. The table can be used to determine controller settings for a single MA6000, MA/RA8000, EMA/ESA12000, or EMA16000 storage system using ACS 8.6, being shared between two or more platforms and operating systems.

Table 4–6: Platform Interoperability for Single Shared MA6000, MA/RA8000, EMA/ESA12000, or EMA16000 Storage System – ACS 8.6

Platform or Operating System	**Compaq OpenVMS** 7.2-2, 7.3 Clusters	**Tru64 UNIX** 4.0F Trucluster Software Products V1.6	**Tru64 UNIX** 5.1, 5.1A TruCluster Server Version 5.1/5.1A	**HP-UX** 10.20, 11.0, 11i MC/ ServiceGuard Clusters	**IBM AIX** 4.3.3, 5.1	**Linux Caldera** 2.3.1 **Redhat** 7.0, 7.1 **SuSE** 6.3, 7.0	**Microsoft Windows NT** 4.0 SP6a **Windows 2000** SP2 MSCS	**Novell NetWare** 4.2, 5.1, 6.0 Clusters 1.01, 1.06	**SGI IRIX** 6.5.11, 6.5.12	**Sun Solaris** 2.6 and 7, 8 (32/64-bit) Sun Clusters 2.2, VERITAS Clusters 1.2.1
Compaq OpenVMS 7.2-2, 7.3 Clusters	Multiple-Bus FABRIC SCSI-3	Requires two storage systems	Multiple-Bus FABRIC SCSI-3	11.0, 11i: Fabric attachment, With Zoning Multiple-BusFABRIC SCSI-3	With Zoning Multiple-Bus FABRIC SCSI-3	Requires two storage systems	Multiple-Bus FABRIC SCSI-3	5.1, 6.0: Multiple-Bus FABRIC SCSI-3 4.2: Requires two storage systems	Requires two storage systems	Multiple-Bus FABRIC SCSI-3
Tru64 UNIX 4.0F Trucluster Software Products V1.6	Requires two storage systems	Transparent FABRIC SCSI-2 or SCSI-3	Transparent FABRIC SCSI-2 or SCSI-3	11.0, 11i: Fabric attachment, With Zoning TransparentFABRIC SCSI-2 or SCSI-3	With Zoning Transparent FABRIC SCSI-2 or SCSI-3	With Zoning Transparent FABRIC SCSI-3	Transparent FABRIC SCSI-2 No CCL or SCSI-3	Transparent FABRIC SCSI-2 or SCSI-3	Transparent FABRIC SCSI-2 No CCL or SCSI-3	Transparent FABRIC SCSI-2 or SCSI-3
Tru64 UNIX 5.1, 5.1A TruCluster Server Version 5.1/5.1A	Multiple-Bus FABRIC SCSI-3	Transparent FABRIC SCSI-2 or SCSI-3	Transparent or Multiple-Bus FABRIC SCSI-2 or SCSI-3	11.0, 11i: Fabric attachment, With Zoning Transparent or Multiple-BusFABRIC SCIS-2 or SCSI-2	With Zoning Transparent or Multiple-Bus FABRIC SCSI-2 or SCSI-3	With Zoning Transparent FABRIC SCSI-3	Transparent or Multiple-Bus FABRIC SCSI-2 No CCL or SCSI-3	Transparent or Multiple-Bus FABRIC SCSI-2 or SCSI-3	Transparent FABRIC SCSI-2 No CCL or SCSI-*3*	Transparent or Multiple-Bus FABRIC SCSI-2 or SCSI-3
HP-UX 10.20, 11.0, 11i MC/ ServiceGuard Clusters	11.0, 11i: Fabric attachment, With Zoning Multiple-Bus FABRIC SCSI-3	11.0, 11i: Fabric attachment, With Zoning Transparent FABRIC SCSI-2 or SCSI-3	11.0, 11i: Fabric attachment, With Zoning Transparent or Multiple-Bus FABRIC SCIS-2 or SCSI-2	11.0, 11i: Fabric attachment, Transparent or Multiple-BusFABRIC SCSI-2 or SCSI-3 10.20, 11.0, 11i: FC-AL attachment Transparent or Multiple-Bus LOOP_HARD SCSI-2 or SCSI-3	11.0, 11i: Fabric attachment, With Zoning Transparent or Multiple-Bus FABRIC SCSI-2 or SCSI-3	11.0, 11i: Fabric attachment, With Zoning Transparent FABRIC SCSI-2 or SCSI-3	11.0, 11i: Fabric attachment, With Zoning Transparent or Multiple-Bus FABRIC SCSI-2 No CCL or SCSI-3	11.0, 11i: Fabric attachment, With Zoning Transparent or Multiple-Bus FABRIC SCSI-2 or SCSI-3	11.0, 11i: Fabric attachment, With Zoning Transparent FABRIC SCSI-2 No CCL or SCSI-3	11.0, 11i: Fabric attachment, With Zoning Transparent or Multiple-Bus FABRIC SCSI-2 or SCSI-3
IBM AIX 4.3.3, 5.1	With Zoning Multiple-Bus FABRIC SCSI-3	With Zoning Transparent FABRIC SCSI-2 or SCSI-3	With Zoning Transparent or Multiple-Bus FABRIC SCSI-2 or SCSI-3	11.0, 11i: Fabric attachment, With Zoning Transparent or Multiple-Bus FABRIC SCSI-2 or SCSI-3	Transparent or Multiple-Bus FABRIC SCSI-2 or SCSI-3	With Zoning Transparent FABRIC SCSI-3	With Zoning Transparent or Multiple-Bus FABRIC SCSI-2 No CCL or SCSI-3	With Zoning Transparent or Multiple Bus FABRIC SCSI-2 or SCSI-3	With Zoning Transparent FABRIC SCSI-2 No CCL or SCSI-3	With Zoning Transparent or Multiple-Bus FABRIC SCSI-2 or SCSI-3

Table 4–6: Platform Interoperability for Single Shared MA6000, MA/RA8000, EMA/ESA12000, or EMA16000 Storage System – ACS 8.6 (Continued)

Platform or Operating System	**Compaq OpenVMS** 7.2-2, 7.3 Clusters	**Tru64 UNIX** 4.0F Trucluster Software Products V1.6	**Tru64 UNIX** 5.1, 5.1A TruCluster Server Version 5.1/5.1A	**HP-UX** 10.20, 11.0, 11i MC/ ServiceGuard Clusters	**IBM AIX** 4.3.3, 5.1	**Linux** **Caldera** 2.3.1 **Redhat** 7.0, 7.1 **SuSE** 6.3, 7.0	**Microsoft Windows NT** 4.0 SP6a **Windows 2000** SP2 MSCS	**Novell NetWare** 4.2, 5.1, 6.0 Clusters 1.01, 1.06	**SGI IRIX** 6.5.11, 6.5.12	**Sun Solaris** 2.6 and 7, 8 (32/64-bit) Sun Clusters 2.2, VERITAS Clusters 1.2.1
Linux **Caldera** 2.3.1 **Redhat** 7.0, 7.1 **SuSE** 6.3, 7.0	Requires two storage systems	With Zoning Transparent FABRIC SCSI-3	With Zoning Transparent FABRIC SCSI-3	11.0, 11i: Fabric attachment, With Zoning Transparent FABRIC SCSI-2 or SCSI-3	With Zoning Transparent FABRIC SCSI-3	Transparent FABRIC SCSI-3	With Zoning Transparent FABRIC SCSI-3	Transparent FABRIC SCSI-3	Transparent FABRIC SCSI-3	Transparent FABRIC SCSI-3
Microsoft Windows NT 4.0 SP6a **Windows 2000** SP2 MSCS	Multiple-Bus FABRIC SCSI-3	Transparent FABRIC SCSI-2 No CCL or SCSI-3	Transparent or Multiple-Bus FABRIC SCSI-2 No CCL or SCSI-3	11.0, 11i: Fabric attachment, With Zoning Transparent or Multiple-Bus FABRIC SCSI-2 No CCL or SCSI-3	With Zoning Transparent or Multiple-Bus FABRIC SCSI-2 No CCL or SCSI-3	With Zoning Transparent FABRIC SCSI-3	Transparent or Multiple-Bus FABRIC SCSI-2 No CCL or SCSI-3	Transparent or Multiple-Bus FABRIC SCSI-2 No CCL or SCSI-3	Transparent FABRIC SCSI-2 No CCL or SCSI-3	Transparent or Multiple-Bus FABRIC SCSI-2 No CCL or SCSI-3
Novell NetWare 4.2, 5.1, 6.0 Clusters 1.01, 1.06	5.1, 6.0: Multiple-Bus FABRIC SCSI-3 4.2: Requires two storage systems	Transparent FABRIC SCSI-2 or SCSI-3	Transparent or Multiple-Bus FABRIC SCSI-2 or SCSI-3	11.0, 11i: Fabric attachment, With Zoning Transparent or Multiple-Bus FABRIC SCSI-2 or SCSI-3	With Zoning Transparent or Multiple Bus FABRIC SCSI-2 or SCSI-3	Transparent FABRIC SCSI-3	Transparent or Multiple-Bus FABRIC SCSI-2 No CCL or SCSI-3	Transparent or Multiple-Bus FABRIC SCSI-2 or SCSI-3	Transparent FABRIC SCSI-2 No CCL or SCSI-3	With Zoning Transparent or Multiple-Bus FABRIC SCSI-2 or SCSI-3
SGI IRIX 6.5.11, 6.5.12	Requires two storage systems	Transparent FABRIC SCSI-2 No CCL or SCSI-3	Transparent FABRIC SCSI-2 No CCL or SCSI-3	11.0, 11i: Fabric attachment, With Zoning Transparent FABRIC SCSI-2 No CCL or SCSI-3	With Zoning Transparent FABRIC SCSI-2 No CCL or SCSI-3	Transparent FABRIC SCSI-3	Transparent FABRIC SCSI-2 No CCL or SCSI-3	Transparent FABRIC SCSI-2 No CCL or SCSI-3	Transparent FABRIC SCSI-2 No CCL or SCSI-3	Transparent FABRIC SCSI-2 No CCL or SCSI-3
Sun Solaris 2.6 and 7, 8 (32/64-bit) Sun Clusters 2.2, VERITAS Clusters 1.2.1	Multiple-Bus FABRIC SCSI-3	Transparent FABRIC SCSI-2 or SCSI-3	Transparent or Multiple-Bus FABRIC SCSI-2 or SCSI-3	11.0, 11i: Fabric attachment, With Zoning Transparent or Multiple-Bus FABRIC SCSI-2 or SCSI-3	With Zoning Transparent or Multiple-Bus FABRIC SCSI-2 or SCSI-3	Transparent FABRIC SCSI-3	Transparent or Multiple-Bus FABRIC SCSI-2 No CCL or SCSI-3	With Zoning Transparent or Multiple-Bus FABRIC SCSI-2 or SCSI-3	Transparent FABRIC SCSI-2 No CCL or SCSI-3	Transparent or Multiple-Bus FABRIC SCSI-2 or SCSI-3

Booting from the SAN

Table 4–7 indicates the platforms and operating systems that are currently able to boot from SAN storage.

Table 4–7: SAN Boot by OS

Platform/Operating System	Comments
Microsoft Windows 2000, NT 4.0 MSCS	Enterprise Virtual Array: KGPSA-CB
Microsoft Windows 2000, NT 4.0, Datacenter MSCS	MA6000, MA/RA8000, EMA/ESA12000, and EMA16000 Storage: KGPSA-CB
OpenVMS VMS Clusters	Enterprise Virtual Array: KGPSA-BC/CA
	MA6000, MA/RA8000, EMA/ESA12000, and EMA16000 Storage: KGPSA-BC/CA – Requires use of *wwidmgr,* SRM console firmware v5.5 (minimum) KGPSA-DA – Requires use of *wwidmgr* and SRM console firmware v6.1 (minimum)
Tru64 UNIX TruCluster Software Products	Enterprise Virtual Array: KGPSA-BC/CA
	MA6000, MA/RA8000, EMA/ESA12000, and EMA16000 Storage: KGPSA-BC/CA – Requires use of *wwidmgr*, SRM console firmware v5.5 (minimum) KGPSA-DA – Requires use of *wwidmgr* and SRM console firmware v6.1 (minimum)

Storage System Rules

Enterprise Virtual Array Configuration Rules

1. The Enterprise Virtual Array Storage System is supported in all Compaq SAN Fabric topology configurations noted in this guide or in the platform application notes. The Enterprise Virtual Array is compatible in SANs using Compaq SAN Switch 8, 16, 8-EL, 16-EL, and SAN Switch Integrated 32/64 model switches.
2. For SANs with more than 1024 HBAs, use fabric zoning to limit the number of connections visible to each storage system to a maximum of 1024.
3. The supported platforms and operating systems are listed in Table 4–1
4. Shared access and heterogeneous platform zoning requirements are listed in Table 4–6.
5. Supports Multiple-Bus Failover mode only. Multiple-Bus Failover requires a minimum of 2 Fibre Channel HBAs and native operating system or separate, layered multi-path driver functionality.
6. Supports simultaneous access to Enterprise Virtual Array and MA6000, MA/RA8000, EMA/ESA12000, and EMA16000 storage systems from the same Server/HBA when using the supported Enterprise Virtual Array supported operating system and HBA driver

and firmware versions. This requires common HBA, driver, and multi-path software version support for the Enterprise Virtual Array and the MA6000, MA/RA8000, EMA/ESA12000, or EMA16000 storage systems. Refer to the:

Supplemental Tables for the Heterogeneous Open SAN Design Reference Guide.

7. All host table entries must have the proper operating system type parameter set based on the platform type accessing the assigned LUNs.
8. Servers accessing Enterprise Virtual Arrays must not have access to MSA1000 or RA4100 or RA4000 storage systems. Zoning is required when configuring these storage systems in the same physical SAN.
9. SSP/LUN level masking – Use storage system LUN presentation to enable/disable LUN access to specific hosts.

NOTE: Shared access between different servers to the same storage unit (LUN) requires specific application software (i.e., cluster software) to ensure proper data preservation.

Enterprise Virtual Array Maximums

Table 4–8 lists the maximum connections supported by Enterprise Virtual Array HSV110 controller-based storage systems. In addition, Table 4–8 lists the maximum supported storage limits for each hardware platform or operating system. The maximums shown here are for access to a single Enterprise Virtual Array with dual redundant HSV110 controllers. If the connection requirements for the number of servers in a particular SAN exceed the maximums, then deploy multiple storage systems within the SAN.

- Maximum of 1024 Host Bus Adapters (HBA)
- Maximum of 256 LUNs
- Maximum of 256 Servers: A server is defined as one or more HBAs
- The number of LUNs times the number of servers must not exceed 8192.

  Example: 64 LUNs on an Enterprise Virtual Array gives a maximum of 128 Servers that may access that storage system

Table 4–8: SAN/Platform Storage Maximums - Enterprise Virtual Arrays

Platform or Operating System	Host Bus Adapters per Server	Active Controller Ports (Targets) per HBA	LUNs per HBA Target
See **Reference Notes**	1	2	3, 4
Compaq OpenVMS 7.2-1H1, 7.2-2, 7.3	26	128	255
Tru64 UNIX 5.1, 5.1A	64	128	255
Microsoft Windows NT 4.0 SP6a Windows 2000 SP2	8	2/4	8/64
Sun Solaris 7, 8	16	2/4	128

Reference Notes

1. The maximum number of HBAs supported per server is dependent on the specific server model.
2. For Tru64 UNIX and OpenVMS this column typically represents the total number of active controller ports per HBA when accessing all ports of a storage system or ports on multiple storage systems. For all other platforms, this column typically represents 2 ports per storage system, or a total of 4 ports across 2 storage systems. Use of zoning may be required to limit the number of active targets presented to each HBA to the maximums stated for each platform in this column.
3. For Microsoft Windows, supports 8 LUNs per HBA target with the Large LUN feature disabled and 64 LUNs per HBA target with the Large LUN feature enabled.
4. For Sun configurations utilizing the same HBA accessing both an Enterprise and a MA6000, MA/RA8000, EMA/ESA12000, or EMA16000 storage system, the maximum number of LUNs per HBA target is reduced to 64.

MA6000, MA/RA8000, EMA/ESA12000, and EMA16000 Configuration Rules

These storage systems are supported in all Compaq SAN Fabric topology configurations noted in this guide or in the platform application notes. These storage systems are supported using any Compaq Fibre Channel fabric switch model. (Refer to Chapter 3 for rules related to specific switch models.)

1. Limit the number of connections visible to each storage system to the maximum number available by using fabric zoning. The number of connections available is based on the version of Array Controller Software (ACS).
2. The supported platforms and operating systems are listed in Table 4–1.
3. Shared access and heterogeneous platform zoning requirements are listed in Table 4–6. The heterogeneous platform and operating system mix in the SAN determines the appropriate controller topology attachment, SCSI mode, and Command Console LUN settings for shared storage systems. Refer to Table 4–1
4. Single or dual redundant controller configurations are supported. For dual redundant controllers, the available failover modes are Transparent and Multiple-Bus. Multiple-Bus failover requires native operating system or separate multi-path driver functionality.
5. All host connection table entries must have the proper operating system type parameter set based on the platform type accessing the assigned LUNs.
6. Supports simultaneous access to MA6000, MA/RA8000, EMA/ESA12000, or EMA16000 storage systems and the Enterprise Virtual Array from the same Server/HBA when using the supported Enterprise Virtual Array supported operating system and HBA driver and firmware versions. This requires common HBA, driver, and multi-path software version support for the Enterprise Virtual Array and the MA6000, MA/RA8000, EMA/ESA12000, or EMA16000 storage systems.

Refer to the
Supplemental Tables for the Heterogeneous Open SAN Design Reference Guide.
Refer to Chapter 3 for information about configuring the SANworks SAN Appliance to manage Enterprise Virtual Arrays and MA6000, MA/RA8000, EMA/ESA12000, EMA16000 storage systems in the same SAN.

7. Servers accessing MA6000, MA/RA8000, EMA/ESA12000, and EMA16000 storage systems must not have access to MSA1000, RA4000, or RA4100 storage systems. Zoning is required to prevent access from servers to multiple storage system types when configuring these storage systems in the same physical SAN.

8. SSP/LUN level masking – Use storage system Selective Storage Presentation to enable/disable LUN access to specific connections. Use the unit offset feature to provide needed LUN numbering for host connections. The default LUN numbering for Transparent Failover mode is 0 to 99 for controller port 1 and 100 to 199 for controller port 2. For Multiple-bus Failover mode the default LUN numbering is 0 to 199 on all controller ports.

NOTE: Shared access between different servers to the same storage unit (LUN) requires specific application software (i.e., cluster software) to ensure proper data preservation.

9. F-Port fabric attachment to the SAN is available through all Compaq Fibre Channel 8, 16, SAN Switch 8, 16, 8-EL, and 16-EL models. Controller setting is FABRIC topology.

10. FL-Port fabric loop attachment to the SAN with QuickLoop is available through Compaq SAN Switch 8 and 16 switches. Controller port topology set to "LOOP_HARD".

11. All controller ports must be set to the same topology type.

Maximum Paths or Maximum LUNs

For MA6000, MA/RA8000, EMA/ESA12000, and EMA16000 storage systems use the HSG60/80 controller Unit Offset feature to maximize path accessibility or to maximize the number of LUNs.

- **For Maximum Controller Path Accessibility to the same set of LUNs**

  Use a common unit offset value for all 4 controller ports. Access to a common set of LUNs through all 4 controller host ports is provided by using the same unit offset value on all controller host port connections for each server. For example, set the unit offset value for connections on all 4 controller ports to zero (0) for a given server. The server will be capable of accessing one set of LUNs beginning with LUN 0 from all 4 controller host ports. This method provides for the highest number of paths to a given set of LUNs.

- **For Maximum LUN Count**

  Use distinct controller port unit offsets for each port pair. Access one set of LUNs with controller port 1 of each controller and access a different set of LUNs with controller port 2 of each controller. For example, set the unit offset value for connections on controller port 1 of each controller to zero, and then set a unit offset value for connections on controller port 2 of each controller to 100 for a given server. The server will be capable of accessing one set of LUNs beginning with LUN 0 through controller port 1 on each controller, and a second set of LUNs beginning with LUN 100 through controller port 2 of each controller. This method provides the highest number of LUNs accessed through a maximum of 2 paths. It also allows for the highest number of servers. Refer to Table 4–9.

MA6000, MA/RA8000, EMA/ESA12000, and EMA16000 Maximums

Table 4–9 lists the maximum supported storage limits for each hardware platform or operating system. The maximums shown are for access to MA6000 storage systems with dual redundant HSG60 controllers, and MA/RA8000, EMA/ESA12000, or EMA16000 storage systems with dual redundant HSG80 controllers. If the maximums listed are below the requirements for the number for servers required, deploy multiple storage systems within the SAN.

Table 4–9: Platform Maximums - MA6000, MA/RA8000, EMA/ESA12000, EMA16000 Storage Systems Using ACS 8.6

Platform or Operating System	Host Bus Adapters per Server	Active Controller Ports (Targets)per HBA	LUNs per HBA Target	Servers or HBA Connections per Active Controller Port		Servers or HBA Connections per Storage System	
Controller Failover Mode				TF	MB	TF	MB
Compaq OpenVMS 7.2-1H1, 7.2-2, 7.3	26	128	128 (255)	32	24	64	48
Tru64 UNIX 4.0F	32	4	8	4	4	8	8
Tru64 UNIX 5.1, 5.1A	64	128	128 (255)	48	24	96	48
HP-UX 10.20, 11.0, 11i	16	2/4	8	8	8	16	16
IBM AIX 4.3.3, 5.1	4	2/4	16	6	6	12	12
Linux Caldera Alpha/Intel 2.3.1	2	2/4	64	2	2	4	4
Linux Redhat Alpha/Intel 7.0, 7.1	2	2/4	64	2	2	4	4
Linux SuSE Alpha /Intel 6.3, 7.0	2	2/4	64	2	2	4	4
Microsoft Windows NT 4.0 SP6a Windows 2000 SP2	8	2/4	8/64	8	8	16	16
Microsoft Datacenter	4	2/4	64	4	4	8	8
Novell NetWare 4.2, 5.0, 6.1	4	2/4	32	8	8	16	16
SGI IRIX 6.5.11, 6.5.12	4	2/4	64	3	3	6	6
SUN Solaris 2.6, 7 & 8 (32/64 bit)	16	2/4	64	8	8	16	16
Heterogeneous SAN Reference Notes				8		8	
Reference Notes	1	2	3, 4, 5	6, 7		7	

Reference Notes

1. The maximum number of HBAs supported per server depends on the specific server model.
2. For Tru64 UNIX and OpenVMS this column typically represents the total number of active controller ports per HBA when accessing all ports of a storage system or ports on multiple storage systems. For all other platforms this column typically represents 2 ports per storage system, or a total of 4 ports across 2 storage systems. Use of zoning may be required to limit the number of active targets (controller ports) presented to each HBA to the maximums stated for each platform in this column. A minimum of OpenVMS 7.2-2 is required for the indicated maximum.
3. Numbers in this column are reduced by one if the command console LUN is enabled.

4. Microsoft Windows NT supports 8 LUNs per HBA target with Large LUN feature disabled and 64 LUNs per HBA target with Large LUN feature enabled. Windows 2000 supports Large LUN by default.
5. For OpenVMS and Tru64 5.1, 5.1A, the operating system maximum is 255 LUNs per target. The single storage system maximum is 128 LUNs.
6. These are the maximum number of HBAs that can be configured for access to an active controller port. Assumes 1 HBA per server for single path using controller transparent failover or 2 HBAs per server for multi-path using controller multiple-bus failover. For transparent failover, the limit is specified by controller port pair–1 active and 1 standby controller port. For multiple-bus failover, the limit is specified per single active port.
7. The maximums specified for each platform are the result of one or more of following limiting conditions:
   - — A qualification limit
   - — Command flow queuing characteristics of specific HBA drivers
   - — Connection table size in the array controller software in conjunction with the number of HBA to controller port paths.

   For maximum server or HBA connectivity, limit the number of active HBA to controller port paths to one per server for controller transparent failover mode (Figure 4–1) and two per server for controller multiple-bus failover (Figure 4–2). The use of zoning is required to limit the number of active controller ports visible to each HBA.

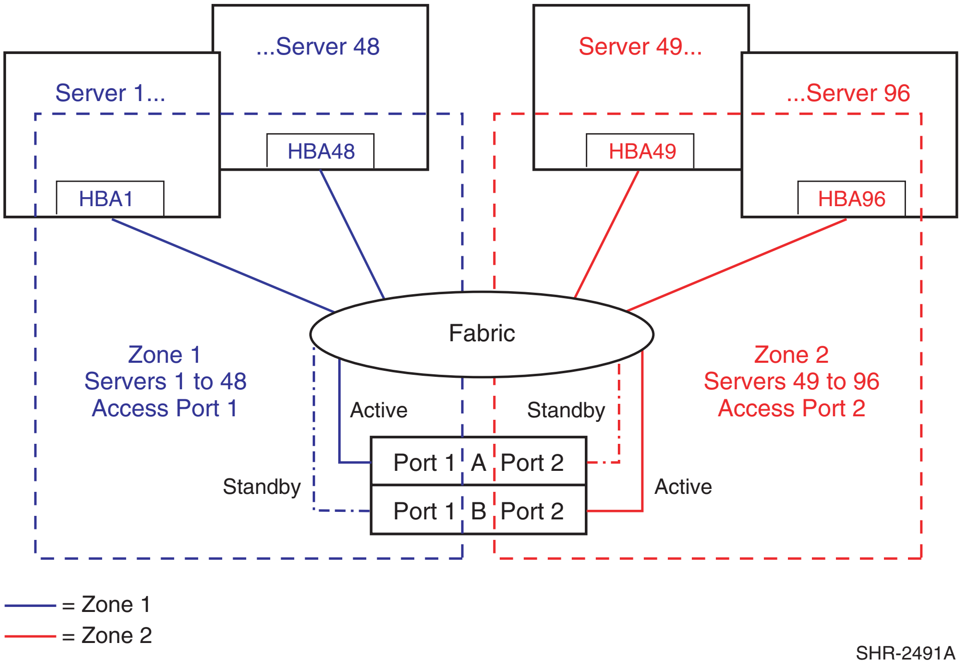

Figure 4–1: Maximum server example for Tru64 UNIX with transparent failover using 96 connections

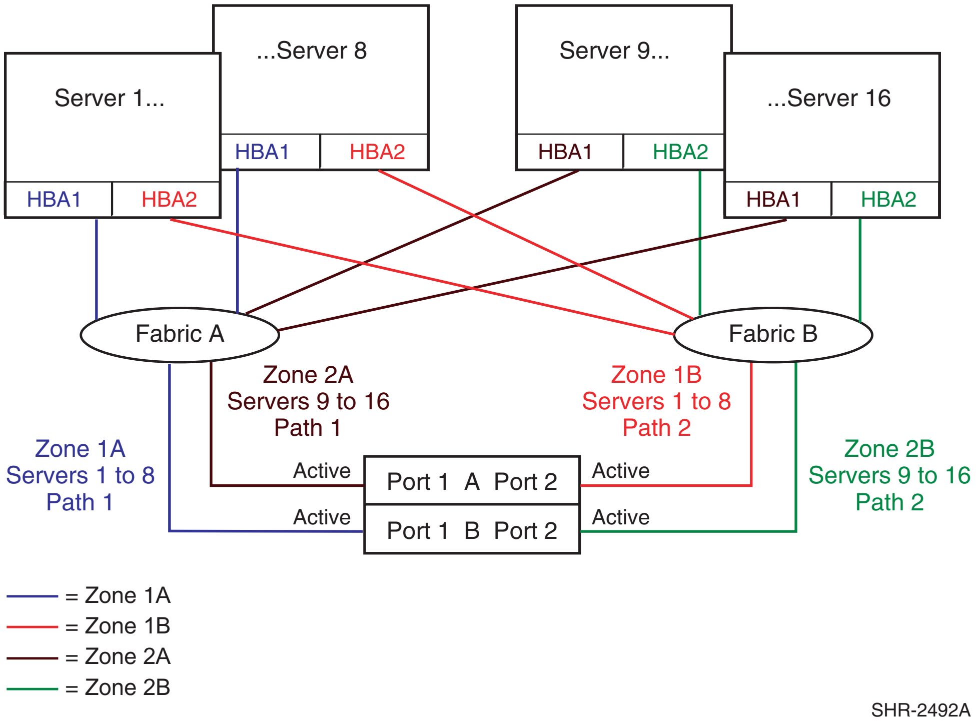

Figure 4–2: Maximum server example for 16 servers with multiple-bus failover

8. In a heterogeneous SAN, the maximum number of servers or HBAs is equal to the lowest maximum listed in these columns for the operating systems that are sharing the storage system. Microsoft Windows 2000 Datacenter is not supported for shared access to a storage system with other heterogeneous servers

Specific Platform/Operating System Rules – MSA1000, RA4100, RA4000

Table 4–10 with the supported storage system types. The table lists the platforms, HBAs, storage system types, and the preferred method of storage attachment for each platform type. The sections following the table describe additional rules not listed in the table.

For the most current information on HBAs, firmware, and drivers see this document:

Supplemental Tables for the Heterogeneous Open SAN Design Reference Guide

Table 4–10: Platform/Storage System SAN Attachment Summary

Platform or Operating System, Storage System in ( )	Platform HBA SAN Attachment	Multi-Path Support	MSA1000 Storage System SAN Attachment	RA4100, RA4000 Storage System SAN Attachment
Linux Redhat 7.0, 7.1	RA4000/4100: FL-Port 120186-B21 223180-B21	N/A	N/A	FL-Port to SAN Switch
Linux SuSE 7.1	RA4000/4100: FL-Port 120186-B21 223180-B21	N/A	N/A	FL-Port to SAN Switch
Microsoft Windows 2000 SP2 (MSA) Microsoft Windows NT 4.0 SP6a (MSA) Microsoft Windows 2000 SP1 (RA) Microsoft Windows NT 4.0 SP5, SP6a (RA) MSCS	MSA1000: F-Port FCA-2101 RA4000/4100: FL-Port 120186-B21 223180-B21	Secure Path	F-Port to SAN Switch	FL-Port to SAN Switch FC-AL Switch cascaded to a SAN Switch
Novell NetWare 5.1	FL-Port 120186-B21 223180-B21	Secure Path	N/A	FL-Port to SAN Switch FC-AL Switch cascaded to a SAN Switch

Heterogeneous SAN Platform Interoperability for MSA1000 Storage

This section specifies the rules for shared access to a single MSA1000 storage system. MSA1000 storage systems are supported for shared access subject to the following guidelines:

- Server with Windows 2000 SP2
- Server with Windows NT 4.0 SP6a
- Clusters can not share MSA1000(s) with other clusters
- Clusters can not share MSA1000(s) with standalone servers
- Servers can only support a single or redundant path (not both)

MSA1000 Configuration Rules

- The MSA1000 storage systems can be configured in a SAN directly using the Compaq SAN Switch 8, 16, 8-EL, and 16-EL models. Table 4–10 lists the platforms and operating systems that are supported using these storage systems. Supports single or redundant controllers with Active/Passive controllers
- Use ACU to enable/disable LUN access to specific connections
- Servers accessing MSA1000 storage systems must not have access to Enterprise, MA6000, MA/RA8000, EMA/ESA12000, EMA16000, RA4000, or RA4100 storage systems. Zoning is required to prevent access from servers to multiple storage system types when configuring these storage systems in the same physical SAN.
- Zoning is required with multiple clusters when using the Compaq Fibre Channel Switch 8, 16, 8-EL, and 16-EL.

MSA1000 Maximums

The following are the maximum configurations for MSA1000 systems:

With 8 ports:

- 6 servers
- 2 two-node clusters
- 5 MSA1000

With 16 ports:

- 12 servers
- 5 two-node clusters
- 10 MSA1000

With 8 port cascaded to 8 port:

- 10 servers
- 4 two-node clusters
- 10 MSA1000

With 16 port cascaded to 8 port:

- 15 servers
- 15 MSA1000

With 16 port cascaded to 16 port:

- 20 servers
- 10 MSA1000

Heterogeneous SAN Platform Interoperability for RA4100 and RA4000 Storage Systems

This section specifies the rules for shared access to a single RA4100 or RA4000 storage system. RA4100/4000 storage systems are supported for shared access with any combination of the following operating systems:

— Linux Redhat 7.0, 7.1

— Linux SuSE 7.1

— Microsoft Windows 2000 SP1

— Microsoft Windows NT 4.0, SP5, SP6a

— Novell NetWare 5.1

- RA4100/RA4000 systems can not be shared by more than one cluster when using Microsoft Windows NT 4.0 or Microsoft Windows 2000
- RA4100/RA4000 systems owned by a Microsoft Windows NT or Microsoft Windows 2000 cluster can not be shared with a standalone server or server

RA4100 and RA4000 Configuration Rules

These storage systems can be configured in a SAN directly using the Compaq Fibre Channel SAN Switch 8, 16, 8-EL, 16-EL, and indirectly through the Compaq FC-AL Switch 8 cascaded to Compaq SAN Switch 8, 16, 8-EL, or 16-EL models. Table 4–10 lists the platforms and operating systems that are supported using these storage systems. Refer to Table 4–10 for the specific shared access rules when different platforms or operating systems need shared access to the same RA4100/4000 storage system. FL-Port fabric attachment to the SAN is available through Compaq SAN Switch 8, 16, 8-EL, 16-EL models.

The Compaq FC-AL Switch 8 is supported for cascaded attachment to the SAN through a single FL-Port on a Compaq SAN Switch 8, 16, 8-EL, and 16-EL.

1. Use single or redundant controllers with Active/Passive controllers.
2. Use ACU to enable/disable LUN access to specific connections.
3. For RA4100/RA4000 SAN configurations with *heavy I/O traffic*, it is necessary to increase the fabric switch buffer capacity from the default value of 16 to 27.
4. Servers accessing RA4100 or RA4000 storage systems must not have access to Enterprise, MA6000, MA/RA8000, EMA/ESA12000, EMA16000, or MSA1000 storage systems. Zoning is required to prevent access from servers to multiple storage system types when configuring these storage systems in the same physical SAN.
5. Zoning is required with multiple clusters when using the Compaq Fibre Channel Switch 8, 16, 8-EL, and 16-EL.

RA4100 and RA4000 Maximums

The following are the maximum configurations for RA4100/RA4000 systems.

With 8 ports:

- 6 servers
- 3 two-node Clusters
- 6 RA4100/RA4000

With 11 ports:

- 6 servers
- 3 two-node clusters
- 9 RA4100/RA4000

With 16 ports:

- 10 servers
- 5 two-node Clusters
- 14 RA4100/RA4000

SAN/DRM Integration

The Compaq Data Replication Manager (DRM) is approved for use within a larger Heterogeneous Open SAN provided the following additional rules are followed: All DRM implementations require Level 4 NSPOF SANs using two separate fabrics. Refer to Chapter 2, "Data Availability in a SAN". Several special purpose DRM configurations are also supported as defined in the DRM Features and Benefits Guide which is available on the Web at:

> DRM Technical Documentation
> http://www.compaq.com/products/sanworks/drm/documentation.html

Each shared storage array must adhere to the DRM sharing rules as defined in the DRM Design Guide. These sharing rules may be more restrictive than those in this guide due to the requirements for DRM, for example, the operating system must support multiple-bus failover. In addition, the current DRM solution supports a sub-set of those operating systems listed in this guide.

1. Shared usage of the DRM configured storage systems by non-DRM configured servers (e.g. running in transparent failover) or non-DRM supported operating systems is not supported.

2. All servers sharing the same storage sub-system must share a compatible SCSI command mode as shown by a yes in the following table:

Table 4–11: Heterogeneous DRM Operating Systems

Operating System	Versions	SCSI-2	SCSI-3
Compaq OpenVMS	7.2-1H1, 7.2-2, 7.3	No	Yes
Compaq Tru64 UNIX	5.1, 5.1a	Yes	Yes
HP-UX	11.0	Yes	Yes
IBM AIX	4.3.3	Yes	Yes
Microsoft Windows NT	4.0	Yes	Yes
Microsoft Windows 2000	Svr., Ad. Svr., DC	Yes	Yes
Novell NetWare	5.1, 6.0	Yes	Yes
Sun Solaris	2.6, 7, 8	Yes	No (2.6) Yes (7, 8)

3. Each DRM solution may contain up to 20 switches supporting, for example, up to 96 servers and 8 storage arrays per site. In some cases the actual limit will be smaller due to restrictions imposed by the intersite link.   DRM supports the limit of up to seven hops between devices, with the understanding that there are three links involved. There is the host to local storage link, the local storage to remote storage link, and the local host to remote storage link. Each of these links must not exceed seven hops. All active/standby host-to-storage links as well as local-to-remote storage links must conform to the 7-hop limit

4. Each DRM solution instance may contain up to 12 servers per storage system per site. At one remote copy set per server and a maximum of 8 storage systems per site per instance, a single instance can support up to 96 servers per site.

5. DRM over ATM configurations are limited to two Fibre Channel switches per fabric, for a total of 4 switches, one at each end of each fabric (2 fabrics, times 2 switches per fabric equals 4 switches). Cascaded switches are not supported. This no-cascaded switch

restriction also includes non-support for the SAN Switch Integrated/32 or SAN Switch Integrated/64 port switches due to the fact that the bigger switch is built up internally with six, 16-port switches that are cascaded together.

6. The DRM Link supports mixed heterogeneous SAN, DRM, and Host Based Shadowing traffic.
7. StorageWorks Command Console (SWCC) and the SANworks Management Appliance (SWMA) element manger can be used for initial setup of Data Replication Manager (DRM) storage sub-systems. However, neither of these tools should be used for DRM failover and failback operations. Therefore to prevent any potential inference by SWMA polling of the HSG80 when running DRM scripts, it is recommended that the SWMA be removed from all DRM zones before running the scripts..
8. Please see the DRM release notes for current information on any hop count restrictions between devices.

SAN/OpenVMS Host Based Shadowing Integration

Compaq OpenVMS servers implementing Host Based Shadowing are supported, integrated in a heterogeneous SAN with remote shadowset distances of up to 160 km. The long distance link supports mixed heterogeneous SAN, DRM, and Host Based Shadowing traffic.

StorageWorks CSS 2105 Storage System Interoperability and Integration

Compaq provides support for heterogeneous multi-vendor online storage interoperability on a common SAN. This support includes both the StorageWorks by Compaq Centralized Shared Storage 2105 (CSS 2105) and the StorageWorks by Compaq Enterprise RAID Array.

The initial integration support represents the first phase or level of interoperability. This level of support provides for:

1. Coexistence of Compaq and IBM storage systems in a common heterogeneous Open SAN. The Compaq and IBM storage systems operate in separate fabric zones within the same physical SAN.
2. Data migration support between Compaq and IBM storage systems using a shared server running either Windows 2000/NT, IBM AIX, or Sun Solaris.
3. Multi-path failover capabilities using Compaq Secure Path for the Compaq storage and the IBM Subsystem Device Driver on a single shared server running Windows NT, IBM AIX, or Sun Solaris. Each storage system is connected to the server using independent HBA pairs.
4. Simultaneous enterprise backup support from both the Compaq storage and IBM storage utilizing a single shared server and the Compaq Enterprise Backup Solution with VERITAS NetBackup to a common tape library for Windows 2000/NT.

Refer to the Compaq Technical Note "Compaq StorageWorks Centralized Shared Storage 2105 Interoperability" for additional information. Future phases will provide additional levels of interoperability over time.

Table 4–12 lists the specific products and versions supported for integration of the CSS 2105 and the StorageWorks Enterprise RAID Array.

Table 4–12: CSS 2105 Interoperability Support

Operating System	Window NT V4.0 SP6	Window 2000 SP1	IBM AIX V4.3.3	Sun Solaris V7.0
Host Bus Adapter (CSS 2105)	Qlogic QLA 2200F v7.05.02 BIOS v1.61	Qlogic QLA 2200F v7.05.02 BIOS v1.61	Emulex LP7000E - 380574-001 (KGPSA-BC) IBM 6227	JNI 64bit SBUS - 123503-001 (DS-SWSA4-PC)
Host Bus Adapter (MA Zone)	Emulex LP8000 64Bit/33Mhz PCI to FC HBA for Windows -176479-B21 (DS-KGPSA-CB)	Emulex LP8000 64Bit/33Mhz PCI to FC HBA for Windows -176479-B21 (DS-KGPSA-CB)	Cambex PCI Adapter for RS/6000 FC-AL/FC-SW - 197819-B21 (DS-SWIA 1-PD	JNI 32 PCI Bus-380576-001 (SWSA4-PC)
Switch	Brocade Fibre Channel Switch Silkworm 2250 Type 5.4, Kernel v5.3.1, Fabric OS v2.1.7	Brocade Fibre Channel Switch Silkworm 2250 Type 5.4, Kernel v5.3.1, Fabric OS v2.1.7	Brocade Fibre Channel Switch Silkworm 2250 Type 5.4, Kernel v5.3.1, Fabric OS v2.1.7	Brocade Fibre Channel Switch Silkworm 2250 Type 5.4, Kernel v5.3.1, Fabric OS v2.1.7
Compaq Solutions Platform Kit	StorageWorks Command Console v8.5c	StorageWorks Command Console v8.5c	StorageWorks Command Console v8.5c	StorageWorks Command Console v8.5c
Compaq Multipath	Secure Path v3.1	Secure Path v3.1	Secure Path v3.1	Secure Path v3.1
IBM Multipath	Subsystem Device Driver (SDD) v4.3	Subsystem Device Driver (SDD) v4.3	Subsystem Device Driver (SDD) v4.3	Subsystem Device Driver (SDD) v4.3
Compaq-IBM MultiPath Coexistence	Supported	Not yet supported	Supported	Supported
Path Fail	Supported	Supported	Supported	Supported
Cluster Software	Microsoft Enterprise Edition Clustering Software (MSCS)	Microsoft Enterprise Edition Clustering Software (MSCS)	Not yet supported	VERITAS Cluster Server v1.1.1
Disk-to-Disk Migration	Supported	Supported	Supported	Supported
Single Server Common EBS	VERITAS NetBackup v3.4	VERITAS NetBackup v3.4	Not yet supported	VERITAS NetBackup v3.4
Dual Server Common EBS	VERITAS NetBackup v3.4	VERITAS NetBackup v3.4	Not yet supported	VERITAS NetBackup v3.4
Volume Management	Enterprise Volume Manager v1.1a	Enterprise Volume Manager v1.1a	Not supported	Enterprise Volume Manager v1.1a

NOTE: Compaq supports a maximum of four switches in a fabric using switch firmware 2.1.7.

High Availability Configuration Considerations

Cabling Scheme Options

This section describes cabling scheme options for implementing high availability multi-path configurations for Enterprise Virtual Array, MA6000, MA/RA8000, EMA/ESA12000, and EMA16000 storage systems. Figure 4–3 and Figure 4–4 show cabling options when implementing a Level 4 high availability no single point of failure configuration. Figure 4–5 and Figure 4–6 show the cabling and associated zoning requirements when implementing a level 3 high availability configuration. Refer to Chapter 2, "Levels of Availability" for a description of the availability levels.

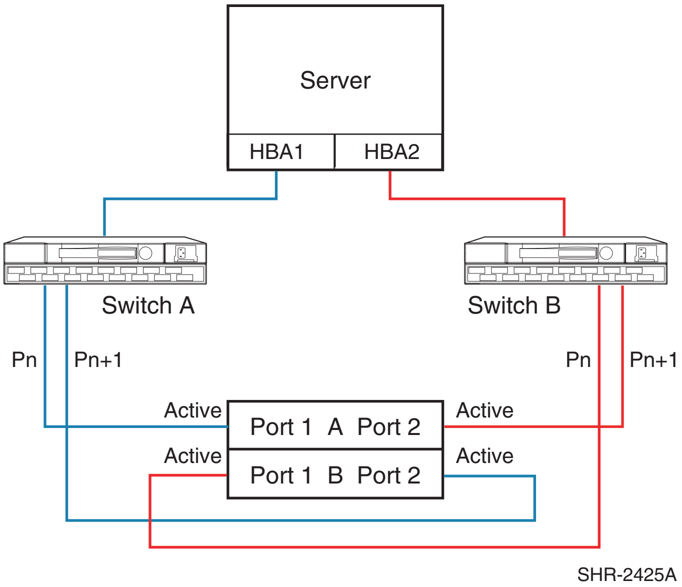

Figure 4–3: Cross-Cable High Availability NSPOF Configuration

Figure 4–3 shows the physical connections for a cross cable, high availability, no single-point of failure configuration for storage systems using two separate fabrics. The advantage of this cabling scheme is that it allows for both controllers to be active and in use should any path failure occur (other than a controller failure).

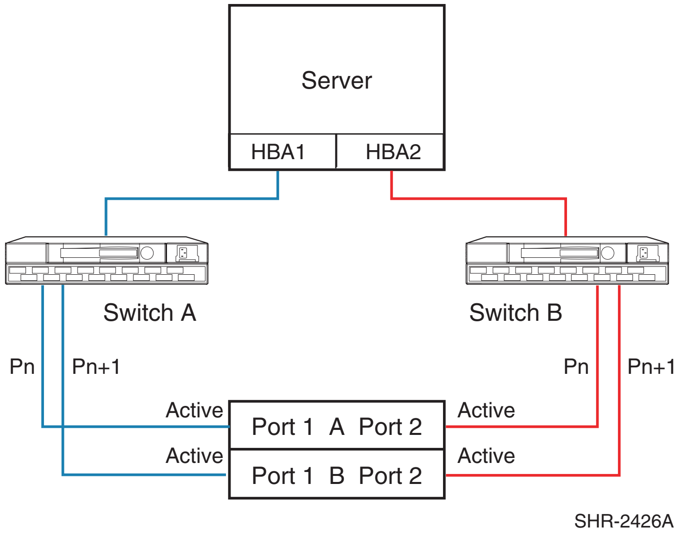

Figure 4–4: Straight-Cable High Availability NSPOF Configuration

Figure 4–4 shows the physical connections for a straight cable, high availability, no single point-of-failure configuration. The advantage of this cabling scheme is that it is the same cabling scheme used in Transparent failover mode for MA6000, MA/RA8000, EMA/ESA12000, and EMA16000 storage systems. This allows you to migrate from Transparent failover mode to Multiple-Bus failover mode without the need to re-cable the controller connections.

Figure 4–5 and Figure 4–6 below specify the logical path zoning that may be required for cross cable and straight cable configurations when implementing a level 3 single fabric high availability configuration. The requirement to zone separate logical paths in single fabric high availability implementations is O/S and platform specific. The zoning specified enforces and effectively results in the same configuration as physically depicted in Figure 4–3 and Figure 4–4. Single fabric cross cable implementations require cross port zoning, straight cable implementations require straight port zoning. In order to provide high availability, ensure each HBA is cabled to a different switch and configured for access to specific controller ports.

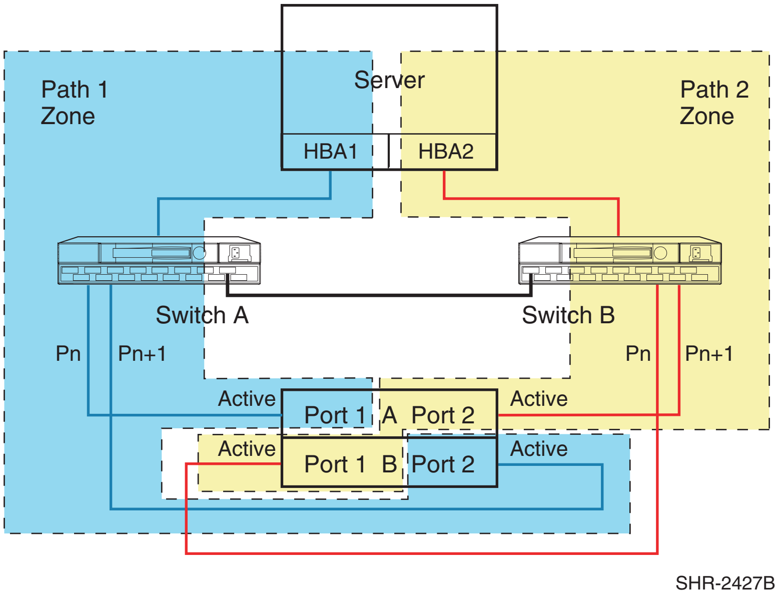

Figure 4–5: Cross-Cable High Availability Single Fabric Zoned Configuration

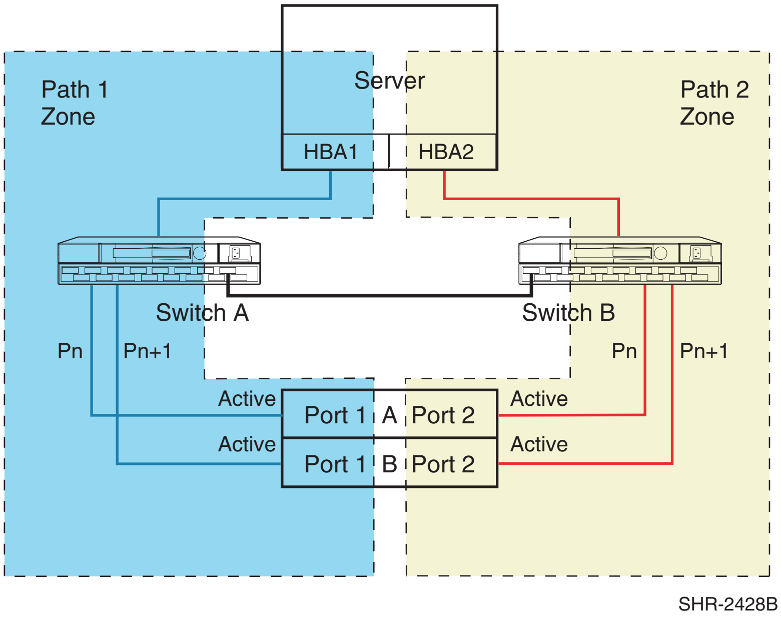

Figure 4–6: Straight-Cable High Availability Single Fabric Zoned Configuration

For two or more high availability server configurations, it is suggested that the first adapter in each server be connected to the first (same) Fibre Channel switch, the second two adapters to the second switch, etc. For example:

- Server 1 Fibre Channel HBA 1 to Fibre Channel Switch 1 - Switch Port 1
- Server 1 Fibre Channel HBA 2 to Fibre Channel Switch 2 - Switch Port 1
- Server 2 Fibre Channel HBA 1 to Fibre Channel Switch 1 - Switch Port 2
- Server 2 Fibre Channel HBA 2 to Fibre Channel Switch 2 - Switch Port 2

It is highly recommended that the cabling scheme shown in each Secure Path multiple-bus configuration be followed as depicted. This is not required; however, it does aid in understanding logical to physical LUN and path mapping for maintenance purposes.

5

Enterprise Backup Solution

Overview

The StorageWorks by Compaq Enterprise Backup Solution (EBS) is an integration of Independent Software Vendor (ISV) backup and restore application software. EBS uses StorageWorks by Compaq SAN hardware, including tape libraries, providing a complete enterprise-class backup solution. It supports the functionality and management of Heterogeneous StorageWorks by Compaq SANs, using both Fibre Channel switches and hubs, as described in this guide. This chapter describes the EBS configurations currently supported, the rules and recommendations for each ISV backup application, and how to efficiently and effectively provide shared tape library backup in a heterogeneous SAN environment.

Solution Benefits

Better Price/Performance

- Allows a high degree of scaling by adding more tape drives to an existing tape library, or by adding a new tape library containing one or more tape drives.
- Provides multiple tape drives streaming for maximum performance.
- Eliminates the data burden on the corporate communication network.
- Provides resource storage sharing and increased performance to all servers on a switched fabric.

Centralized Management

- All backup and restore jobs can be scheduled and managed from a single server.
- Management of shared tape libraries compared to many stand-alone tape libraries.
- Centralization of the tape libraries and tape drives will optimize and streamline the management and administration of the physical tape media.

Lower Total Cost of Ownership and Investment Protection

- Integrate legacy digital linear tape (DLT) libraries.
- Redeploy existing infrastructure.
- Leverage existing software investments.
- Leverage existing training.
- Lower cost per server by sharing central backup devices.
- No need to purchase separate switches for different operating systems.
- Shared tape library management, compared to many stand-alone servers.

Backup Applications

As of the writing of this document, EBS has certified the following Independent Software Vendors’ (ISV) backup applications:

- VERITAS NetBackup v3.4.1 for Windows NT/2K, Sun Solaris, Tru64 UNIX, AIX
- VERITAS Backup Exec v8.6 for Windows NT/2K, and v9.0 for NetWare
- Legato NetWorker v6.1 for Windows NT/2K, Sun Solaris, Tru64 UNIX
- Tivoli Storage Manager v4.2 for Windows NT/2K, Sun Solaris, AIX
- CA ArcServe v2000 for Windows NT/2K, and v7.0 for NetWare
- CommVault Galaxy v3.1 for Windows NT/2K, Sun Solaris

For the complete EBS configuration support information, refer to the Compaq web site at:

http://www.compaq.com/products/storageworks/ebs/EBScompatmatrix.html

Basic Storage Domain

The following diagram depicts a basic EBS storage domain or data zone. This diagram does not represent the actual cabling, or is not intended to be a complete representation of a SAN configuration. For more information please refer to the specific component/EBS references listed at the end of this chapter

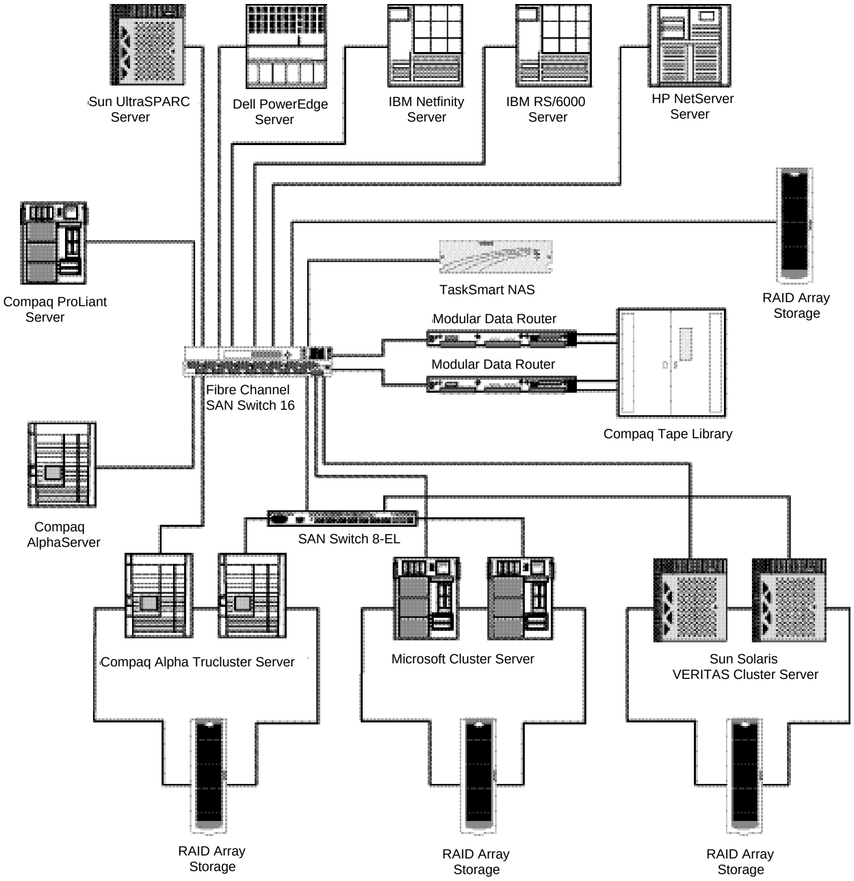

Figure 5–1: EBS Storage Domain

Supported Components

The following components make up the EBS:

- Servers
- Fibre Channel Host Bus Adapters
- SAN Switch, Fibre Channel Arbitrated Loop (FC-AL) Switch, or Fibre Channel Hub
- Fibre Channel to SCSI Bridge
- Tape Library

Servers

Compaq ProLiant Server

Operating System

- Windows NT 4.0 (w/SP6a or higher)
- Windows 2000 (w/SP2.0 or higher)
- NetWare 4.1x (w/SPIWSP6A or higher)
- NetWare 5.0 (w/SP5.0 or higher)
- NetWare 5.1 (w/SP2a or higher)

ProLiant Servers

- PL ML330, ML350, PL800, PL ML360, PL1500, PL ML370, PL1600
- PL DL380, PL1850R, PL CL380, CL1850, PL 2500
- PL ML530, PL 3000, PL 4500, PL 5000, ML570, PL 5500
- PL 6000, DL580, PL 6400 PL 6500, PL 7000, ML750, PL8000
- DL750 8500, ML770

Compaq AlphaServer

Operating Systems

- Tru64 UNIX 4.0f
- Tru64 UNIX 5.0a, 5.1, 5.1a

Compatible Compaq Tru64 UNIX Servers

- AS800, AS1000A, AS1200, AS4x00, AS8x00
- DS10, DS20
- ES40
- GS60, GS80, GS140, GS160, GS320

Sun UltraSPARC Server

Operating Systems

- Solaris Version 2. 6, 7, or 8 (with the latest patches)

Compatible Servers

- Sun ULTRASparc
- Enterprise:
  - — 10000, 6500, 5500, 4500, 3300
  - — 450, 420R, 250, 220R
  - — U10S,U5S

Third Party x86 Servers

Operating Systems

- Windows NT 4.0 (w/SP5.0 or higher)
- Windows 2000 (w/SP2.0 or higher)
- NetWare 4.1x (w/SPIWSP6A or higher)
- NetWare 5.0 (w/SP5.0 or higher)
- NetWare 5.1 (w/SP1.0 or higher)

Compatible Servers

- Dell PowerEdge Series
- HP Netserver Series
- IBM Netfinity Series

TaskSmart N-Series

- TaskSmart N2400 in a SAN Switch environment.

NOTE: EBS does not currently support Open File Option with the TaskSmart N2400.

Host Bus Adapters

- Fibre Channel Host Bus Adapter (KGPSA-BC, KGPSA-CB, and KGPSA-CA) for Windows NT, Windows 2000 and Tru64 UNIX
- Fibre Channel Host Bus Adapter (SWSA4-SB, SWSA4-SC, and SWSA4-PC) for Sun Solaris
- Compaq Host Bus Adapters (FC HBA, PCI, 32bit, 64bit, FCAL, and FC- AL Switch) for Windows NT, Windows 2000 and NetWare
- PCI-FC 197819-B21 (PC1000) Host Bus Adapter for AIX
- QLA2200 Host Bus Adaptor for Linux

NOTE: Follow the instructions in the Compaq Fibre Channel Host Controller Installation Guide for installing and configuring the Host Bus Adapter in each server. For servers running Microsoft Windows, the KGPSA-BC and KBPSA-CB topology parameter must be set = 1 to enable switched fabric (F-Port) support. For servers running Sun Solaris, the host bus adapter must be set to fabric mode during driver installation. This can be done by installing HSG80 platform kit or manually by editing the *fcat.conf* file. For more information, refer to the driver Readme file.

Switches and Hubs

- StorageWorks Fibre Channel SAN Switch 16 and Fibre Channel SAN Switch 8
- StorageWorks Fibre Channel SAN Switch 16-EL and Fibre Channel SAN Switch 8-EL
- StorageWorks Fibre Channel Arbitrated Loop (FC-AL) Switch
- StorageWorks Fibre Channel Hub 12
- StorageWorks Fibre Channel Hub 7
- Gigabit Interface Converter - Short Wave (GBIC-SW)
- Gigabit Interface Converter - Long Wave (GBIC-LW)

NOTE: For FC-AL configuration using long wave GBICs, total loop length cannot exceed 2800 meters.

Fibre Channel to SCSI Bridge

- Modular Data Router (MDR)
- FCTC II

Tape Library Support

- StorageWorks ESL9326 Enterprise Tape Library
- StorageWorks ESL9198 Enterprise Tape Library
- StorageWorks MSL5026 DLT Library
- StorageWorks TL895 DLT Library
- StorageWorks TL891 DLT MiniLibrary System
- StorageWorks SSL2020 AIT Tape Library

NOTE: Any DLX or SSL2020 Library requires an MDR with LVD SCSI modules.

Configuration Rules & Recommendations

The following configuration rules & recommendations are made based on the solution integration testing conducted by EBS. Certain limitations apply to each ISV and are noted where applicable. For additional EBS installation and configuration information please refer to the Enterprise Backup Solution User Guide documentation at: www.compaq.com/ebs.

Maximum EBS Configurations

The following identifies the maximum number of servers and tape drives supported in a single EBS storage domain/data zone by ISV.

- TSM – 16 Servers x 16 Tape Drives
- Netbackup – 32 Servers x 32 Tape Drives
- NetWorker – 16 Servers for Windows, 32 Servers for UNIX x 32 Tape Drives
- CA – 32 Servers x 32 Tape Drives for Windows 2000
- CA – 20 Servers x 20 Tape Drives for NetWare
- Backup Exec – 32 Servers x 32 Tape Drives (NetWare is limited to 32 Servers x 27 Tape Drives)
- CommVault Galaxy - 16 Servers x 16 Tape Drives

NOTE: For SAN configurations exceeding the maximum supported servers and tape drives, multiple data zones can be implemented.

Zoning

To facilitate SAN management, and minimize server boot times in an EBS configuration, host centric zoning is recommended. Host centric zoning is implemented by creating a specific zone for each server or host, and adding only those storage elements to be utilized by that host. This prevents a server from detecting any other devices on the SAN, to include other servers, and simplifies the device discovery process.

Specific ISV Requirements

- CA Arcserve - Computer Associates ARCserve does not require the use of indexed addressing for device persistency.
- VERITAS BackupExec – No special requirements.
- Veritas NetBackup – Requires professional installation and configuration by authorized Compaq/Veritas personnel.
- Tivoli Storage Manager - Requires professional installation and configuration by authorized Compaq/Tivoli personnel.
- CommVault Galaxy – No special requirements.
- Legato NetWorker - Requires professional installation and configuration by authorized Compaq/Legato personnel.

NOTE: Reserve and release of the tape devices is only utilized by Veritas Backup Exec application.

SCSI Bridge Configuration Rules

Each SCSI bus should be limited to 2 SCSI tape drives.

The tape library SCSI IDs must be set as follows.

Bus 0:

- SCSI ID 1 = Robot (If applicable)
- SCSI ID 2 = Tape drive

- SCSI ID 3 = Tape drive

Bus 1: (and all succeeding SCSI buses)

- SCSI ID 1 = Robot (If applicable)
- SCSI ID 2 = Tape drive
- SCSI ID 3 = Tape drive

IMPORTANT: All Fibre Channel Tape Controller-II SCSI Buses, active or inactive, must be terminated (one terminator per unused bus).

Modular Data Router

The Modular Data Router (MDR) may have a single or dual port Fibre Channel module and up to two Quad SCSI Bus modules with Very High-Density Cable Interconnect (VHDCI) connectors, that can be attached to a maximum of eight tape drives per module. The SCSI Bus modules are available in High Voltage Differential (HVD) and Low Voltage Differential (LVD) configurations. The MDR using modular technology supports up to four SCSI channels per Fibre Channel connection. Each SCSI bus can be connected to a maximum of two tape drives.

Fibre Channel Host to SCSI Target Configuration

The MDR's default configuration allows it to act as a target to a Fibre Channel Initiator and to pass Fibre Channel Protocol (FCP) requests to SCSI target devices. In order to map SCSI targets to Fibre Channel Hosts, the MDR supports two Fibre Channel-to-SCSI addressing methods: Progressive Persistent Device Discovery Addressing and Indexed Addressing.

Progressive Persistent Device Discovery Addressing

Progressive Persistent Device Discovery (PPD) address mapping is the default mode during the new SCSI device discovery process that is initiated upon power-up or remapping. During the MDR discovery process on a SCSI bus, the Index table is filled with adjacent FCP LUNs referencing each subsequent SCSI device. The host system will then detect every attached device without voids, allowing full device discovery to the host. Mapping tables are saved in persistent memory and are loaded each time the MDR is power cycled. The MDR creates new entries in the address mapping tables for newly attached devices. This is the default setting of the MDR.

Device Discovery

The MDR can perform SCSI device discovery in two ways:

- Target ID priority - This mode fills the table according to ascending SCSI Target ID order. This is the MDR's default discovery mode.
- Bus Number priority - This mode fills the table in ascending SCSI Bus Number (SCSI port number) order. Use the Application Management Console (AMC) `setFcLunPriority` command to choose which mode to use (refer to Appendix D, "Application Management Console in Windows HyperTerminal" in the *Compaq Modular Data Router Reference Guide* (part number 133834-001).

Indexed Addressing

Indexed Addressing mode is recommended for environments where users want more flexibility in mapping Fibre Channel to SCSI addresses than provided by the default Progressive Persistent Device Discovery Addressing (PPD) Mode. The user can edit the entries in the Fibre Channel to SCSI Mapping Table with indexed addressing using an in-band management tool. The user can then select a table entry by FCP LUN and specify the associated BUS:TARGET:LUN. The MDR saves all changes to the Mapping table in persistent memory and loads them at the next power cycle. It also allows editing of an address mapping table. SCSI Targets are selected by mapping the appropriate values into the FCP LUN field, and comparing a Fibre Channel LUN value to a SCSI Bus:Target:LUN value. The MDR acts as a single initiator on each SCSI bus, defaulting to ID 7. All commands passed through to a SCSI bus originate from this SCSI ID.

RemapFcSCSI Command after Hardware Configuration Change

When there is a change to the SCSI hardware configuration attached to the SCSI connectors on an MDR, the user must reconfigure the router by executing the REMAPFCSCSI command from the router's Application Management Console (AMC). The AMC is accessed using Hyperterminal (or any suitable serial communications application) through the serial port on the MDR.

NOTE: For instructions on accessing and configuring the MDR, refer to the Compaq StorageWorks Modular Data Router Reference Guide.

NOTE: Tru64 UNIX 4.0F & AIX 4.3.3 supports 8 LUNs per MDR FC link.

IMPORTANT: The Compaq StorageWorks Modular Data Router and FCTC-II are interchangeable in HVD SCSI library configurations. Replacing one type of tape controller with a different model tape controller will necessitate a restart of all servers in the SAN. A restart is required to update all servers' device maps and insure consistency between the operating system and the application.

Selective Storage Presentation

SCSI device targets are presented to a Fibre Channel environment as Fibre Channel Logical Unit Numbers (LUNs). In its default configuration, the MDR presents its Fibre Channel LUN mapping table to all devices on the SAN. Known as an Open Model, a default map (FC_SCSI_MAP) containing all FC LUNs is presented to all initiators. Selective Storage Presentation (SSP) provides the ability to restrict access to a given Fibre Channel LUN. Utilizing a Host Map Administrator List, access to a particular device can be granted or denied based upon that host's World Wide Name (WWN). This functionality is available in MDR firmware revision v1170 or later.

Additionally, Selective Storage Presentation provides the ability to utilize multiple Fibre Channel paths connected to the same Fibre Channel Switch or Fibre Channel Hub. In multi-path configurations without Selective Storage Presentation, the host would issue a REPORT LUNS command to each Fibre Channel port in the MDR. In turn, the MDR would report identical sets of FC LUNs behind each port back to the host, thus causing the host to believe that there were twice as many devices as actually exist. This condition is known as *device ghosting* and is common in multi-path configurations. Selective Storage Presentation allows for the creation of multiple device maps. Each Fibre Channel port or server can then be assigned a unique device map, preventing the ghosting problem and allowing the configuration to maximize the bandwidth of the Fibre Channel ports. This type of configuration is known as a Closed Model because access to a given device is granted to a specific host; there is no default device map.

Large LUN Support

In their default configurations, Windows NT4 and Windows 2000 do not support large LUN values. Specifically, those operating systems do not support LUN values greater than 7 (meaning integers 0, 1, 2, 3, 4, 5, 6, and 7 are valid). In practical terms, this means that without large LUN support enabled, the operating system would not correctly identify all of the tape library devices of a large library (more than eight devices, including robots and tape drives) behind a single MDR.

Windows NT and Windows 2000 bus scan operations are performed in one of two ways. If the windows SCSIPORT.SYS Driver believes a device is capable of supporting a LUN space greater than seven, it will issue a REPORT LUNS command to FC LUN 0. LUN 0 would in turn provide an enumerated list of known devices. In the absence of evidence of large LUN support, SCSIPORT sends inquires to LUNs 0 - 7. Keys in the system registry enable or disable large LUN support.

Compaq has provided a free Windows utility for making the necessary large LUN Registry modifications. In addition to the large LUN Registry entries, this utility inserts a device entry in the Registry for Windows 2000 systems. Without this second entry, each time a Windows 2000 system is started, the New Hardware Discovery Wizard will prompt the user for installation of a device driver. This Large LUN utility will detect the version of the operating system and make the appropriate registry modifications, eliminating the need for the user to make manual edits.

The Large LUN utility is available in two versions: ACS version 8.5 and ACS version 8.6. ACS is the HSG80 Controller firmware found in the StorageWorks Raid Array product line. If the target system's configuration does not include a Raid Array Storage system, then either version of the utility may be used.

NOTE: The Large LUN utility is available for download from www.compaq.com/ebs or can be found on the user guide documentation CD (this CD) located in the Enterprise Backup Solution Software Solution Kit.

Sizing and Performance Tuning

Sizer Tool

The Compaq StorageWorks Enterprise Backup Solution Sizing Tool is a Microsoft Windows-based tool used to determine a backup solution based upon the information supplied by the user. Before beginning the sizing process, the user must have a thorough understanding of the network, the type of data to be backed up, and the backup window parameters.

The StorageWorks Enterprise Backup Solution Sizing Tool:

- Accepts user input.
- Allows the user to select the options that are offered by the ISV.
- Performs calculations for realistic performance.
- Performs the calculations necessary for backup and tape retention.
- Configures a solution for a single data zone or a larger solution or domain.
- Generates both a report and a proposal for a solution.

Download, install, and run the StorageWorks Enterprise Backup Solution Sizing Tool to configure an EBS. The SAN Sizer can be found on the Compaq web site at http://www.compaq.com/ebs

Performance

To analyze speed and performance, it is necessary to examine the entire backup process as a system of components. The backup process can be divided into a set of five components that can affect performance. Each of these components must be thoroughly understood and factored into the backup equation in order to determine the maximal performance in any specific situation. The five components of the EBS are as follows:

- Feed Source

  This is usually the hard disk Primary Storage system, but it can be network-connected storage or even a remote system.

- Storage Connection

  For the EBS, this is Fibre Channel connection.

- File Block Size

  EBS supports up to a 32-KB transfer block size and 64-KB transfer block size for Windows NT 4.0 or Windows 2000.

- File (Data) Compression Ratio

  The amount of compression has a direct impact on the rate at which a DLT tape drive can read/write data.

- Tape Drive (Secondary Storage) System

  For the EBS, these are Compaq StorageWorks Libraries.

Feed Source, Primary Storage, and Controller Type (DLT only)

The type of controller that is used has a direct effect on the speed at which the server can send data to the tape device. Compaq tests show that it is necessary to read from the primary storage device at a speed at least three times the backup rate (3:1) of each DLT drive. This allows the data to stream to the DLT drive, achieving optimal performance results. The base rate for 1:1 backups is approximately 15 GB/h per drive. The base rate for 2:1 backups is approximately 26 GB/h per drive. Therefore, if the controller cannot feed data to the DLT drive at a fast enough rate, the drive performance will slow down due to idle time on the DLT drive.

Storage Connection

The EBS environment is made up of a 100 MB per second (MB/s) fiber storage network that supports simultaneous transfer of many different data protocols, including SCSI, IP, and others. The EBS has been tested to support 200 meter lengths of 62.5 micron multi-mode fiber and 500 meter lengths of 50 micron multi-mode fiber. It also supports Long-Wave 9 micron single mode cable up to 10 km in length.

File Block Size

The use of the largest block size will provide the optimal data transfer rate to a DLT tape drive.

File (Data) Compression Ratio

Compaq tests show that not all data can be compressed equally. The compression ratio will affect the amount of data that can be stored on each tape cartridge, as well as the speed at which the tape drives can read or write the data. As the data compression ratio increases, tape storage capacity increases. For example: At 1:1 compression, a tape can store 35GB of data; at 2:1 compression, it can store 70GB of data (using a DLT 35/70 drive).

Tape Drive Solution

The tape drive solutions is the fifth piece in determining backup and restore performance. Use of the MDR or Compaq Fibre Channel Tape Controller-II and its connections to Compaq StorageWorks tape libraries is a simple way to scale backup performance.

Troubleshooting

The following items have been identified as potential issues, and should be used to assist in troubleshooting problems in an EBS configuration:

Windows NT:

- **Unable to see any SCSI devices that should be available through the LP7/8k HBA.**

  KGPSA cards come pre-configured for arbitrated loop from the factory. Make the following registry entry:

  "Topology=1" to the end of the data string "DriverParameter" located in:

  `HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\`
  `lp6nds35\Parameters\Device`.
  Then reboot the server. Be sure to use the latest driver for HBA.

- Tape devices displayed in wrong order as viewed in the SCSI Adapters applet.

  FCTC to switch port configuration is incorrect per solution installation. Insure that the bridge controlling the robot and first 4 tape drives is plugged in to a lower switch port number than the bridge controlling the last three drives.

- Unable to see devices after changing HBA from AL to fabric and switch port light is flashing.

  FC switch port in bad state. Re-seat/replace GBIC in switch; reset switch port through http: connection; last resort, reset entire switch.

Solaris 7:

- The ov_mount command fails, and the tape is left in the drive unmounted

  According to the DCP output, the test_unit_ready command times out too quickly...the operation errors out There is a libovscsi.so.2 patch file. The customer fix will be a new DCP in the next release of SM (1.5)

Tru64 UNIX:

- Not able to use all drives in the TL895 OS limitation.

  Tru64 UNIX v4.0f does not recognize target ID's above 7 Reset SCSI ID's in the TL895 library according to the Amadeus Functional Spec (This limitation is eliminated in Tru64 UNIX v5.x)

- ASE Tape Service does not failover - NetWorker Tape Service will go into an "Unassigned" status when attempting to failover. (This condition occurs due to a limitation on all current fibre-to-SCSI bridges).

The required ASE Tape Service – for NetWorker Server - will fail to relocate to an alternate node in the event of the primary node's failure.

1. The Tape Service must be restarted.
2. Manual Relocation of the Tape Service. (The Compaq Modular Data Router should correct this problem in the next release).

- No device files exist for TL895 tape devices.

  OS limitation. Tru64 UNIX v4.0F does not automatically create device files for devices with a SCSI ID greater than 7. Follow the instructions in the KGPSA driver information file /etc/emx.info to bind the tape controller WWNs to SCSI IDs from 0 to 7.

Tru64 UNIX and Solaris:

- TL895 tape device files are not in sequential order as ordered from top to bottom in the tape library.

  WWN/SCSI ID bindings are incorrect. The tape controller with the robotic arm should have the smallest SCSI ID. Redo the WWN/SCSI ID bindings in the appropriate order per the driver instructions, remove existing devices (/dev/rmt*h and /dev/nrmt*h for Tru64, /dev/rmt/* and the files they are soft linked to for Solaris), and reboot the server (use the -r option when booting Solaris).

Solaris:

- TL895 tape device files are not consistent across reboots.

  JNI driver limitation. WWNs bind to the first available SCSI ID in the order the devices are encountered. Follow the instructions in the fca man page to bind tape controller WWNs to specific SCSI IDs.

- Secure Path 2.1 will not failover to a certain controller and my host and HSG controller are both set up to spec.

  It is possible for everything to be set up properly and not respond accordingly if the Compaq FC switches do not have the right firmware on them. Read the Solaris Secure Path 2.1 Installation Guide. It specifies f/w 1.6d for the Compaq FC Storage Switch. F/w for the SAN switch needs to be 2.0.3a.

- Disrupting data path may cause panic/core dump on Solaris. Resetting the bridge while Solaris is up.

  Shutdown Solaris, reset bridge, then bring Solaris back up.

SAN Switch 8, SAN Switch 16:

- Unable to run rshd f/w update utility - Returns the following: *** [0] ERROR: Cannot determine port number for the rshd daemon. [0] Winsock error: Error number = 11004. rshd utility does not function properly in dual NIC environment.

  Perform switch f/w update from server with single NIC.

Management Tools

Storage Utility Software Kit

The Storage Utility Software (SUS) Kit is included with all EBS tape libraries. This solution requires SUS version 1.4 or greater. The following CD-ROMs are included in the SUS Kit:

- Tape Drive Supplemental Driver CD
- Compaq SmartStart and Support Software CD
- Tape Storage Management Console (TSMC) CD
- Compaq Management CD.

  The Compaq Management CD contains the following utilities:

  — Compaq Insight Manager

  — Compaq Management Agents for Servers

  — Compaq Management Agents for Clients

  — Compaq Survey Utility

  — Compaq Systems Management Toolkit

  — Compaq Integration Tech Notes

  — Documentation

  — Compaq Power Management

References

Compaq

For comprehensive online support, refer to

http://www.compaq.com

For international information, refer to

http://www.compaq.com/corporate/overview/world_offices.html

For more information about product compatibility, refer to

http://www.compaq.com/ebs

Important Telephone Numbers

Consumer Direct 1-800-888-0220

Compaq DirectPlus 1-800-888-5858 (U.S.)

Compaq Partner Direct 1-800-888-5874

Compaq Reseller Locator 1-800-345-1518 (Option 3)

Compaq Canadian Reseller Locator and Product Literature 1-800-567-1616

Diskette Fulfillment (backup diskettes for pre installed software) 1-800-952-7689 (U.S.)

1-800-349-8498 (Canada)

Compaq Product Information 1-800-345-1518 (U.S.)

1-800-567-1616 (Canada)

Compaq Technical Support 1-800-OK-COMPAQ (U.S. and Canada)

1-800-652-6672

Legato Support

For comprehensive online support, refer to:

www.legato.com/support

For more detailed information on Legato NetWorker™, refer to:

www.legato.com/

Tivoli Support

For comprehensive online support of Tivoli Storage Manager, refer to

http://www.tivoli.com/support

For more detailed information about Tivoli Storage Manager, refer to

http://www.tivoli.com/products

Computer Associates Support

For comprehensive online support, refer to:

support.cai.com/

For more detailed information on Computer Associates ARCserve 2000 or ARCserve*I*T, refer to:

www.cai.com/

VERITAS

For comprehensive online support, refer to:

support.veritas.com/

For more detailed information on VERITAS NetBackup and Backup Exec™, refer to:

www.veritas.com/

6

SAN Management

Overview

With the advent of Storage Area Networks (SANs) and Fibre Channel technology, Compaq is rapidly transitioning from the traditional server, storage and component level-based management to a SAN-level application architecture and implementation using the Compaq SANworks Management Appliance (SWMA).

Just as important as the quality and feature set of the SAN's hardware is the effectiveness of the SAN management applications in tying these devices together and simplifying the complexity of the storage network. Whether using a Compaq standard topology or a custom design using the StorageWorks by Compaq SAN design rules, IT managers need to configure, monitor, and maintain the SAN, as well as plan for, and accommodate, growth.

The Compaq Open SAN management strategy is to:

- Simplify storage management using standardized web-based graphical user interfaces (GUIs) residing on easy-to-use, easy-to-implement management appliances.
- Centralize the management of multi-vendor Heterogeneous Open SANs in distributed and consolidated environments.
- Automate policy-based management.
- Optimize functionality by exploiting all currently available management levels such as appliances, SAN fabrics, and servers/storage.

This chapter describes the currently available StorageWorks by Compaq SAN management tools.

SANworks Management Appliance

Key to the StorageWorks by Compaq SAN management strategy is the use of the SWMA. Compaq SANs can be designed for local, centralized, or distributed data access. Regardless of the arrangement or location of the storage components and preferred data access method, storage environment management can be centralized using a Management Appliance.

The Management Appliance connects directly to the storage network through a Fibre Channel switch providing full access to all supported devices in the storage environment. Strategically located out of the SAN data path, the appliance allows data transfers to proceed independently between computers and storage devices. The appliance optimizes SAN availability and performance while streamlining manageability.

NOTE: For more information about using an SWMA SAN, see Chapter 3, "SANworks Management Appliance Rules and Recommendations".

Appliance Features / Functionality

- Simple, unintrusive management of SAN elements
- High SAN performance since the appliance is located out of the data path
- High SAN availability, since data transfers occur independent of the appliance
- Support for multiple management and monitoring applications
- A web-based, centralized user interface
- No console operations for increased SAN management security
- Support for heterogeneous platforms attached to the SAN
- Higher utilization for processing applications on host servers
- Rack mountable, ease of installation and administration

Open SAN Manager

Open SAN Manager (OSM) is included with and resides on the SANworks Management Appliance, giving you access to the management appliance. Logging into OSM anywhere over the web provides a single aggregation point to launch a variety of Compaq's SAN management applications to monitor and manage your storage network.

Using the SANworks Management Appliance in a Heterogeneous Server Environment

Whenever a management appliance is placed in a fabric with heterogeneous servers, a dedicated storage management zone must be created. This zone is specifically for the management appliance and the elements it is to monitor and manage.

For example, create a zone called SANAPP_#_ZONE that would contain the appliance host bus adapter WWID and the WWIDs of all the HSG or HSV controllers managed by this SWMA. Because fabric devices can be in multiple zones, this will have no effect on other zones containing the same HSG and HSV controller WWIDs.

Currently, the management appliance communicates with HSG or HSV controllers in-band, within the Fibre Channel fabric itself. It is not necessary or recommended to include either the switch WWIDs or server HBA WWIDs in this zone. Communications to these devices are done out-of-band; outside the fabric via TCP/IP.

SAN Management Categories

SAN management is wide ranging, covering many aspects of the day-to-day activities used for monitoring and managing, as well as simplifying, the complexity of the storage network.

This section classifies SAN management into four major categories:

- Fabric management
- Storage management
- Data management
- SAN usage and monitoring

SAN Fabric Management

SAN Fabric management can be thought of as the control of the SAN infrastructure or "traffic flow" within the SAN. This pertains to control and management of device communication or access within the SAN, such as switch zoning, or LUN level Selective Storage Presentation (SSP). This also includes managing SAN interconnect components, individually and collectively, throughout the fabric.

SAN Storage Management

Storage management allows control of the specific storage system configuration such as redundant paths, creation and management of storage sets (LUNS), setting of RAID levels, and the setting of platform specific SAN interface characteristics and parameters.

SAN Data Management

SAN data management applications help ensure that data is available and accessible. The data being stored on the SAN is part of a company's assets. It is imperative to keep this data available to system applications with minimal, if any, downtime. Techniques such as cloning, snapshots, data replication, and backups protect the data from disasters.

SAN/Storage Usage & Monitoring

SAN and storage usage and monitoring applications are necessary to provide SAN event notification and fault/failure information for service before SAN anomalies can adversely impact the enterprise. They may also provide reporting and billing information for determining the amount of storage and quality of service delivered.

SAN Management Application Deployment

Within the different categories of management tools, individual tools are implemented either on the management appliance, within fabric interconnect components, or within servers/storage systems. Table 6–1 lists the management tools by category, and identifies where the specific tools reside.

NOTE: Some applications may reside in more than one category.

Table 6–1: SAN Management Tools & Location

SAN Management Application	OSM Based	Fabric Based	Server Based	Storage Based
SAN Fabric Management				
StorageWorks Fabric Watch	No	Yes	No	No
SAN/Fibre Channel Switch Management	No	Yes	No	No
SANworks Network View	Yes	No	No	No
SAN Storage Management				
SANworks Element Manager for HSG	Yes	No	No	No
SANworks Element Manager for HSV	Yes	No	No	No
SANworks Network View	Yes	No	No	No
StorageWorks Command Console	No	No	Yes	No
Storage System Array Controller Software (ACS) Command Line Interface (CLI)	No	No	No	Yes
Storage System Scripting Utility (SSSU)	No	No	Yes	No[1]
RA4000/4100 Array Configuration Utility (ACU)	No	No	Yes	Yes
MSA1000 (ACU, ACU-XE, ACU-XE(Offline)	No	No	Yes	Yes
SANworks Secure Path	No	No	Yes	No
SAN Data Management				
SANworks Enterprise Volume Manager (EVM)	No	No	Yes	No
SANworks Virtual Replicator	No	No	Yes	No
SANworks Data Replication Manager (DRM)	No	No	No	Yes[2]
SANworks Command Scripter	No	No	Yes	No[3]
SAN/Storage Usage & Monitoring				
SANworks Network View	Yes	No	No	No
SANworks Open SAN Manager (OSM)	Yes	No	No	No
SANworks Storage Resource Manager (SRM)	No	Yes	Yes	No
SANworks Storage Resource Manager for Exchange	No	Yes	Yes	No
SANworks Resource Monitor	Yes	No	No	No
SANworks Storage Allocation Reporter	Yes	No	No	No

[1] This product is a character cell interface to configure and control an Enterprise Virtual Array

[2] DRM requires ACS Version 8.xP Software

[3] This product is a front-end to the Storage System CLI

SAN Fabric Management Tools

SANworks Network View

SANworks Network View is a SANworks Management Appliance application providing at a glance views of SAN configuration and availability. SAN devices, their fiber channel interconnects, and associated status are automatically discovered and represented in an intuitive topographical display. SAN device management is made easy by double clicking on a device icon. A Java based design allows remote SAN management from any web-enabled console having Internet Explorer browsing capability. Network View serves as a SAN management consolidation point.

Software Features / Functionality

- Simplified SAN management from one application
- At a glance SAN visualization
- SAN administration from remote locations
- Automated SAN availability monitoring and notification when faults arise
- A consolidation point and launch pad for device specific device and storage management tools
- Fibre link connection mapping for established SANs or for future planning
- Scalable for future SAN growth
- A host independent solution

SANworks Network View spans three SAN management categories by providing:

- SAN Fabric Management - Network View can view, monitor and manage FC Switches, Tape Routers as well as Inter Switch Links (ISL) right from the topology map. By either clicking on the device icon or the device folder Network View will automatically launch the device's web GUI.
- SAN Storage Management - HSG elements can also be viewed, monitored and managed from the Network View topology map. By clicking on the element icon or device entry, Network View will call up the respective Element Manager application.
- SAN monitoring - Network View can monitor the condition of the fabric hardware by displaying and reporting the condition of server HBAs, Fabric Interconnects, and HSG elements via E-mail, pager or SNMP traps. Performance can be monitored either in real time or a SAN history may be maintained to playback at anytime.

Network View Setup in a Large SAN

Network View discovers, monitors and manages various FC devices either through in-band or out-of-band communications with that device.

HSG Elements

Network View uses in-band communication to discover the HSG controllers or elements. Network View will automatically populate its database and topology map based on the HSG elements discovered by the appliance. Currently, HSG elements are displayed by the controller serial numbers discovered

It is recommended to change the properties of the HSG Element to a more intuitive name for display and error reporting reasons. Right-click on the element icon or device list to change properties. A suggestion would be: *location-controller type-fail over mode- ACS version*

For example: RACK05TOP-G80-T-V8.6F would indicate the element resides in the top of Rack 5, and are HSG80 controllers running in Transparent Failover with ACS code V8.6F.

Fibre Channel Switches/ Fibre Channel Tape Controllers

Network View uses out-of-band communication (TCP/IP) to discover Fibre Channel Switches and Fibre Channel Tape Controllers. When Network View is launched for the first time a Configuration pop-up window will appear indicating the database is empty. You may then add a range of FC Switch/Tape Controller IP addresses for Network View to map and monitor.

You may also at any time add additional IP addresses by clicking on the Configure button on the topology map, adding the new IP addresses then Start discovery.

By default, Network View will display the DNS name of the IP address, if any, or the IP address itself. Right-clicking on the device icon or name will allow you to edit the name within the properties box that is displayed in the topology map.

A suggestion would be: *FCtopology-device-xx*

For example: RING-SWITCH-01 would indicate the first FC switch in the RING topology.

NOTE: Use at least 2 characters for numbers to keep the display sorted properly.

Server Host Bus Adapters

Network View does not initially discover server HBAs until a Device Manager Agent is installed on the server. During the agent install it will prompt you to input the appliance name running Network View. It is this device manager service that "pushes" the HBA information over TCP/IP to the appliance. Currently, server device managers are available on Window NT, Windows 2000, Sun Solaris and Tru64 UNIX.

Network View will display the DNS name of the server in the topology map and it cannot be renamed. However it is suggested to append the HBA name located under the Host name for monitoring and error reporting purposes.

A suggestion would be to prefix the existing entry with: *servername-topology*

For Example: SERVER04-RING-Emulex-LP8000-Port0 would indicate this is SERVER04's RING topology adapter.

For further information, including agent O/S versions, please read the Network View QuickSpec and release notes.

SANworks Fabric Watch

Fabric Watch allows the SAN manager to monitor key fabric and switch elements, making it easy to quickly identify and escalate potential problems. It monitors each element for out-of-boundary values or counters and provides notification when any exceed the defined boundaries. The SAN manager can configure which elements, such as error, status, and performance counters within a Compaq SAN Switch, are monitored.

Fabric Watch can be accessed through a web GUI, a telnet interface, an SNMP-based enterprise manager, or by modifying and uploading the Fabric Watch configuration file to the switch.

Fabric Watch monitors the following elements:

- Fabric events (such as topology reconfigurations, zone changes)
- Switch environment (fans, power supplies, and temperature)
- Ports (state changes, errors, and performance)
- GBICs

With Fabric Watch, each switch continuously monitors error and performance counters against a set of defined ranges. This and other information specific to each monitored element is made available by Fabric Watch for viewing and, in some cases, modification. This set of information about each element is called a *threshold*, and the upper and lower limits of the defined ranges are called *boundaries*. If conditions break out of acceptable ranges, an *event* is considered to have occurred, and one or more alarms (reporting mechanisms) are generated if configured for the relevant threshold.

SAN/Fibre Channel Switch Management

The Compaq Fibre Channel SAN Switches are high performance, scalable switch fabrics designed for creating large SANs. The management functions let you control and monitor fabric topology, frame throughput, error statistics, fans, cooling, media type, port status, and a variety of other information to aid in system debugging and performance analysis.

The administrative and diagnostic functions of the SAN switch are accessible from IP over the RJ-45 10/100BaseT Ethernet port or any Fibre Channel port. You can use any Simple Network Management Protocol (SNMP)-based management product to access the SNMP agent. You can also use any supported web browser to use the Java Web Management Tools.

Supported management methods include:

- SNMP
- Telnet
- Web-based Management Tools launched via StorageWorks Command Console or Network View
- Telnet command subset via switch front panel display (Fibre Channel SAN Switch/16 only)

SAN Storage Management Tools

Element Manager for HSG

Element Manager for HSG is a SANworks Management Appliance application to configure and monitor HSG80/60 controllers. For each controller pair, Element Manager for HSG enables you to:

- View existing virtual disk, controller, physical disk, and host properties
- Make changes to these properties for different configurations
- Configure Remote Copy sets and add associations (with supported firmware)
- Dynamically expand volumes for operating systems that support dynamic volume expansion
- Make temporary snapshots of volumes for backup purposes (with supported firmware)

HSG Element Manager Restrictions

Current restrictions of the HSV Element Manager must be enforced:

- Maximum of 25 HSG Storage Systems can be managed by one SWMA
- An HSG Storage System must not be visible to more than one SWMA

Management Appliance and HSG storage system Communication

When configuring an HSG controller in a Management Appliance storage environment, you will need to enable CCL (i.e. setting the controllers to SCSI-3 or SCSI-2 with CCL Enabled) or provide a dedicated LUN (i.e. setting the controllers to SCSI-2 CCL Disabled) for the SANworks Management Appliance. If you are using a dedicated LUN instead of CCL, verify that the LUN is presented to the Management Appliance through each controller host port connection for the Management Appliance (there are two host ports per controller). The LUN may be a partition. For more information, see your HSG controller user guide.

General HSG Storage System Configuration Process

The following steps highlight the configuration process for storage systems. Refer to your storage system user guide for more information.

1. Set up the storage system according to the product user guide.
2. Ensure that you have connected the Fibre Channel from HSG controller ports to optical interfaces found on the Fabric Switch.
3. Ensure that the appliance and the HSG controllers WWIDs are in a SAN appliance zone.

NOTE: The following steps are only necessary if the HSG controller is configured for SCSI-2 CCL Disabled.

   a. Connect to the HSG controller via the serial interface. Refer to your HSG controller manual for further information on the serial connection.

   b. Start terminal session. Refer to your HSG controller manual.

   c. At the prompt in the terminal session, enter the Command Line Interface (CLI) command SHOW CONNECTION.

   d. Verify, via HSG connection table, that the Fibre Channel HBA of the Management Appliance is online. Use CLI command SHOW CONNECTION.

   e. Create Logical Unit Number (LUN) and enable access from the LUN to the Management Appliance.

4. Verify that the HSG controllers are discovered. To verify HSG status, you will need to configure and launch Element Manager for HSG and click on OPTIONS. Enable controllers that are discovered.

Currently, HSG elements are displayed by the controller serial numbers discovered. It is recommended to change the properties of the HSG Element to a more intuitive name for display and error reporting reasons.

1. In the navigation pane click on a controller serial number displayed.
2. In the Content Pane edit the ALIAS field and save changes.

A suggestion would be: *location-controller type-fail over mode- ACS version*

For example: RACK05TOP-G80-T-V8.6F would indicate the element resides in the top of Rack 5, and are HSG80 controllers running in Transparent Failover with ACS code V8.6F.

HSG Storage System Array Controller Software /Command Line Interpreter

HSG Array Controller Software (ACS) for Fibre Channel Arbitrated Loop and Switched Fabrics provides storage controller software capability for the StorageWorks HSG60 and HSG80 Array Controllers in Fibre Channel arbitrated loop and switched fabric environments. HSG Array Controller Software is designed to be common across multiple operating system platforms. However, there may be operational differences between platforms, and there may also be features that are not supported on every platform.

Management of storage systems based on the HSG60 or HSG80 is provided directly through the controller serial port using a terminal or a terminal emulator (such as Microsoft Windows NT HyperTerminal) using the CLI interface. The CLI provides all the commands necessary to configure controller fail over modes and parameter settings, controller and host connections to the SAN, storage set creation, SAN LUN access (SSP), RAID levels, and cache settings. The CLI also provides access to the array controller utilities. The utilities are used to monitor controller functions and statistics, and to allow storage system component replacement procedures to be conducted while the storage system is active.

Selective Storage Presentation

Selective Storage Presentation (SSP) provides a way to control SAN access at the storage set or LUN level. SSP allows each server or HBA's storage sets (LUNs) to be presented exclusively to those that are allowed access. Additionally, SSP allows the setting of host modes and LUN offsets for each HBA connected to the storage system. The host mode is specially tailored to the storage communication techniques of the operating system type. The LUN offset feature of SSP allows higher numbered LUNs in a storage array to be presented in a range required by specific operating systems. The SSP feature also provides a way to track the numerous Fibre Channel HBAs within servers attached to a SAN by identifying each by name and WWN.

ACS Features / Functionality

- Host Interconnect and Protocol Services
- Microsoft Cluster Server (MSCS) Support
- Dual Redundant Controller Operation
- Testing and diagnosis of the HSG array controller
- SCSI device control
- Transparent Controller Failover Support
- Multiple-Bus Failover Support
- Asynchronous Disk Swap (Hot Swap)
- ACS system management services
- Local program support
- Mirrored Write-Back Cache support
- Read Ahead Cache support
- Disk Mirroring capability (RAID 1)
- Disk Striping capability (RAID 0, 0+1)

- RAID capability (RAID 3/5)
- Storage set Expansion
- Disk Partitioning capability

Supported management methods include:

- Terminal emulation through the HSG's serial port using the CLI
- SANworks Command Console
- SANworks Command Scripter

ACS works in a heterogeneous environment that includes Tru64 UNIX, Compaq OpenVMS, Microsoft Windows NT and Windows 2000, Novell NetWare, Sun Solaris, HP-UX, SGI IRIX, IBM AIX, Linux x86, and Linux Alpha. This application is at the storage system level.

Element Manager for HSV

Element Manager for HSV is a SAN management application to configure and monitor HSV110 controllers. For each controller pair, Element Manager for HSV enables you to:

- Initialize an Enterprise Virtual Array and create a pool of disks drive
- View, configure and upload code to the controllers and disk drives
- View and configure virtual disks, and host properties
- Dynamically expand volumes for operating systems that support dynamic volume expansion
- Make temporary snapshots of volumes for backup purposes (with supported firmware, requires license)
- View Enterprise Virtual Array event logs

VCS Features and Functionality

- Support for up to 240 disk drives per controller pair
- Management of up to 256 virtual disks per disk pool ranging in size from 1GB to 2TB per virtual disk
- Dynamic capacity expansion and virtual disk data load leveling
- Distributed sparing of disk capacity
- Virtually Capacity-Free Snapshot (Vsnap)
- Virtually Instantaneous Snapclone
- Dual redundant controller operation for increased fault tolerance
- Multiple Bus Failover Support
- Battery Back-up
- Asynchronous Disk Swap (Hot Swap)
- Clustered Server Support
- Mirrored Write-Back Cache Support
- Read-Ahead and Adaptive Read Caching Support
- Virtual RAID Arrays (Vraid0, Vraid1, Vraid5)

- Non-disruptive software upgrade capability
- Initially supports connection of up to 256 hosts
- Multi-Vendor Platform Support
- Controller Password Protection for Configuration Control
- Selective Storage Presentation and SAN-based data zoning
- GUI Interface for management and monitoring

Supported management methods include:

- SANworks SSSU
- SANworks Management Appliance

VCS works in a heterogeneous environment that includes Tru64 UNIX, Compaq OpenVMS, Microsoft Windows NT and Windows 2000, and Sun Solaris. This application is at the storage system level.

HSV Element Manager Restrictions

Current restrictions of the HSV Element Manager must be enforced:

- A maximum of 16 Enterprise Storage Systems can be managed by one SWMA
- An Enterprise Storage System must not be visible to more than one SWMA
- Use of more than one browser session or one browser session and an SSSU session at the same time to a single SWMA is not supported

General HSV Storage System Configuration Process

The following steps highlight the configuration process for storage systems. Refer to your storage system user guide for more information.

1. Set up the storage system according to the product user guide.
2. Ensure that you have connected the Fibre Channel from HSV controller ports to optical interfaces found on the Fabric Switch.
3. Ensure that the appliance and the HSV controllers WWIDs are in an SAN appliance zone.

Initially, Enterprise Storage Systems are display as "UNINITIALIZED" on the HSV Element Manager browser window. It is recommended that, when the storage system is initialized, a intuitive name is used for display and monitoring convenience.

1. In the navigation pane, click on a controller icon.
2. In the Content Pane, click the INITIALIZE button.

   A pop-up displays that confirms you are initializing the storage system. It also states that any data associated with the selected system will be lost, and then asks if you wish to proceed with the initialization procedure.

   If you have not previously entered the license key for the storage system, you will be prompted to do so.
3. Enter a name for the storage system.

   A suggestion would be:

   location-controller type- VCS version -

For example: RACK05-V110-V1.0 would indicate the element resides in Rack 5, and contains HSV110 controllers running VCS code V1.0.

4. Specify the number of disks in the default group. Enter from 8 up to the total number of drives in the subsystem.
5. Click on the **Advanced Options** button and set the date and time option. It is recommended that you synchronize the time with the SANworks Management Appliance time. If this practice is used with all controllers, then you will have synchronized times on all event log entries.
6. Leave the Console LUN ID set to "**0**"

StorageWorks Command Console

StorageWorks Command Console (SWCC) is a feature-rich, graphical user interface providing local and remote management of StorageWorks HSG60 and HSG80 array controllers. It is a user-friendly tool for monitoring, configuring, and troubleshooting Compaq storage arrays and controllers.

SWCC can be connected to your StorageWorks controller in several ways. Once connected, the program issues commands and interprets the responses sent by the controller. The user interface displays the logical and physical layout and status of a selected subsystem in graphical form. Command Console consists of two major components: the Client and the Agent. The Client, which includes the user interface and some additional services, provides a window into your storage subsystems. The Agent is a host-resident program that is an interface between the Client and the host's storage controller to interpret and transfer information.

The Agent acts as the Client's assistant in controlling your storage subsystem. The Agent continuously monitors the subsystem and notifies the Client of changes. Commands sent from the Client are received by the Agent and are routed to the storage subsystem via the subsystem's Fibre Channel bus. Subsystem status is transmitted back to the Client from the Agent via the network connection.

Software Features / Functionality

- Easy, graphical configuration of the storage subsystem using an interface similar to Windows Explorer.
- Graphical view of the controller and its physical and logical storage elements.
- Status monitoring of the storage subsystem using intuitive icons.
- Fault notification by pager, electronic mail, and event log entries.
- Management of multiple host systems through a TCP/IP network connection.
- Direct serial port connection.
- Direct SCSI port connection (Windows NT and Windows 2000 Only).
- Robust security that prevents unauthorized access to configuration capabilities.
- The Client supports Microsoft Windows NT 4.0 and Windows 2000.
- The Agent supports Tru64 UNIX, Compaq OpenVMS, and multi-vendor platforms.
- This application is at the server level for both the Client and the Agent.

Array Configuration Utility for RA4000/4100/MSA1000

The Compaq Array Configuration Utility (ACU) software (for Smart Array products, StorageWorks RAID Array 4100/4000 systems, and MSA1000 systems) makes it easy to configure and expand your disk drive arrays. This graphical tool is very intuitive: by using its Configuration Wizards, you have the ability to configure your array controller, add additional disk drives to an existing configuration, or completely reconfigure your disk drive array.

Software Features/ Functionality

- Selective Storage Presentation: allows RA4100 and MSA1000 array sets to be partitioned to multiple servers for SAN access
- Online RAID Level Migration: allows for online post-configuration change to RAID level without destroying data or volume information.
- Online Capacity Expansion: lets you add storage to an operational RA4100 or MSA1000, reducing expensive server downtime.
- Online Volume Extension: allows for the capacity growth of existing logical volumes.
- Global Online Spare: reduces the risk of data loss by facilitating automatic rebuilds after a drive failure.
- Logical Drive Capacity Extension: allows the user to increase the size of existing logical drives online under Windows NT and offline for other operating systems.
- Pre-Failure Warranty: Drives installed in an RA410 or an MSA1000 and monitored under Compaq Insight Manager are supported by a Pre-Failure (replacement) Warranty.

NOTE: Pre-Failure Warranty allows for the replacement of designated drives in an RA4100 before they actually fail when using Compaq Insight Manger on Compaq servers.

NOTE: Some operating systems may not support all of these features.

SANworks Secure Path Multi-Path Software

Depending on the platform or operating system, high availability functionality may or may not be embedded in the operating system I/O drivers. Tru64 UNIX V5.0A/V5.1 and OpenVMS operating systems have the ability to create and maintain multiple paths over the SAN to the same LUN, with support for these functions embedded. For those operating systems that do not support multi-pathing, Compaq provides this capability using Compaq SANworks Secure Path.

The Compaq SANworks Secure Path provides continuous data access for RAID storage systems accessed by operating systems that are not Compaq-based. When combined with the inherent fault-tolerant features of the RAID Array, this configuration effectively eliminates single points of failure in the storage system.

When a host bus adapter, cable, or controller in a path fails, the failure is detected and I/O is automatically re-routed to the functioning, alternate path. This process, called fail over, requires no resource downtime and ensures high availability of data. Storage units that have experienced fail over may be configured to failback automatically after a path is restored. Failback can also be done manually through the use of the Secure Path GUI.

Software Features / Functionality

- Switched fabric and loop support
- Automatic path fail over
- I/O load distribution
- User-selectable failback
- Supported on the SANworks Management Appliance

SANworks Secure Path works in a heterogeneous environment. See Chapter 4, "Heterogeneous SAN Platform and Storage System Rules".

Secure Path Element Manager on the SANworks Management Appliance

Managing Windows NT and Windows 2000 Secure Path servers throughout the enterprise is now available using the SANworks Management Appliance. Secure Path Element Manager uses the same easy-to-use SANworks Management Appliance Web GUI interface to manage and monitor hosts and HSG60/80 storage subsystems as well as integrates with the SWMA's notification utility.

The notification utility provides centralized functionality. The web-based Notification console is used to provide a single, modular, networked software unit that has the ability to handle Event Logging, SMTP, SNMP and command line launching operations.

Secure Path Element Manager uses TCP/IP to communicate with Secure Path servers. Adding the SANworks Management Appliance server name to the Client list on the Secure Path Server will allow Secure Path Element Manager to discover the Secure Path server and add it to a profile.

SAN Data Management Tools

SANworks Enterprise Volume Manager

Compaq SANworks Enterprise Volume Manager (EVM) is web-based application software that manages controller-based clone and snapshot operations. Cloning is a mirroring copy function that allows you to create an exact copy of a LUN; snapshot provides an instantaneous point-in-time copy function. Both that can be used to minimize downtime required for system backups and data migration activities. EVM can be used to meet business continuance requirements by minimizing application downtime required for system backups and data migration activities. EVM automates the creation of command files that control the cloning or snapshot operation. EVM also allows you to mount the clone or snapshot on a second host on the same controller.

EVM is a host-based tool that can be accessed by the user directly via a GUI or remotely via a Browser. The tool then interacts with the requested application(s) to stop new I/O and flush pending I/O to disk. Once the I/O is stable, EVM instructs the storage via the in-band CLI to perform a snap or clone operation. Finally, on completion of that operation, EVM restarts the application.

EVM automates the creation of command files that control the cloning or snapshot operations. EVM also allows users to mount the clone or snapshot to a new host. The new host can then act as a dedicated backup server or data warehouse server. All operations are performed on the clone or snapshot, minimizing performance impact on the production system.

Software Features / Functionality

- Web-based application
- Supported on the SANworks Management Appliance
- Easy management of complex cloning and snapshot operations
- Supports LAN-less backup
- Simplified, centralized storage management

Enterprise Volume Manager works in a heterogeneous host environment that includes Tru64 UNIX, Microsoft Windows NT, Windows 2000, and Sun Solaris. This application is at the server level.

Enterprise Volume Manager on the SANworks Management Appliance

Managing Windows NT and Windows 2000 Enterprise Volume Manager (EVM) servers throughout the enterprise is now available using the SANworks Management Appliance.

Enterprise Volume Manager uses TCP/IP to communicate with EVM servers. Adding the SANworks Management Appliance server name during the EVM server agent setup will allow the EVM application on the SWMA to discover the EVM server.

SANworks Virtual Replicator

The Compaq SANworks Virtual Replicator (VR) combines a rich set of innovative capabilities that enhances and simplifies storage management for Microsoft Windows NT and Windows 2000 environments. Through virtualization, online volume growth, snapshot and management features, the software complements the standard capabilities within the operating system.

SANworks Virtual Replicator utilizes industry-standard server, storage, and network-interconnect components, protecting an organization's current and future storage investments.

SANworks Virtual Replicator 2.0 provides the ability to create instant, virtual snapshots of production data-without having to physically copy it. A snapshot, which looks exactly like the original disk from which it was copied, takes seconds to create and allows customers to backup and restore data with minimal impact to users and applications. Customers can schedule automated snapshot backups using the integrated policy-based scheduling and scripting features.

Software Features / Functionality

- Virtualization:

  Allows companies to respond quickly to rapidly changing storage capacity requirements. With storage virtualization, multiple storage arrays can be grouped into a pool of disk space for individual or clustered systems to use. Multiple virtual disks, up to 1 terabyte in size, can be created from a pool for users and their applications. System administrators can tailor disk space to specific requirements.

- Online volume growth:

  Enables easy, non-disruptive growth for Windows 2000 with zero downtime. Online Volume Growth allows a system administrator to grow an existing volume on a SANworks Virtual Replicator virtual disk and also on a Windows 2000 basic disk. The system will remain online, and the data on the volume will remain intact.

- Snapshots:

Enable the instant creation of multipurpose virtual replicas of production data without the requirement of a physical copy. Snapshots function identically to ordinary physical disks with both read and write capability. Whenever a quick copy of production data is needed, snapshots can be used with minimal disruption to running applications. For example, the snapshot can be the source for backup using standard backup tools. Snapshots can remain online for restore operations, testing, and data mining.

- Management:

  Simplification through easy-to-use interfaces using Microsoft Management Console or a command line. Interactive wizards are available to guide the administrator through all management tasks and create automatic schedules of operations.

SANworks Virtual Replicator is supported on Microsoft Windows NT, Windows 2000, Professional, Server, and Advanced Server. This application resides at the server level.

SANworks Data Replication Manager

The HSG80 Data Replication Manager (DRM) Software is the software component of the HSG80 array controller used in switched fabric environments for remote data replication. Data Replication Manager Software is a storage-based disaster tolerance and workload migration solution that provides the ability to copy data, in real time, to a remote location, up to 100 km away using direct Fibre Channel or further using ATM links. This is done without any host involvement. The HSG80's dual host port design, when used in DRM configurations, allows for long distance mirroring in a switched No Single Point of Failure (NSPOF) Fibre Channel topology.

The DRM software executes in the HSG80 array controller and processes I/O requests from the hosts, performing the local and remote device-level operations required to satisfy the requests. This is done through the use of a pair of initiator and target controllers sharing a switched NSPOF Fibre Channel fabric. Host generated Reads are performed on the local copy of the data. Host generated Writes to the local storage are copied by DRM from the local controller directly to the remote controller automatically. This capability provides the ability to maintain the same data at both locations, providing disaster tolerance protection.

Software Features / Functionality

- Online, real-time data replication to a local or remote site
- Data replication over a Fibre Channel SAN
- Cloning at Initiator and Target sites
- Snapshot support at Target site
- Cascaded switches support
- Full Fibre Channel-to-ATM connectivity with line speeds of T1 through OC3
- Replicate up to 100 km(~63 miles) with Very Long Distance GBIC
- Asynchronous and synchronous transfer modes
- Write History Logging and "Mini-Merge" reconstruction
- Stretched Clusters capabilities for Microsoft Windows NT and Compaq OpenVMS
- Association sets
- Non-RCS LUN support
- Switch Zoning support
- Wave Divisional Multiplexing

Data Replication Manager works in a heterogeneous host environment that includes Compaq OpenVMS, Compaq Tru64 UNIX, HPUX, IBM AIX, Microsoft Windows NT, Windows 2000, Novell NetWare, and Sun Solaris. The application is at the storage system level.

SANworks Command Scripter

Compaq SANworks Command Scripter is application software that provides command-level control of Compaq StorageWorks systems equipped with HSG60, HSG80, HSZ70, and HSZ80 Array Controllers. With Command Scripter, you can create, edit, and run script files that contain StorageWorks Command Line Interpreter (CLI) commands. This allows automation of frequently performed StorageWorks operations.

Two interfaces are included in Command Scripter: a command line interface for local, direct connection to StorageWorks controllers and a web-based interface, which requires StorageWorks Command Console (SWCC)) for centralized, remote connection via browser.

Software Features / Functionality

- Web-based interface for centralized, remote connection to StorageWorks array controllers
- Command line interface for local, direct connection to array controllers
- Select agent host and StorageWorks subsystem
- Create and edit CLI script files
- Run saved CLI script files
- Execute a single CLI command
- Display CLI command history

Command Scripter works in a heterogeneous host environment that includes Tru64 UNIX and OpenVMS, Microsoft Windows NT, Windows 2000, and Sun Solaris. This application is at the server level.

Storage System Scripting Utility

SSSU is the character cell interface for a user. Host based application that needs to access the HSV element manager should use the EMClientAPI. That API will transport SOAP/XML requests over the wire to the element manager, handling security and communication. The EMClientAPI provides an efficient machine interface to the HSV element manager, specifically designed for host-based applications.

SAN Storage Usage & Monitoring Tools

SANworks Automation Manager

SanWorks Automation Manager provides a tool from which a storage administrator can automate the management of a storage area network. SanWorks Automation Manager runs, controls, and manages predefined policies the storage administrators can configure for their environment. Predefined policies are provided with the product as Perl scripts. In addition, you can create and import your own management scripts.

Automation Manager also provides the following utilities to assist in managing storage operations:

Reports – View and print status reports about storage operations.

Agents – View and download an agent to hosts on which scripts resides. Agents enable Automation Manager to communicate with and run batch jobs on hosts systems.

Notification – Set up different notification types for Automation Manager events.The notification utility provides centralized functionality. The web-based Notification console is used to provide a single, modular, networked software unit that has the ability to handle Event Logging, SMTP, SNMP and command line launching operations.

SANworks Resource Monitor

Residing on the SANworks Management Appliance, SANworks Resource Monitor provides continuous and accurate event notification for StorageWorks switches and arrays to reduce mean time-to-repair and increase overall SAN availability. Customizable remote notification options include e-mail, alphanumeric page and SNMP traps. Authorized users can access event information anywhere, anytime via the intuitive Web-based user interface.

SANworks Resource Monitor allows the user to monitor, identify, and address SAN anomalies and events before they can adversely impact the enterprise. This application delivers reliable, continuous monitoring and event notification for supported devices, reducing overall storage management costs while increasing SAN uptime.

Software Features / Functionality

- Scalable Host-Independent Solution
- Preemptive Broadcast Notification.
- Multiple notification options
- Web-based user interface.
- SAN-scale monitoring of Compaq Storage Arrays and switches.
- SANworks Resource Monitor is a host-independent solution.
- This application is at the SAN Appliance level.

Event Notification Guidelines - Resource Monitor and Element Manager for HSG

Resource Monitor and Element Manager for HSG log separate occurrences of an event based on these attributes:

- Specific device that generated the event
- Event Name that is displayed in the Event Log
- Message that is displayed in the Event Log

NOTE: Notification setup is linked to both Resource Monitor and Element Manager for HSG.

Changes in the settings affect both applications.

In Resource Monitor, event notifications are displayed in the *Device Event Log*. In Element Manager for HSG, event notifications are displayed in the *HSG Event Log*.

NOTE: If you would like to view events for a particular subsystem, access that information through Resource Monitor.

For the first occurrence of an event, you will receive the appropriate notification as configured in the Resource Monitor and Element Manager for HSG event notification screens. For subsequent occurrences of an event, you will receive notification based on the following guidelines:

- You will receive an event notification for the notification severity level plus all levels above. For example, if you set event notifications at the Information level, you will receive event notifications for Information, Warning, and Failure.
- If the occurrence was *acknowledged* through the event log view, you will receive one (1) notification per hour and only if one of the remaining conditions applies to the occurrence (see rules below):
  - — **Failure** - severity events will be sent every occurrence.
  - — **Warning** - severity events will be sent every 5th occurrence.
  - — **Information** - severity events will be sent every 20th occurrence.

NOTE: All occurrences of an event display in the Resource Monitor and Element Manager for HSG event logs.

Setting Event Notification on a Fibre Channel Switch

Although Resource Monitor automatically detects HSG controllers for event notification, Compaq Fibre Channel switches must be configured to 'push' SNMP events to the appliance. These settings can be configured either through the switch GUI or a switch telnet session. Using a telnet session the command would be *agtcfgset*. You will then be prompted step-by-step to make any changes.

For example, a typical Telnet session to enable switch events to be sent to an appliance with the IP address of 16.117.74.80 would be:

```
Switch admin> agtcfgset
Customizing MIB-II system variables ...
At each prompt, do one of the followings:
    <Return> to accept current value,
    enter the appropriate new value,
    <Control-D> to skip the rest of configuration, or
    <Control-C> to cancel any change.
```

To correct any input mistake:

```
<Backspace> erases the previous character
<Control-U> erases the whole line,
sysDescr: [Fibre Channel Switch.]
sysLocation: [End User Premise]
sysContact: [Field Support.]
swEventTrapLevel: (0…5) [0] 4
authTrapsEnabled (true, t, false, f): [false] T
 SNMP community and trap recipient configuration:
Community (rw): [Secret C0de]
Trap Recipient's IP address in dot notation: [0.0.0.0]
Community (rw): [OrigEquipMfr]
Trap Recipient's IP address in dot notation: [0.0.0.0]
Community (rw): [private]
Trap Recipient's IP address in dot notation: [0.0.0.0]
Community (ro): [public]
Trap Recipient's IP address in dot notation: [0.0.0.0] 16.117.74.80
Community (ro): [common]
Trap Recipient's IP address in dot notation: [0.0.0.0]
Community (ro): [FibreChannel]
Trap Recipient's IP address in dot notation: [0.0.0.0]
Committing configuration...done.
```

NOTE: Resource Monitor will not initially see the FC switch until the switch generates an event. An event can be forced by either removing a cable or GBIC.

SANworks Storage Resource Manager (SRM)

SANworks Storage Resource Manager is a reporting and event management solution that provides storage data analysis to detect trends, foresee problems, and balance resources by providing automated reporting on storage capacity, consumption, and availability. It automatically scans the enterprise storage topology and collects capacity, consumption, availability, and configuration statistics. It then correlates the data in a Microsoft SQL Server database, and provides out-of-the-box alerts, reports, historical trends, and policies that enable easy, efficient problem detection and correction. Out-of-the-box storage graphs and historical trends provide the knowledge needed to accurately plan for growth, avoid down time, and justify storage capacity.

Software Features / Functionality

- Automated storage monitoring-like having a 7x24 administrator watching over every file, directory, share point, partition and disk on your network.
- Threshold-based alerts and events issued via SNMP, e-mail, NT event log, and Alerts page.
- Consolidated, network-wide reports inventory all storage assets and identify trouble spots.

- Historical planning trends automatically created to eliminate guesswork in capacity planning.
- 100% web architecture for use anywhere on Intranet, WAN, or dial-up.
- Free space alerts, thresholds, and reports for every disk partition and directory prevent server and application crashes caused by insufficient free disk space.
- Reports on files not backed up identify backup holes to prevent data recovery failures.
- File access trends identify unbalanced and overloaded file servers.
- Incremental and full backup sizing reports keep backups inside backup windows.
- Wasted space and largest file reports identify storage that can be reclaimed, deleted, or archived. Custom filters let you identify specific file types.
- Network-wide user and directory disk space quotas alert the administrator (and the user) when users consume more disk space than they really need.
- Computer, user, and directory groups enable user or project-based chargeback on disk space, so departments and business units pay for what they use.
- Capacity planning trends show when, where, how much disk space to add.
- Directory and share point reports show where to add new users and how to load balance existing file servers.
- Customizable storage filters make SAN planning and server consolidation easy.

SANworks Storage Resource Manager uses Server/Agent architecture with SRM Server running Microsoft Windows NT. The SRM Agents may run on Microsoft Windows NT, Windows 2000 Server, Windows 2000 Advanced Server, Tru64 UNIX, IBM AIX, Sun Solaris, HP-UX, and Red Hat Linux. The SRM Server and SRM Agents reside at the server level.

SANworks Storage Resource Manager for Exchange

Compaq SANworks Storage Resource Manager for Exchange is web-based software that provides network-wide reports, alerts, planning trends, and policies for managing Exchange storage capacity, consumption, and availability. SANworks Storage Resource Manager for Exchange provides the centralized reports and policies needed to manage critical, fast-growing Exchange information stores.

Software Features / Functionality

- Alerts provide advanced warning of out-of-space conditions, such as spikes and full partitions, to ensure availability of Exchange services.
- Exchange Storage Capacity Planning Trends enable avoidance of Exchange Server downtime caused by capacity shortages.
- With backup alerting and trending capabilities, backups do not run into production time.
- Exchange Storage Resource Inventory Reports provide a complete view of Exchange storage topology and eliminate the need for time-consuming, manual inventories.
- Mailbox Consumption Policies eliminate the need for labor-intensive Exchange mailbox quota overrides.
- Automate manual daily tasks, such as backup and partition space checks, on a server by server basis.
- Exchange Storage Consumption Reports and Policies identify Exchange users who are sources of inappropriate or unnecessary disk space use for reclaiming.

- Locate and delete "orphaned mailboxes" that are no longer in use.
- Exchange Storage Capacity Planning Trends enable the delay of storage purchases, until necessary, to save capital and depreciation expenses.
- Exchange Storage Capacity Planning Trends enable accurate planning for growth of Exchange storage, and justification of new purchases to management.

SANworks Storage Resource Manager for Exchange uses Server/Agent architecture with both the Server and the Agent supported on Microsoft Windows NT. This application resides at the server level.

SANworks Storage Allocation Reporter

SANworks Storage Allocation Reporter is the industry's first accounting tool that allows SAN storage to be billed as a utility service. Storage Service Providers and IT organizations can track and assign a cost to the storage allocated to an internal department or external storage customer. The application reports on the amount and RAID protection level of storage reserved, quality of storage services delivered; and a pricing model calculates charges on a per gigabyte per time period basis for billing and cost-recovery purposes. This application provides a powerful tool for reporting allocated storage trends, predicting problems, anticipating increased demand, and recovering costs for storage.

Information about storage usage can be accessed anywhere, anytime using the web-based user interface. The Reporter provides the capability for delivering storage as a utility service.

Compaq SANworks Storage Allocation Reporter is a web-based application used for reporting and billing of storage and associated Quality of Service (QOS) reserved by customers in a SAN environment. It provides a leading-edge SAN management tool to report allocated storage trends and costs, foresee problems and provide a cost recovery method for quality of service delivered.

SANworks Storage Allocation Reporter automatically provides reporting and billing of allocated LUN capacity from RAID-level storage subsystems and tracks levels of service to bill against consumers.

Software Features / Functionality

- Configuration of customer IDs, custom LUN attributes (QOS), and price sheets.
- Billing reports capture costs of reserved storage by customer.
- Alerts and notifications issued via SNMP and forwarded to enterprise management frameworks, e-mail, and the Microsoft Windows NT event log.
- Consolidated, LUN-level reports inventory and track all reserved storage assets.
- Historical trends automatically created to eliminate the guesswork in storage planning.
- One hundred percent web architecture for use anywhere on Intranet, WAN, or dial-up.
- Reports on the QOS metrics, such as RAID type, backup, remote mirror, or JBOD, associated with reserved storage capacity.
- Billing reports allow for cost recovery of reserved storage by consumer.
- LUN-level cost recovery independent of storage consumers' platform.
- Dates and costs in billing reports displayed in local format.
- No host agent or software required; supports heterogeneous host operating systems (Microsoft Windows NT, Windows 2000, Tru64 UNIX and OpenVMS, Sun Solaris, HP-UX, IBM AIX, and Linux).

- Supports MA6000, MA/RA8000 and EMA/ESA12000 storage systems.
- Microsoft SQL Server database provides reliable and industry-standard access.
- Web-based online help documentation.
- SANworks Storage Allocation Reporter is a Host-Independent solution.
- This application is at the SAN Appliance Level.

7

SAN Security

Information Security Overview

Information security is a fundamental issue that must be dealt with while managing any data center. Compaq understands the importance and complexity of establishing and maintaining a secure information storage environment. Compaq storage products are designed to make it easy to protect the availability, integrity, and confidentiality of the customer data that they hold.

Compaq is working with other storage vendors in the Storage Networking Industry Association to develop enhanced SAN security technology. Refer to www.snia.org for additional information.

This chapter describes storage aspects of information security in a StorageWorks by Compaq SAN environment.

Basic Security Model

The ideal mass storage system provides fast storage and retrieval of information for a number of servers.

This one line summary leaves unspoken a number of additional expectations: It is expected that data written to the storage system today will be available tomorrow. It is expected that the data will be the same when it's read as it was when it was written. And it's expected that the data is not available to any server or any person not specifically authorized to have access. These three possibilities are covered under the general headings of availability, integrity, and confidentiality.

These additional expectations form the basis for defining the availability and security of the data in the mass storage system. For example, the data should be available even if a hardware or software component in the storage system fails: RAID and remote mirroring technology are methods used to maximize data availability.

Three types of attacks, corresponding to the three aspects of information security, can be made on a computer system. Data can be made unavailable for access. Data can be deleted or modified without permission. Data can be examined without permission. Any computer security system must deal with these types of attacks.[1]

The security of a computer system is the responsibility of a Security Manager. This person defines the operational rules and procedures that are required to maintain the desired security level. To achieve the desired security level in a Compaq SAN system, the operational rules and procedures should incorporate the guidelines discussed in this chapter.

The basic approach to making a system secure is to define one or more security domains. A security domain is a logical grouping of related components in the storage system, along with a set of rules that specify the amount of communication that is allowed between the

[1] An excellent introduction to computer security may be found in "Computer Security Basics", by Deborah Russell and G.T. Gangemi Sr, published by O'Reilly.

components. Devices such as servers and storage systems that are within a given security domain are allowed to communicate with each other. The security manager defines the communication–if any–that is allowed between domains. The security system works by controlling every possible communication path between the security domains, so that data cannot be moved between domains without authorization.

The boundaries of the security domains are physical barriers that control access to the components. The boundaries also control communication between domains through the network or storage bus connections. Any potential path between security domains must be reviewed to make sure that only approved access is permitted. This can be an extremely complex undertaking.

Summary of SAN Security Practices

StorageWorks by Compaq SAN hardware and software components incorporate features that can be used to implement a secure data storage system. The following table shows the appropriate use of these security features in various environments.

Table 7–1: How to Use SAN Security Features

SAN Storage Security Feature	Departmental Storage System	Enterprise Storage System	Service Provider Storage System
Physical security of SAN environment.	Suggested. All personnel are employees.	Suggested. All personnel are employees.	Essential. Many personnel are competitors.
Controlled physical access to switch management using Front Panel Controls.	Suggested. Reduces risk of accidental problems.	Suggested. Reduces risk of accidental problems.	Optional. Switches are to be kept in a secure area.
Password protection on switch management using Telnet via in-band or out-of-band connections.	Essential. Avoid potential of remote access attempts over your network.	Essential. Avoid potential of remote access attempts over your network.	Essential. Avoid potential of remote access attempts over your network.
Disable switch management using SNMP via in-band or out-of-band connections.	Optional. SNMP is useful for system management, and SNMP only allows monitoring of system.	Optional. SNMP is useful for system management, and SNMP only allows monitoring of system.	Essential. Disable by use of license management on switch.
Disable switch management using SCSI Enclosure Services (SES) via in-band connections.	Optional. This tool is useful for system management.	Suggested. This tool is useful for system management, but it increases the inter-departmental risk.	Essential. Disable SES by use of license management on switch.
Disable web browser management interface.	Optional. Password protected.	Optional. Password protected.	Essential. Disable Web OS by use of license management on switch.
Use of zones.	Optional. Use soft or hard zoning as required to manage Operating System conflicts.	Optional. Use soft or hard zoning as required to manage Operating System conflicts.	Optional. Use hard zoning as required to manage Operating System conflicts.
Use of SSP.	Essential. Use as required to manage access to data.	Essential. Use as required to manage access to data.	Essential. Use as required to manage access to data.

Table 7–1: How to Use SAN Security Features (Continued)

SAN Storage Security Feature	Departmental Storage System	Enterprise Storage System	Service Provider Storage System
Controlled access to storage system management using serial line interface.	Optional. Risk from local access is low.	Suggested. Limit physical access to machine room.	Optional. Storage systems are physically secure in this environment.
Controlled access to storage system management using in-band interface.	Optional.	Optional.	Optional.
Restricted use of multiple switches.	Optional. No additional risk is added.	Optional. No additional risk is added.	Optional. No additional risk is added.
Restricted use of multiple storage systems.	Optional.	Optional.	Essential. Each customer must be located on a different HSG80 controller pair.
Restricted use of Management Appliance.	Optional. Appliance applications are password protected.	Optional. Appliance applications are password protected.	Optional. Appliance applications are password protected.
Use of logical unit visibility control on Modular Data Router tape controller.	Essential. Use as required to manage access to data.	Essential. Use as required to manage access to data.	Essential. Use as required to manage access to data.
Event logging enabled.	Essential. Needed to track possible intrusion attempts.	Essential. Needed to track possible intrusion attempts.	Essential. Needed to track possible intrusion attempts.

Security Features of StorageWorks by Compaq SAN Components

The components of a StorageWorks by Compaq SAN are shown in the figure below.

Hardware components include the Host Bus Adapter (HBA) residing in each Application Server, the Fibre Channel Switches that make up the SAN fabric (or fabrics in a multiple fabric SAN), the Disk Storage Systems (including their RAID controllers, cache memory, disks, and related management components), the Tape Storage Systems (including their Modular Data Router gateways), the SAN Management Server, the SANworks Management Appliance, and various communication cables.

Software components include the server operating systems, the StorageWorks Command Console software, the SANworks SAN Management Software, Web Browser and Terminal Emulator interfaces to the Fibre Channel Switch and Storage System management tools, and the MDR management interface.

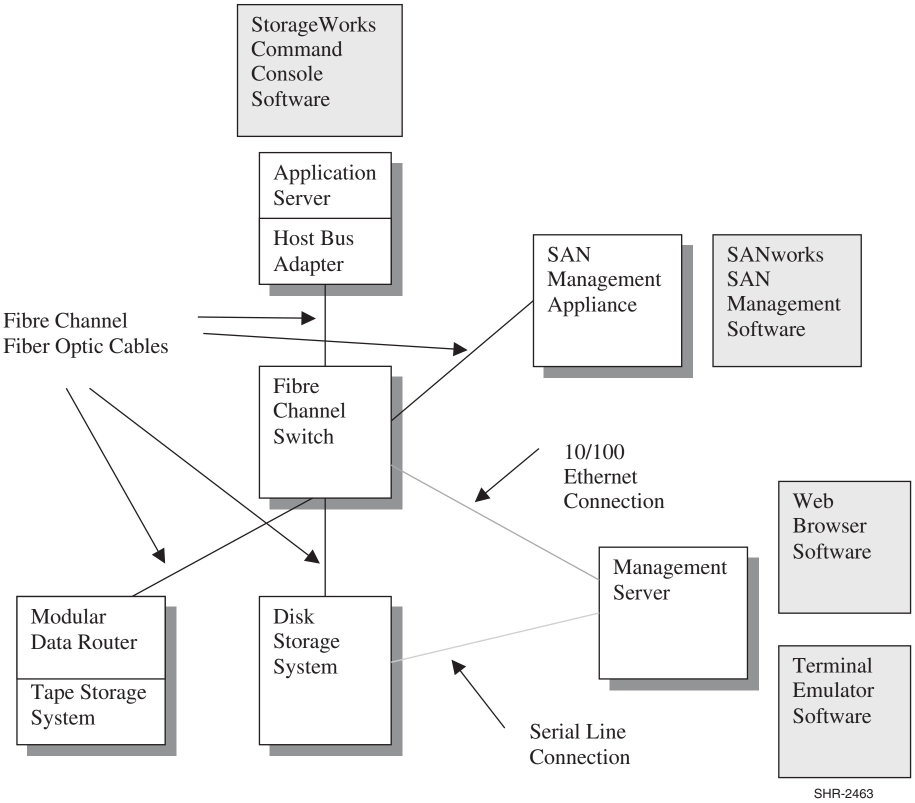

Figure 7–1: SAN Components

The current and future security features of each SAN component are listed below. SAN security is a rapidly developing technology, and the information in this chapter reflects the status of the technology as of the date of publication.

Fibre Channel Fiber Optic Cables

Fiber optic cables used for Fibre Channel communication do not emit electromagnetic radiation. This reduces the risk of security intrusion by means of remote sensing. However, it is not particularly difficult to make a physical tap into an active fiber optic cable. To maximize communication security, the cables should be kept within a secure area.

If a Fibre Channel cable is disconnected, the loss of signal is detected by the connecting devices and is logged in the devices' event logs. Verify that event logging is enabled on all connecting devices that have this feature.

10/100 Ethernet

Because of the difficulty of securing a distributed system, the security of IP LANs is low. The security manager should verify that good passwords are in use on all the SAN components that are connected to a LAN, including the application servers, the SANworks Management Appliance, the management server, and the Fibre Channel switches.

Serial Line

Serial line interfaces are used to connect a terminal (with its associated keyboard and display) to a server or other SAN component. Serial line connections are made using RS-232 physical interface. The EIA-423 protocol is used, and the connection runs at a low speed (typically 9600 baud). The serial line protocol itself does not have any provision for access security.

The security manager should verify that good passwords are in use on all the SAN components that have serial line connections, or that these connection points are in a secure area.

Host Bus Adapter

The host bus adapter (HBA) is the basic interface between the SAN and each server. The microcode in an HBA can be changed by using a utility program. In the case of Windows NT, a microcode load can be done on an active system, and the server does not need to be re-booted to resume normal I/O activity. A new host bus adapter may be installed in an operational server.

In Fibre Channel there is no equivalent functionality to the "promiscuous" mode of operation that historically could be used on 10 Mbps CSMA/CD Ethernet networks. The security risk associated with HBAs in a Fibre Channel environment is low because the switches filter all traffic. Only traffic intended for a given server is communicated between the switch and that server's HBAs.

If the operating system driver is changed, then the system must be rebooted. This minimizes the likelihood of undetected changes to driver software.

Fibre Channel Switch

Fibre Channel switches are connected together to form a SAN fabric. The switches are the foundation of the SAN system.

Switch Management Interfaces

The Fibre Channel switches in a Compaq Storage Area Network support several management interfaces. The interfaces and their security aspects are shown in the following table.

Table 7–2: SAN Switch Management Interface Security

Management Interface	Security Control Method
Front Panel Controls.	Interface supports only basic performance features, not data access management. Physical security required to give complete protection
Telnet via in-band or out-of-band connections.	Password (except for initial power-on) with several levels of access control
SNMP via in-band or out-of-band connections, limited management capability.	Interface supports only monitoring and traps. Physical security required to give complete protection
SCSI Enclosure Services (SES) via in-band connections.	Requires license to enable, physical security required
Web browser via in-band or out-of-band connections.	Requires license to enable, password protected

To maximize security in a SAN fabric that uses SAN Switches, the security manager should verify that the switches are in a secure area, and that the SES management interface is disabled. This prevents the use of the in-band SES interface for intrusion attempts.

Switch Zones

The switches in a fabric cooperate to enforce data access zones. Servers are identified either by the switch port to which they are connected, or by their WWID. These two methods are called "hard zoning" and "soft zoning", respectively.

The advantage of hard zoning is that it is enforced on a port-by-port basis by the switches in the fabric. The disadvantage is that if a server is moved from one port to another, the zone configuration must be changed to reflect the new connection topology. The advantage of soft zoning is that it is independent of port, so servers may be moved from one port to another without changing the zone settings. The disadvantage is that an HBA in a server could, at least theoretically, take on the WWID of another HBA and thus gain unauthorized access to the wrong zone.

The purpose of zones is to manage the interaction of servers in a SAN, preventing interference between the operating system drivers. In heterogeneous configurations the drivers may interfere with each other, and in homogeneous operating system environments the capacity of certain driver data tables may be exceeded. Zoning is used to manage these operational factors.

The security manager should verify that event logging is enabled to record unintended and unauthorized changes to the SAN configuration.

Storage System

Products in the Compaq HSG and HSV series of storage systems incorporate security controls on all the interfaces to the storage system.

Each storage system consists of a pair of HSG or HSV controllers, along with assorted supporting hardware.[1] The storage system is connected to one or more servers, and presents logical disks to those servers. Each logical disk has a logical unit number (LUN).

The Selective Storage Presentation (SSP) feature allows visibility of logical units to be restricted to a subset of the servers connected to the storage system.

Physical Access Control

The storage system is typically housed in a standard Compaq rack with locking front and rear doors. The locks for these cabinets all use the same key, so the security aspect of the locks is only sufficient to deter the most casual intrusion. The locks can be changed to provide additional physical security if desired.

Controller Management

Basic control of the storage system is performed using various buttons and lights on the front and rear panels of the RAID controller shelf. These controls allow the controllers to be halted or restarted. The controller microcode is stored on PCMCIA cards that are inserted into these panels. Physical access controls to the controller shelf must be maintained to prevent unauthorized manipulation of these controls and to prevent unauthorized replacement of the controller microcode.

[1] Refer to the HSG80 and HSV110 documentation for a complete description of the features of the Compaq family of storage systems.

One option for initial setup of the storage system as well as for ongoing operation is to use a serial line connection to each RAID controller. This connection is typically made between a controller and a terminal emulator program running on a nearby computer. All storage system management operations can be done using this interface. Physical access to the controller shelf must be maintained to avoid unauthorized use of this interface.

Another option for the initial setup and ongoing operation of the storage system is to use the in-band Fibre Channel management system. This system sends SCSI commands to logical units on the storage system to control the logical unit definitions and the SSP settings. A server may send these commands to any logical unit to which SSP allows it access.

Data Access Control

The Selective Storage Presentation feature of the storage system is the method used to control access to user data. Access is allowed to each logical unit by one or more servers.

The SSP settings may be controlled by any server having access to any logical unit on the storage system. This includes the SWCC agent, the SSSU tool, and the SANworks Management Appliance, and could include a purpose-built intrusion application running on a server connected to the SAN. If a computing environment has multiple security domains then the domains must not coexist on a single storage system.

For example, consider the configuration shown in the following figure. Server A and Server B have access to logical unit D and logical unit E respectively. Server A and logical unit D are in one security domain, and Server B and logical unit E are in a separate security domain. Since both have access to Storage System C, then Server A may change the SSP settings to prevent Server B from accessing any logical units on the storage system.

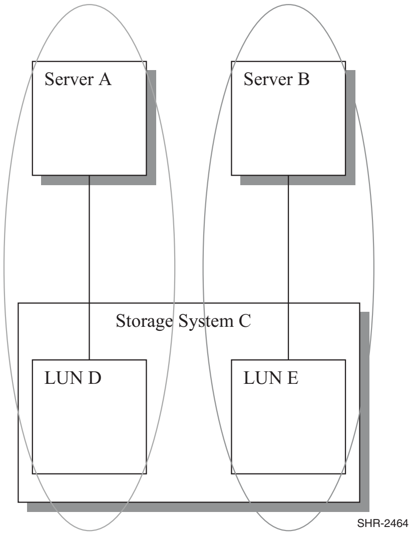

Figure 7–2: Multiple Security Domains on One Storage System

A future version of the HSG Array Controller Software will include a security enhancement that restricts management access to the controller. With this feature, the ability to make configuration changes to an HSG controller is restricted to those servers who are specifically authorized. This will allow multiple servers in multiple security domains to be connected to a single controller (or controller pair).

HSV110 Management Access Control

A management agent can control many storage systems, and many management agents can control a storage system. Without password protection, any management agent on the fabric can access any storage system on the fabric. A password is used to increase the security within your storage subsystem. Specifically, password protection:

- Allows a management agent to control only certain storage systems.
- Allows only certain management agents to control a storage system.

All management functions for Enterprise Virtual Array storage subsystems is done via the SANworks Management Appliance (SWMA). There are two levels of security that are implemented for the Enterprise Virtual Array to control unauthorized access to the storage subsystem.

> The first level controls access to the SWMA itself. User access to the SWMA is controlled by a username and password method that uses the WEBM security model. Without the correct username and password an unauthorized user cannot access the SWMA.
>
> Secondly the storage subsystem has an optional password protection to control which SWMA can manage which storage subsystems. The password is established by entering a password into the operator control panel (OCP) of one of the controllers. Use the Element Manager for HSV Management Agent options to enter the password used by that SWMA to access particular Storage Subsystems.

In addition to the optional storage subsystem zoning on the fabric, this should prevent someone from putting an unauthorized SWMA on the fabric and attempting to manage a HSV storage subsystem.

StorageWorks Command Console Management Software

StorageWorks Command Console (SWCC) is a client-server storage management software product that supports in-band management of Compaq storage systems. An agent program runs on a server and communicates with any storage system attached to that server. The SWCC client program runs on a second, remote server to provide the GUI. The two servers communicate by using a TCP/IP connection between the two servers.

The Command Scripter tool also uses the SWCC agent to communicate with storage systems.

User access to the SWCC agent is controlled by a username and password. Any SWCC client accessing the agent to perform management tasks will be asked for this password. The communications between the management station and the host servers connected to the storage controllers is protected by single-use key encryption. Also, remote configuration can be optionally disabled.

Communication between the agent and the controller is done by using SCSI commands on the Fibre Channel connection between the server and the controller. The agent communicates with a logical unit on the controller.

Storage System Scripting Utility

Storage System Scripting Utility (SSSU) is a character cell interface that allows a user to configure and control Storage Controllers (SCs) generically on a Storage Area Network (SAN). Simple or initial configuration requests can be handled easily and expediently through this simple character cell interface, such as the initial creation of LUNs presented to the host. SSSU meets this requirement with an interface that allows the user to issue simple, terse commands.

SSSU uses the SANworks Management Appliance (SWMA) to communicate with HSV storage subsystems. User access to the SWMA is controlled by the username and password method that uses the WEBM security model.

SANworks Management Appliance

Compaq offers an optional integrated SAN management system that uses an appliance connected to the Fibre Channel fabric. The SANworks Management Appliance hosts web-based Open SAN Storage Management software. This software provides a wide variety of management tools.

Access to the Open SAN Management applications is controlled by a username and password method that uses the WEBM security model.

Storage Security in a Departmental Environment

A departmental computer system handles the computing and data storage requirements for a group of people under one local management umbrella. Since all the people in the department are involved in the same set of projects, the work and the data needed to do the work are both shared widely between the people in the department. This minimizes the data storage security requirements, because everyone involved has access to some or all of the data in the storage system.

This is not to say that there is completely open data sharing. Account passwords, private directories, and routine access controls are used in this environment. But it is not considered a critical security violation if one person accidentally or unintentionally gains access to data that they would not routinely view.

Physical access to the computer system is controlled by the overall security environment of the building or office area where the people work. There may be physical access to the storage system cables, for example. There is not a concern that employees will make a sophisticated and malicious attempt to break into the computer system. Users have physical access to the servers, which implies that the security requirements are not particularly rigid.[1]

Security Expectations

This is an environment with a relatively low requirement for storage system security. Protection against accidental access to the wrong information can be prevented by routine operational practice and default settings on the storage system.

SAN Component Security Attributes

The following features are used to provide security in this environment.

Traditional user account security is in effect in the servers. This protects each user account against accidental access by an unauthorized user. Disk quotas are enabled for each account. This prevents a user from consuming all of the storage capacity of the storage system.

The HBAs pass user I/O requests to a Fibre Channel switch.

The switch is shared by all the users and servers in the system. Configuration management of the switch is done by the system manager using the web management interface. The interface is protected by password to prevent unauthorized changes to the switch configuration.

[1] With physical access to a server, the server's security may be easily compromised.

Data is stored in several StorageWorks storage systems. Access to each logical unit is controlled by the Selective Storage Presentation feature of the array controller.

Response to Attacks

Two attack scenarios are possible in this situation. Accidental inappropriate data access requests might be made by any user, and malicious attempts to make an inappropriate data access requests might be made by a user.[1]

Inappropriate read and write requests by system users are routinely handled by the operating system. Disk mounting requires a privileged account, and directories are protected by access control lists. The benign server environment puts little stress on the security capabilities of the storage system.

Denial of service attacks initiated by system users, whether accidental or purposeful, are not fully protected against. A user could write a program that issues a useless but very high I/O load and consume most of the I/O operation capability of the storage system.

Checklist

For a SAN storage system that requires a low level of security, the following steps are required.

- Good employment practices to minimize malicious attacks.
- Computer system security awareness training for all personnel.
- Routine user account management at the server.
- Disk quotas enabled for all users.
- Passwords enabled on all switch configuration ports.
- Selective Storage Presentation for all logical units.
- Routine periodic security audits.

Storage Security in an Enterprise Environment

In a business enterprise, computer systems are shared between two or more departments. The systems may be managed and operated by an Information Systems organization which has enterprise-wide responsibility for the computing environment. All the people in the enterprise work towards a common business goal, but the day-to-day interests of the departments may vary widely depending on the business climate, time of year, or product development issues. Each department has specific computing requirements that must be met by the IS organization.

The security requirements in the enterprise environment are more stringent than in the departmental environment. There may be wide differences in the need for data security. For example, a typical accounting department has strict security guidelines, while the marketing department may be willing to tolerate more risk.

The IS organization may try to achieve efficiency by placing the computer equipment in a single location. A considerable amount of computer and storage hardware is required for an enterprise of moderate size. This discussion assumes that the storage for all the departments is located in a single SAN storage system. Servers are distributed throughout the facility.

[1] We've ruled out serious attempts to break into the storage system, but unsophisticated attempts to read another’s data are possible in any computer environment.

The IS organization must implement a computing system that meets the security and capacity requirements of all the departments to which it provides service, and the IS security manager must implement a security plan that is suitable for the needs of the enterprise.

To meet the security requirements, many security managers specify a centralized machine room located in a secure area. This substantially reduces the security risk for the storage system, because the ordinary users of the system do not have physical access to the machines.

Security Expectations

This is an environment with a requirement for a high level of storage system security. Protection is needed against unauthorized accidental and malicious data access attempts. The required security level is set by the department with the most strict security needs.

SAN Component Security Attributes

The following features are used to provide security in this environment.

Traditional user account security is in effect in the servers. This protects each user account against accidental access by an unauthorized user. Disk quotas are enabled for each account. This prevents a user from consuming all of the storage capacity allocated to the server.

The HBAs pass user I/O requests to a Fibre Channel switch. Communication is done using Fibre Channel fiber optic cables. These cables pass from the servers into the secure area that holds the storage systems.

The SAN switches are shared by all the users and servers in all departments in the system, and are located in the secure area. Configuration management of the switches is done by the system manager using the web management interface. The interface is protected by password to prevent unauthorized changes to the switch configuration.

Data is stored in several StorageWorks storage systems. Access to each logical unit is controlled by the Selective Storage Presentation feature of the array controller.

Response to Attacks

Two attack scenarios are possible in this situation. Accidental inappropriate data access requests might be made by any user, and malicious attempts to make an inappropriate data access requests might be made by a user.[1]

Inappropriate read and write requests by system users are routinely handled by the operating system. Disk mounting requires a privileged account, and directories are protected by access control lists. The benign server environment puts little stress on the security capabilities of the storage system.

Since the storage systems are located in a secure area, the risk of inappropriate access to the array controllers is limited. There is some risk that the fiber optic cables might be tapped, but this requires a technical approach that is unlikely in this scenario.

Denial of service attacks initiated by system users, whether accidental or purposeful, are not fully protected against. A user could write a program that issues a useless but very high I/O load and consume most of the I/O operation capability of the storage system.

[1] We’ve ruled out serious attempts to break into the storage system, but unsophisticated attempts to read someone else’s data are possible in any computer system environment.

Checklist

For a SAN storage system that requires a moderate level of security, and where the storage systems and Fibre Channel switches are located in a secure area, the following steps are required.

- Good employment practices to minimize malicious attacks.
- Computer system security awareness training for all personnel.
- Routine user account management at the server.
- Disk quotas enabled for all users.
- Locate storage systems and Fibre Channel switches in a secure area.
- Passwords enabled on all switch configuration ports.
- Selective Storage Presentation for all logical units.
- Disable SES management interface to Fibre Channel switches.
- Routine periodic security audits.

Storage Security in a Service Provider Environment

Some organizations provide computing services to their customers on a lease or contract basis. The services may include general-purpose office applications such as Microsoft Exchange or file and print service, or they may be specialized. One example of the latter is the Storage Service Provider, which provides storage capacity to some other organization. In all service provider situations, the service provider is the Compaq customer, and the service provider has second level customers of its own who purchase the service.

These second level customers are the users of the systems.

These users may be competitors of each other, and it is essential that they be protected against security breaches—accidental or intentional—by other users in the computer system. The security plan must take into account the possibility of aggressive attacks.[1] This is probably the most difficult environment for a storage system security manager.

Physical access to the storage system is controlled by placing it in a secure area. The servers are in separate secure areas, segregated by user so that each user has a unique secure server area.

Security Expectations

The requirement is for high security. Each user wants a separate security domain because there is no trust between competitors. Protection against accidental or intentional unauthorized access to data must be provided, and protection against unauthorized changes to the configuration of the storage system is also required. Sophisticated attacks are not expected, but intentional attacks may occur.

At the same time, services providers are very sensitive to cost. There is a desire to share equipment between users to minimize hardware and management cost. This must be balanced against the security requirements.

[1] Denial of Service attacks are not considered to be a problem, because the comparatively small number of users on a SAN makes it easy to identify and eliminate this sort of aggressor.

SAN Component Security Attributes

The following features are used to provide security in this environment.

Traditional account security is in effect in the servers. This protects each user from accidental unauthorized access. Disk quotas are enabled for each account. This prevents I/O from one account from consuming all of the storage capacity allocated to the server.

If one user attempts an intentional attack on another user's data, it may be expected that this would be done from a privileged account on a server. Account security does not protect against this sort of attack, but the exposure from a privileged account is to the data of other accounts on that system, not other users—because they are on their own servers.

The HBAs pass I/O requests to a Fibre Channel SAN switch. Communication is done using fiber optic cables. These cables pass from the servers into the secure area that holds the storage systems.

The Fibre Channel switches are shared by all of the service provider's customers, and are located in the secure area. Configuration management of the switches is done by the service provider's system manager using the serial line interface.

Data is stored in several StorageWorks storage systems. Access to each logical unit is controlled by the Selective Storage Presentation feature of the array controller.

A user may attempt to access a competitor's data. To protect against this possibility, it is important to provide a separate storage system for each of the service provider's customers. While the risk associated with sharing a single storage controller between customers is small,[1] distributing them onto private storage controllers eliminates the risk.

Response to Attacks

Two attach scenarios are possible in this situation. Accidental inappropriate data access requests might be made by any user, and malicious attempts to make an inappropriate data access requests might be made by a user.

Inappropriate read and write requests by system users are routinely handled by the operating system. Disk mounting requires a privileged account, and directories are protected by access control lists.

The storage systems are located in a secure area, but there is some risk of intentional attacks. This is prevented by providing a separate storage controller for each user.

There is some risk that the fiber optic cables might be tapped. While this risk is minimal, the Fibre Channel switch logs must be examined regularly and the configuration change alarms on the switches enabled. These will notify the security manager if this sort of activity occurs.

Depending on the service provider environment, it may be possible for a sophisticated attack on the SAN to take place. This could involve equipment such as frame grabbers or phantom switches. It is extremely difficult to protect against a sophisticated attack against any network system, and Fibre Channel is inherently exposed because the data is sent as clear text. If this level of risk is expected, refer to the following section.

Denial of service attacks initiated by system users, whether accidental or purposeful, are not fully protected against. A user could write a program that issues a useless but very high I/O load and consume most of the I/O operation capability of the storage system.

[1] It requires special knowledge and equipment to successfully complete an unauthorized access to data on an array controller.

Checklist

For a SAN storage system that requires a high level of security, and where the storage systems and Fibre Channel switches are located in a secure area, the following steps are required.

- Good employment practices to minimize malicious attacks.
- Computer system security awareness training for all personnel.
- Routine user account management at the server.
- Disk quotas enabled for all users.
- Locate storage systems and Fibre Channel switches in a secure area.
- Passwords enabled on all switch configuration ports.
- Selective Storage Presentation for all logical units.
- Disable SES management interface to Fibre Channel switches.
- Disable SNMP management interface to Fibre Channel switches.
- Disable web browser management interface to Fibre Channel switches.
- Each user (that is, each customer of the service provider) must have a separate HSG array controller.
- Routine periodic security audits.

Storage Security in a Secure Environment

Some system environments require extremely high levels of security. These are cases of national security or where the data is so sensitive that the owner is willing to make substantial functionality trade-offs to maintain the desired security level. These systems are safe in the face of the worst cases of overt attempts to break into the system by any means possible.

Security Expectations

It is expected that the system will have no exposure to security intrusions. This corresponds to the highest levels of information security.[1] Network systems generally are not able to be audited for compliance with the highest levels of security, because network software is too complex for a comprehensive evaluation. To obtain the highest possible levels of information security, the entire system must be enclosed in a secure environment.

SAN Component Security Attributes

The following features are used to provide security in this environment.

- The system is enclosed in a secure area.

[1] Computer system security ratings are set by NIST/NSA in the US and ITSEC in Europe, and “Common Criteria” is a newly-adopted US standard. Within these classifying bodies, a product can be evaluated at various security levels. Currently, most operating systems are classified at ITSEC's E3 rating or Common Criteria's CAPP protection profile EAL4. The OpenVMS SEVMS product has a E3/B1 rating. Tru64 UNIX has an E2/C2 rating. See http://csrc.nist.gov/cc/.

Checklist

For a SAN storage system that requires the highest level of security, enclose the entire system in a secure area.

- Perform routine periodic security audits.
- Follow other appropriate actions based on the required system security level.
- Place machines in a secure area.

8

Business Considerations

SAN Business Considerations

Why a SAN?

Storage Area Networks (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 increased, changed, or re-assigned more easily.

In a SAN, the servers can access a common storage pool and the backup devices can access the same 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.

Value

A Storage Area Network provides value for the corporation in several ways. A SAN improves 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.

Storage Area Networks can reduce application time to market, help facilitate processing needs, such as report writing and rapid restores, and introduce new concepts such as zero backup time. SANs support remote data copies at nearly unlimited distances and can form the basis for improved business continuance scenarios using disaster recovery/disaster tolerance configurations. They support the latest security measures and can be managed by Web-based tools from a single location or different places, regardless of the physical location of the SAN.

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. Because of this flexibility, a SAN can grow and adapt to the corporate business needs. SANs also enable economies of scale that were previously unavailable to open systems in areas of backup, management, growth, and performance.

Improved Staff Efficiency

A SAN allows you to manage more storage per person compared to server-attached storage. Industry averages indicate that each staff member can manage four more times the amount of storage when using a SAN. How? A SAN reduces costs and administration effort by presenting a common storage pool and a unified, comprehensive toolset for management. The common storage pool supports rapid reconfiguration and allocation of storage without the need for cumbersome server-by-server connections and evaluations.

Flexibility

Compaq SANs support a wide variety of heterogeneous environments and platforms, including Compaq Alpha OpenVMS and Tru64 UNIX, Hewlett Packard HP-UX, IBM AIX, Linux, Microsoft Windows NT and Windows 2000, Novell NetWare, SGI IRIX, and Sun Solaris.

Configuration Choices

Compaq SANs can be configured for performance, capacity, and growth. Each of the StorageWorks components can be made fully redundant, so no single point of failure (NSPOF) exists. This means that all servers sharing storage can be assured that the data will be ready for them at any time.

Time to Market

Compaq SANs can also create data snapshots and copies (clones) for rapid, parallel development efforts. For example, Enterprise Volume Manager (EVM) is the SANworks tool used for creating copies on Windows NT, Windows 2000, Tru64 UNIX, and Sun Solaris systems. EVM allows copies of production data to be used by development groups in a safe, predictable manner. When the testing is finished, the copies can be discarded. The data in the clones is accessible by other servers.

Backup

Compaq SANs create an environment for centralized backup. Compaq's Enterprise Backup Solution (EBS) relieves the corporate LAN of backup traffic. The SAN uses a central location and common tape drives, reducing training and operational costs. Each server accesses the tape drive(s) where its backup occurs. The Automated Tape Library (ATL) and backup control software coordinate all backups. Backup speed is improved as much as 20x faster than LAN based backup.

Storage Area Networks also introduce the concept of distance for remote backup procedures. On a SAN, the ATL can be physically removed from the application servers, on a campus, or even miles distant from the actual data. If necessary, zero application downtime is also possible. EVM can be used to create a copy of the data; the copy is then accessed by a backup server and written to tape without affecting the application server. This allows near-continuous application processing.

User Satisfaction

In general, system uptime equals user satisfaction. Compaq SANs offer new and proven ways to keep the applications running, regardless of component outages or even failures. The ability to configure Compaq storage for disaster recovery/disaster tolerance means that eCommerce and eBusiness applications can be designed and built for continuous operation. While customers and users do not see any loss of service, SANs, and the underlying Compaq "Best Practices" that make them operational, create new levels of service while reducing costs.

SAN Technology Benefits

Storage Area Networks offer the benefits of saving time, saving money, and offering new features. This "technology leap" is evident when SANs are compared to traditional methods of data storage.

SANs may cost more initially, but offer greater scaling and sheer size in terms of the number of servers that can attach to the Storage Area Network and in terms of the storage capacity. In addition, SANs offer a combination of features such as centralized backup, heterogeneity, and flexible use. The flexibility in design possibilities is referenced elsewhere in this document.

The scaling capabilities of a SAN are ideal for large or growing corporations that require rapid deployment, frequent changes, and mixed server environments. Because of the economies of scale provided by the various SANworks and StorageWorks products, Compaq SANs offer true cost savings in operations and training, which allows the corporation to focus on its core competencies.

In addition to these benefits, Compaq SANs offer investment protection in the form of universal disk drives. The disks are packaged in a form factor that is common to all Compaq servers and StorageWorks storage arrays. This means that Compaq customers can extend the useful life of the disks by moving them from the enterprise class arrays to departmental arrays and even to individual servers. Of course, all Compaq StorageWorks products are designed to be user maintainable.

Scenarios

There is a wide variety of Storage Area Network topologies and designs that are possible. Each design should take into account the customer requirements for growth, scaling, data traffic, and performance. The following examples serve to illustrate some of the ways that Compaq SANs can be designed to address various business needs.

1. A customer has multiple Microsoft Windows NT and Novell NetWare servers each with locally attached storage. The environment is dynamic—perhaps a new application is being introduced, perhaps they are migrating from Novell to NT. They spend a lot of time re-configuring storage. Sometimes they run out of disk space on one server while another server has lots of space. The applications are mission-critical; they need to be available.

   The proposed SAN design supports Microsoft Windows NT and Novell NetWare servers on a single StorageWorks storage system. The design is based on a High Availability topology, and has two Fibre Channel 16 port switches. Backup is achieved through a Fibre Channel Tape Controller and an Automated Tape Library. A SAN Appliance monitors the entire SAN for failures. RAIDsets can be easily reconfigured as needed and presented to the various servers.

2. The customer has a back-up window problem. Multiple open system servers are connected to dedicated tape drives (difficult to manage) and/or corporate LAN and the back-ups are not getting done in a timely manner.

   The StorageWorks by Compaq SAN supports centralized backup from a number of hosts. The data is backed up in an Automated Tape Library that connects to the SAN.

3. A Service Provider wants to use SANs to deploy and manage a rapidly growing customer base. A SAN topology supports rapid growth and a variety of data traffic patterns. It is modular and extensible. Servers and storage can be added to the SAN as business needs dictate.

4. A customer is located near a coastline that is subject to severe weather and has mission critical data and other less critical data. The mission critical data must be available even if the primary data center is shut down due to storms, while the non-critical data may be unavailable for up to a week. The plan is to implement a SAN for data management, and to integrate both the critical data and the non-critical data into the storage network. They

want to use the controller based Data Replication Manager solution to protect the mission critical data. There are multiple operating systems involved including Microsoft NT 4.0, Microsoft Windows 2000, Tru64 UNIX, OpenVMS and SUN Solaris.

Data Replication Manager co-exists with other heterogeneous storage and servers in the same SAN, allowing an integrated SAN with private zones for each unique operating system with its own unique DRM solution.

9

Best Practices

Overview

This chapter describes "best practices" for implementing heterogeneous Storage Area Networks. The information contained in this chapter should be used as a guide for constructing your SAN. Although every attempt has been made to provide a best practice recommendation, some aspects of SAN implementation are a matter of preference. Also, the physical location of servers, storage, computer labs, or specific building layout and location may dictate particular aspects of your SAN implementation. In part, this is an expected reality and is often easily accommodated, given the inherent flexibility in implementing SANs and Fibre Channel technology.

Rather than just present a list of best practices, the information has been organized into these sections:

- Planning a SAN
- Configuring a SAN
- Upgrading a SAN
- Migrating SAN Topologies
- Merging SAN Fabrics
- Troubleshooting

Much of what is presented here is the result of the actual experiences of building large SANs within the internal Compaq engineering environment and at customer sites.

Although this chapter does describe portions of the design process in the planning phase below, it is not meant to convey the entire SAN design process. Contact a Compaq Enterprise Storage Consultant or the Professional Services organizations for assistance and consultation on designing SANs. Compaq Storage Services may be contacted through this link:

http://www.compaq.com/services/storage/index.html

Planning a SAN

Proper planning considers both present and future requirements. This can be accomplished by over-planning your initial SAN capacity and connectivity requirements to accommodate expected future needs. Whether using a Compaq standard topology or designing your own topology, select a design that not only offers the best implementation for present usage, but also allows you to expand your SAN over time.

It is important that you allocate an adequate amount of time to plan your SAN. In general, the more detail you can define in the planning phase, the greater the benefit you will realize during the configuration phase.

Consider each of these items during the planning phase:

- Deployment Strategy: You can choose to deploy separate smaller SANs or SAN Islands with the idea of increasing capacity by growing the SANs independently or by interconnecting the independent SANs in the future. Smaller SANs are easier to construct, larger SANs offer economies of scale from an operational standpoint, but take longer and are more complex to build.
- Topology Design: Consider the topology design compared to the ease of migrating to another, higher capacity design. In most cases this can be accommodated; however, it is always preferable to choose an initial design that can grow, without the need to transition to a different topology.
- Experience Level: If you are just beginning deployment of SAN technology, consider starting with a smaller implementation. As you gain experience, deploy larger SANs.
- SAN Management Strategy: Refer to *Chapter 6, "SAN Fabric Management Tools*" and Chapter 6, "SAN Storage Management Tools" for information about SAN management tools. After reviewing this chapter, define the management strategy and the specific tools that you will utilize to manage your SAN.
- Technology Advances: The ideal design considers expected future technological advances, and can easily accommodate the resultant changes. Plan for flexibility in your initial design. Higher port count Fibre Channel switches and faster interconnect speeds are an inevitable evolution of Fibre Channel technology. Ensure that your initial plan addresses and can accommodate expected changes such as these.
- Document the Design: This is one of the most important aspects of the planning process. This allows you to fully review and evaluate the design beforehand, evaluate trade-offs, make changes, and effectively communicate specific plans to all groups affected. The other important benefit of documenting your design is that during the later phases of implementation, the documentation serves as the roadmap for the actual implementation.

Compaq recommends, at a minimum, that you document the following before beginning the actual implementation:

1. Topology Map–Shows the logical SAN topology and fabric interconnect scheme; conveys the overall design from a strategic standpoint, and can also serve to convey how future growth and technological advances will be accommodated.
2. Configuration Layout–Shows the physical layout of the entire implementation. More detailed then the topology map, the layout is used during implementation to verify the correct connectivity. This is also extremely helpful if troubleshooting is required in later phases.
3. Storage Map–Defines the storage system arrangement and configuration in the SAN, and storageset settings such as SSP and RAID levels. This map effectively defines how all of the storage is configured in the SAN.
4. Zoning Map–Defines the inter-node communication access within the SAN. This map defines which nodes or device ports are allowed to communicate with each other in the SAN.

NOTE: A key decision in the zoning implementation process is determining whether you will implement hardware or software zoning. Hardware zoning offers higher security than software zoning but is less flexible because device-to-switch cabling changes require zoning information to be updated.

General Planning Considerations

It is difficult to make general recommendations about the choice of a specific SAN topology. There are so many variables in large installations that each new configuration requires substantial customized design work. The following suggestions provide background information for designs that meet typical large SAN requirements and that are compatible with the future direction of StorageWorks by Compaq SAN technology.

Advantages of Dual Fabric SANs

Most large SANs should have two or more independent fabrics. Each fabric operates independently, and the failure of one fabric does not cause a complete loss of SAN communication.

The reliability of modern electronic hardware is so high that it is difficult to make meaningful predictions of failure rates. Software is used in all components, but it is difficult to estimate the likelihood of software failures. Operator errors are the most likely cause of problems, and the frequency of operator errors depends strongly on operational discipline and employee morale, both of which are very difficult to quantify. All of these potential failure points are minimized by the use of multiple fabrics.

The advantage of dual fabric designs is that they support path failover technology. Path failover is available in most operating systems that are supported in Compaq SANs. Two host bus adapters are used in each server, and if the communication path from one HBA to the storage system fails, then the I/O traffic is re-routed through the other HBA.[1]

The two fabrics should be similar in size and topology. This minimizes the risk of asymmetrical performance under certain workloads, and minimizes the total cost of the SAN. Failover software does not support the concept of primary and secondary fabrics.

It should be noted that there is not an automatic increase in cost caused by the use of two separate fabrics. For example, two switches in a single fabric give about two dozen usable ports (depending on the topology). Two separate fabrics, each with a single switch, gives 32 ports at the same cost.

Many of the SAN illustrations in this document show only a single fabric. This is because most of the design and compatibility requirements apply to each fabric as a complete unit. However, practical SAN designs should have two or more fabrics, each satisfying the configuration rules described in this guide.

Data Access Patterns

There are several supported Compaq SAN topologies, suitable for a wide range of applications from small to very large systems. For small installations, the topology may be chosen to maximize connectivity or to minimize cost. SAN performance is not likely to be an issue for a small installation, because of the very high I/O throughput that is provided by basic Fibre Channel SAN components.

Large installations must be designed to maximize performance and minimize cost, to support current and future connectivity requirements, and to enable eventual migration to new technologies. Several factors must be taken into consideration to meet these requirements. The factors are categorized into three different data access patterns, one-to-one, many-to-one, and any-to-any.

[1] Failover can also be useful in SANs with only one fabric. This protects against HBA failures and certain extremely unlikely potential problems in array controllers. In general, fail over technology should be used in SAN configurations that have two or more fabrics.

- One-to-one

  The communication paths within the fabric are used in different ways, depending on the relationship between the servers and the storage systems. In some cases, each specific server stores data on only one or two storage systems. In this case, only a few specific storage systems service all I/O requests from a server, and there is little or no communication between the servers or between the storage systems. A given fabric port sends requests to one (or two) specific fabric ports. This is the traditional server-storage relationship. Many systems still operate this way today.

  From the viewpoint of the fabric, the I/O traffic has a “one-to-one” pattern, and the traffic pattern is stable. Each server sends I/Os to a small, specific set of storage systems, and each storage system is associated with only a handful of servers. Only significant changes to the configuration by the system manager will change the connection pattern.

- Many-to-one

  Multiple servers accessing data stored in a single centralized pool is another data access pattern. This is a common situation when high performance storage systems have enough capacity to handle a number of servers. In this environment, there is a “many-to-one” I/O traffic pattern on the SAN fabric, and the traffic pattern is stable. Each server sends I/O requests to a small set of storage systems, but each storage system may service a large number of servers. The connection pattern changes only when significant changes to the configuration are made by the system manager.

- Any-to-any (or many-to-many)

  In a third case, application servers access data that is distributed across many storage systems. This case may develop in several situations. The latest Compaq storage arrays may handle a large number of servers. (Refer to the configuration rules in this Guide for detailed information.) A system manager may decide to distribute information over a wide set of storage systems, thus requiring each application to access multiple storage systems. This situation can arise when host-based mirroring is used. Another possibility is that it may be easier to manage the data if it is partitioned and stored on multiple storage systems. For example, Accounting Department data might be stored on one storage system, and Personnel Records data on another. A server requiring access to both data types generates I/O requests to both storage systems.

  Another important situation where data is distributed across a range of storage systems is when the Compaq VersaStor virtualization technology is used. VersaStor distributes data over all the available storage systems in a SAN.[1] In this case, I/O requests from a given application server are handled by one or more storage systems, in a pattern that is controlled by the virtualization management appliance. In this environment, many servers access many storage systems, which is a “many-to-many” pattern. Management traffic may occur between servers, storage systems, and management appliances.

  From the viewpoint of the SAN fabric, any port may send traffic to any other port, which is an “any-to-any” pattern. Furthermore, since the virtualization manager performs dynamic reallocation of storage system capacity, the traffic patterns vary continuously without manual intervention.

The optimum SAN configuration depends on the I/O traffic, whether it be one-to-one, many-to-one, or any-to-any pattern.

[1] The specific configuration details are controlled by management options.

Core and Edge Switch Concept

In the future, most large SANs will support any-to-any traffic patterns. The remainder of this chapter focuses on this problem.

The optimum fabric configuration uses a high performance “core” surrounded by a number of “edge switches.” The core provides roughly equal connection performance between any pair of ports. The edge switches provide port aggregation to match the performance requirements of the servers and storage systems to the performance of the core.

The figure below shows a large configuration that uses the core and edge switch approach. Using 16-port switches, the core is a 32 port fat tree. Four ISLs go between each switch pair in the fat tree. Two ISLs connect each edge switch to the core.

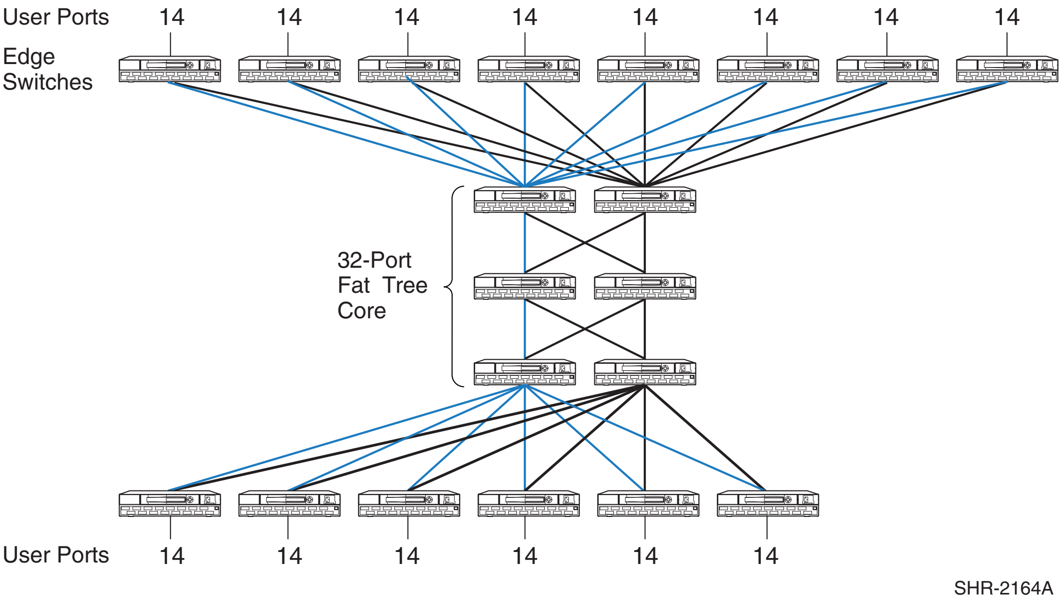

Figure 9–1: Example of Core Switch Plus Edge Switch Configuration

Fabric Core Options

The simplest fabric core is a single switch. Fibre Channel switches support simultaneous full bandwidth connections between any combination of port pairs. A single switch fabric core guarantees support for any-to-any traffic.

Any combination of switches has less performance than a single switch, and the difference depends on the fabric topology. The best-performing topology is the “fat tree”, which has enough Inter-Switch Links (ISLs) to provide, on the average, full bandwidth connections between any combination of port pairs. While it is possible to construct workloads that force traffic contention on the ISLs of a fat tree, which reduces the throughput, fat tree fabric core topologies provide full-bandwidth any-to-any communication, on the average, for random traffic patterns.

A related topology is the “skinny tree”, which has fewer ISLs and fewer switches. This topology introduces an unavoidable performance limit to the fabric. In many cases this limit is beyond what is required by the application servers. The process to upgrade a skinny tree topology to a fat tree topology is fairly straightforward, involving the addition of switches and ISLs to the existing tree

Edge Switch Options

The simplest edge switch is a single switch with one ISL connecting it to the fabric core. Each edge switch provides “User Ports” for connecting servers and storage systems.

The single ISL is a potential bottleneck. All the I/O traffic from the servers or storage systems connected to the edge switch must pass through just one ISL. More ISLs can be provided. Several combinations of ISL and user ports may be used. For example, with sixteen port switches, the ISL to user port ratio could be 1:15, 2:14, 3:13, 4:12, etc. Each of these combinations represents a “port aggregation ratio.” The ratios are 1:15, 1:7, 3:13, 1:3, etc.

The workload of the servers and storage systems attached to an edge switch determines the required port aggregation ratio for the switch. For lightly loaded application servers, a 1:15 port aggregation ratio may be adequate. Heavily loaded servers may require a 1:7 or 1:3 ratio. Extremely high performance servers, such as the Compaq Wildfire system, may be able to completely “fill up” a Fibre Channel connection. In this case, there is no advantage to using an edge switch, and the server should be connected directly to the fabric core. Storage systems may also be able to support a full bandwidth Fibre Channel connection.

To select the appropriate port aggregation ratio, refer to the I/O requirements of your applications and servers. This information is available for many situations by using the Active Answers application sizing tools. In other cases, measurements of an existing system may be required to determine the workload.

Designing a Subsettable SAN

In many cases, the growth pattern for a storage installation is difficult or impossible to predict. Global economic growth, conditions in a given business market, the growth rate of your company, and internal reorganizations or reallocations of computing resources may all have a significant impact on the requirements that must be met by the SAN.

To accommodate this unpredictable variability, the SAN designer should plan for growth within a predefined design. The initial installation should be a subset of a larger pre-designed configuration.

The “core plus edge switch” approach supports this strategy for SAN design.

When the time comes to expand an existing installation, the system manager can make incremental changes to the configuration rather than a complete reconfiguration of the entire Fibre Channel fabric. Changes to the fabric core are isolated from the edge switches, which minimizes the impact of changes required to support core growth. Changes to a given server’s connection to an edge switch are isolated from the core, which minimizes the impact of server-related changes. Furthermore, since two or more fabrics are in use, server I/O traffic may be temporarily forced to a single fabric while the other fabric is undergoing modification.

Start with a single switch core for a moderate sized initial installation,. When needed, the core can be expanded by replacing the switch with one that has more ports, or by reconfiguring the core to a skinny tree or fat tree topology. An existing fat tree core may be expanded by replacing it with a fat tree made up of switches with more ports, or by reconfiguring it to a wider fat tree configuration.

Use a generous estimate of the required I/O performance when selecting edge switches. A port aggregation ratio of 1:7 or 1:3 is adequate for most applications. Increasing bandwidth is a simple, localized modification, if it turns out that more is required.

The initial design should include spare ports on the core to support the future addition of edge switches. For example, consider a configuration that uses sixteen port switches, a single switch core, and edge switches with a port aggregation ratio is 1:3. This design supports up to four edge switches and 48 user ports. This would be a suitable solution for a system where 36 ports are required now, requiring three edge switches. Future growth to 48 ports can be accommodated by adding another edge switch.

SAN Design Summary of Recommendations

Large SANs should include the following features.

- Multiple independent fabrics.
- Core plus edge switch topology.
- Appropriate port aggregation ratio, depending on application server requirements.
- Appropriate core design, depending on number of ports required.
- Subsettable design, with initial installation suitable for current needs.

By following these guidelines for SAN planning, your design will be suitable for supporting future storage technology and future growth in your storage environment.

Configuring a SAN

Once you have completed the planning phase you can begin to configure your SAN. As described in the planning phase, it is important that you document the configuration. During the configuration phase, you should be recording the details of the actual physical configuration.

- Recording. As you construct the SAN, record the cable connections and mark this information on the configuration layout diagram. Record the WWID of all nodes and devices and identify where they physically reside. It is recommended that you place a label on each Fibre Channel HBA with the WWID clearly identified. Compaq storage systems are pre-labeled with this information; however, you may wish to place an additional label on the front of the unit in plain view.
- Cabling. Define a system for cable labeling. Even a small SAN can include a very high number of fiber optic interconnect cables. Label both ends of each cable with the same unique cable number or color code scheme. This will allow you to quickly identify each cable uniquely. Also consider placing a label at each end of the cables that identifies connection points at both ends, such as “TO” and “FROM”. Use label types that are easy to create and read, and ensure they are attached securely to the cable.
- Protect unused or open switch ports with port plugs. Never leave ports exposed.
- Cable Dressing. Use care when routing fiber optic cable and ensure that you do not exceed the recommended minimum bend radius. For single-mode and multi-mode fiber cable the minimum bend radius is 25 mm. Where cables are bundled or hanging unsupported, use velcro tie wraps to group and support the cables. Never use plastic tie wraps as they can damage the internal fiber core if over-tightened.
- Cable Symmetry. When connecting cables, consider slot/port-numbering symmetry. Be consistent across similar servers with cabling in terms of HBA slot placement and cabling to switches. If configuring with two SAN Fabrics and multi-pathing, connect HBA 1 to SAN Fabric 1, HBA 2 to SAN Fabric 2, etc. Cable symmetry is not a requirement, but serves as an aid to troubleshooting if this is eventually required.
- Configure Fibre Channel Switches. Although all Compaq Fibre Channel switches are pre-configured, verify that all Fibre Channel switches in the fabric have the same parameter settings and that each has a unique domain ID.

Label switches using a relevant naming scheme particular to the topology. For example, if implementing a ring topology, label each switch in the ring as Ring1, Ring2. Although not an absolute requirement in all configurations, it is highly recommended that all switches utilize the same switch firmware revision. Different switch code revisions running in the same fabric are supported during a rolling upgrade. This is considered a temporarily acceptable situation for the duration of the code update.

- Configure Servers. For each platform or operating system type, utilize the appropriate Compaq StorageWorks platform kit to ensure that the required server drivers and configuration settings are loaded. Ensure that servers are configured with the proper operating system versions and all required updates.

  Use a numbering type scheme for naming multiple servers of the same type, such as NT01 and NT02 for Windows NT servers.

- Configure Storage. Use the storage map created in the planning phase to configure each of the storage systems. Verify server-to-storage connectivity, and access one server at a time.

  When initially defining storagesets, always disable all access first, and then enable the desired individual access. For Enterprise/Modular RAID Array storage systems, define connection names to be consistent with zoning alias names. Be consistent with connection names relative to storage port and controller connection. Choose a scheme that is easily understood and quickly conveys the physical connectivity.

- Define Zones. Use the zoning map to configure zones. Consider starting with small zones that allow a smaller logical subset of a larger physical SAN to be tested initially.

  Always save old zoning configurations before and after making any zoning change. If possible, it is recommended that no zoning changes be made when an individual switch normally configured in the fabric is temporarily not available.

  You can zone by operating system or by storage system. Zoning by operating systems is useful when the operating systems are accessing storagesets that are localized to specific raid arrays. For example, NT1, NT2 and NT3 have access to storage on ARRAY1, and VMS1, VMS2 and VMS3 have access to storage on ARRAY2.

ZONE NAME	**NT_ZONE**	**VMS_ZONE**
Members	NT1	VMS1
	NT2	VMS2
	NT3	VMS3
	ARRAY1	ARRAY2

ARRAY1 will only have host connections for the NT1, NT2 and NT3 servers and ARRAY2 will only have host connections for the VMS1, VMS2 and VMS3 servers.

Zoning by storage system will limit the connections to the G80 to those systems actually having storagesets on them. This is useful when the storagesets for a specific system are on multiple storage systems.

In the above example, we add 3 more NT servers and another storage system to the NT zone:

ZONE NAME	**NT_ZONE**	**VMS_ZONE**
Members	NT1	VMS1
	NT2	VMS2
	NT3	VMS3
	ARRAY1	ARRAY2
	NT4	
	NT5	
	NT6	
	ARRAY3	

Both Array1 and Array2 will have host connections from all 6 NT systems. This may not be a problem in a small SAN, but as the SAN grows the connections will increase. Also, we do not know which of the NT servers are accessing storage on ARRAY1, and which ones are accessing storage on ARRAY2.

If we zone by storage system we get:

ZONE NAME	**ARRAY1_ZONE**	**ARRAY3_ZONE**	**ARRAY2_ZONE**
Members	NT1	NT4	VMS1
	NT2	NT5	VMS2
	NT3	NT6	VMS3
	ARRAY1	ARRAY3	ARRAY2

Zoning this way also makes it much easier to troubleshoot, especially if servers access storage on multiple arrays. We could have a zone that looks like this:

ARRAY1_ZONE	ARRAY3_ZONE	ARRAY2_ZONE
ARRAY1	ARRAY3	ARRAY2
NT1	NT1	NT4
NT2	VMS2	NT5
VMS2	VMS3	VMS1
VMS3	NT5	NT2
NT6	NT6	NT6

This way it is more apparent that NT1 is only accessing storage on ARRAY1 and ARRAY3. If part of storage can not be seen then it is easy to locate the source of the problem.

Due to some zoning restrictions, you may need more than one zone for a particular ARRAY. If ARRAY1 also has Tru64 UNIX servers, we must zone that separately.

ARRAY1_ZONE1

ARRAY1

Tru64_1

Tru64_2

Zone and Zone Alias Names

When setting up zoning, use meaningful names for zones and zone aliases and be consistent with the naming convention throughout the fabric.

Servers are identified by the WWID of the host bus adapter. Name these by using the system name and the host bus adapter number. For example, server NT1 with one Fibre Channel HBA would have an alias of NT1_HBA1. Server NT1 with a second HBA would have an alias of NT1_HBA2

RA8000 storage systems in a transparent failover configuration will have two WWID's on the fabric, one for port 1 and one for port 2. Give each RA8000 a unique number. RA8000 number 1 could have aliases of R1_P1 (port 1) and R1_P2 (port 2)

For a multiple-bus failover configuration the RA8000 will present 4 WWIDS to the fabric. If you have a multi-path NSPOF configuration, two of the WWID's will be in one fabric, the other two will be in the second fabric. Name the ports using an alias such as R2_A1 (Controller A Port 1), R2_A2 (Controller A Port 2), R2_B1 (Controller B Port1), and R2_B2 (Controller B Port 2).

Ports A1 and B2 will be cabled to the first fabric. Ports A2 and B1 will be cabled to the second fabric. The aliases in fabric 1 will be R1_A1 and R1_B2, the aliases in the second fabric will be R1_A2 and R1_B1. Keep the ports and HBAs the same throughout the setup. For example, always have HBA 1, R1_A1 and R1_B2 in fabric1 and HBA 2, R1_A2 and R1_B1 in the second fabric.

Using this convention conveys the failover mode that the RA8000 is configured for. Any alias with a P1 or P2 is in transparent mode, any alias with A1, A2, B1, or B2 is in multiple-bus mode.

Define RA8000 host connection names for the adapter WWID's in the same manner as you defined the alias name in the fabric. For example, the fabric alias name for NT1, HBA1 will be NT1_HBA1. The host connections on the RA8000 controller should match this as closely as possible.

Example:

Alias NT1_HBA1 in the fabric would have host connection names on the RA8000 of:

```
NT1-P1 WINNT THIS 1 081200 OL this 30
    HOST_ID=2000-0000-C922-8ADC    ADAPTER_ID=1000-0000-C922-8ADC

NT1-P2 WINNT OTHER 2 081200 OL other 130
    HOST_ID=2000-0000-C922-8ADC   ADAPTER_ID=1000-0000-C922-8ADC
```

NOTE: While storage system connection names are not case sensitive, switch alias names are. That means that the switch might have a alias name of TRU64_1 and another alias name of Tru64_1 that refer to two different sets of things.

Upgrading a SAN

Upgrading a Fibre Channel Switch

See the Installation and Hardware Guide for your switch.

Scaling a SAN

The information in this section applies to all SAN topologies, whether a custom design or Compaq defined.

- Replace 8-port switches with 16-port switches.
- Add additional switches, up to the limits specified for a single fabric in Chapter 3, "SAN Fabric Design Rules".
- Add a second fabric as a high availability no single point of failure solution.
- Deploy multiple independent SANs.
- Migrate to a different topology (see below).

Scaling Specific SAN Topologies

The information in this section is specific to the Compaq-defined topologies. Refer to the Fibre Channel switch replacement procedure elsewhere in this chapter for information about preventing fabric segmentation when adding new switches to an existing fabric.

Whenever you are expanding a topology, ensure that the new switch and device connectivity is consistent with the original SAN topology design requirements and goals. Avoid making changes to the topology that may serve to disrupt the original topology design goals. If you need to make topology changes based on a change in data access requirements, consider migrating to a different topology that is better suited to meet these needs. It is important in any expansion that the original data access needs be maintained.

If you have implemented a high availability fabric design (refer to Chapter 2, "SAN Topologies"), it may be possible to expand your SAN in a non-disruptive manner. It is highly recommended, however, as a precaution, that all data be backed up and that I/O activity quiesed when adding new switches to the fabric.

Cascaded Fabric

Expand an existing cascaded fabric by connecting a new switch to an available port on an existing switch. If there are no available ports, remove a device or set of devices from an existing switch, connect the new switch to those ports, and connect the device or devices to the new switch.

Meshed Fabric

Expand an existing meshed fabric by connecting a new switch to available ports on an existing switch. If there are no available ports, remove a device or set of devices from an existing switch, connect the new switch to those ports, and connect the device or devices to the new switch. To maintain the meshed topology, you must ensure that there are multiple paths (ISLs) connecting the new switch to the existing meshed fabric.

Ring Fabric

Expand an existing ring fabric by breaking the ring and inserting another switch into the ring.

Add new switches cascaded off of the ring, up to the maximum number of switches supported in a single fabric. When expanding outside of the ring, ensure that no two devices that need to communicate are more then seven hops apart.

Tree Backbone Fabric

Add edge switches. Expand an existing Tree Backbone SAN Fabric by adding additional edge switches. Connect these edge switches to available ports on the one or two backbone switches.

Add a second backbone switch (if your current design only contains one). Connect all of the edge switches to the new backbone switch.

Migrating SAN Topologies

This section describes how you can convert from one topology type to another if required. Compaq highly recommends that you thoroughly review your initial design to ensure that it meets your present and future requirements in order to avoid having to modify your initial topology design. There may be situations, however, based on changes in business requirements, that require you to consider converting to another topology type. For those circumstances, information is provided below that can help you gain an understanding of how the different topologies can be converted.

As described in the planning phase, it is important that the SAN Fabric topology be well documented. If you are required to change from one topology type to another, use the existing topology diagrams to determine the most efficient manner in which to modify the topology. Create a new diagram that details the desired final connectivity scheme and use this as a map for the topology migration or conversion.

If you have implemented a high availability fabric design, depending on the specific cabling changes required, it may be possible to migrate your SAN in a non-disruptive manner. It is highly recommended, however, as a precaution, that all data be backed up and that I/O activity be quiesed when migrating or reconfiguring any portions of the fabric.

If you have implemented a two-fabric, no single point of failure (NSPOF) SAN, you have the ability to failover over all operations to one fabric while you reconfigure the other fabric. This makes it possible to perform a totally non-disruptive topology migration.

- As a general rule, migrations that only require the addition or re-cabling of ISLs are less disruptive then migrations that require devices be moved from one switch to another. When planning a migration, try to avoid or minimize scenarios that require moving devices from one switch to another.
- Cascaded to a Meshed Fabric. Whether you have implemented a linear cascade or branched cascade of switches from one top switch, additional ISLs are required to connect all switches together as required in a mesh fabric design. Proper planning requires that you carefully calculate the number of additional ports that are needed for the additional ISLs. This may require that devices be moved from one switch to another.
- Cascaded to Ring Fabric. If you have implemented a linear cascade, connect the last switch in the cascade to the first switch to create a ring fabric. For a branched cascade, extensive ISL re-cabling may be required.
- Cascade to Tree Backbone Fabric. Whether you have implemented a linear cascade or branched cascade, determine which switch or switches will be utilized as the backbone (typically no devices) and which switches will be edge switches. Connect all edge switches to the backbone switch or switches; connect all devices to the edge switches. This conversion is less disruptive if you add new switches to the fabric for the backbone and use all of the existing switches as edge switches. In this case, you can simply connect one end of the existing ISLs to the new backbone switches.
- Meshed to Ring Fabric. A meshed fabric can be converted to a ring fabric by simply removing the cross-connected ISLs, leaving the outer connected ISLs connected as a ring. The available ports can be utilized as additional redundant ring ISLs or for additional devices.
- Meshed to Tree Backbone Fabric. Determine which switch or switches will be utilized as the backbone (typically no devices) and which switches will be edge switches. Connect all edge switches to the backbone switch or switches; connect all devices to the edge switches. This conversion is less disruptive if you add new switches to the fabric for the backbone and use all of the existing switches as edge switches. In this case, you can simply connect one end of the existing ISLs to the new backbone switches.
- Ring to Meshed Fabric. If you have implemented two ISLs between all switches in the ring, move one end from an ISL between any two switches to the appropriate switch based on the final mesh design. Repeat this for all of the second ISLs between any two switches. There may be an optimal place to “break” the ring relative to re-cabling. Evaluate different scenarios prior to performing the actual conversion.

- Ring to Tree Backbone Fabric. Determine which switch or switches will be utilized as the backbone (typically no devices) and which switches will be edge switches. Connect all edge switches to the backbone switch or switches; connect all devices to the edge switches. This conversion is less disruptive if you add new switches to the fabric for the backbone and use all of the existing switches as edge switches. In this case, you can simply connect one end of the existing ISLs to the new backbone switches. It is also less disruptive if you have implemented 2 ISLs between all switches in the ring in your original design.

Merging SAN Fabrics

This section describes the process for merging two (or more) independent fabrics into a single, larger fabric. This is typically done when you:

- have grown independent SAN islands to the point where more resources are needed
- wish to share the resources in two or more fabrics
- wish to make information in one SAN available to servers in another SAN

With support for longer distances you may also desire to connect geographically separated SAN islands together into a single SAN, spanning across very long distances.

Although StorageWorks by Compaq SAN designs and components allow versatile configurations, Compaq highly recommends that you thoroughly review all SANs to ensure they will meet existing SAN rules after they are merged into a single fabric. The newly created fabric should not exceed any existing SAN rules.

Merging fabrics can be a complicated process, especially if the fabrics are large. The procedures in the document require a complete understanding of fabrics, zoning commands, and rules. They also require that the user understand how to use the telnet commands as well as the web-based GUI.

It is important to consider not only current SAN configurations but any future SAN needs that may be required. Most difficulties related to merging SANs are due to the fact that not enough planning was put into future SAN considerations at the time the initial SAN was designed and built. Another problem is that the SANs being merged may be implemented differently.

When fabrics discover each other they must go through basic login procedures, or sanity checks, to determine if they are compatible to work as one fabric. If the discovery process determines they are not compatible then the fabric will segment. This means that although they are physically connected, they will still run as separate fabrics.

When zoned fabrics merge they append their zone configuration database to include each fabric's zone configurations. If a non-zoned fabric merges with a zoned fabric, all zoning information is proliferated to the non-zoned fabric switches. If there was a zone configuration enabled at the time of the merge, then that zone configuration will be enabled on the non-zoned fabric switches as well. This means that any devices that were in the non-zoned fabric will be not accessible until they are added into the current enabled configuration.

Please review these causes of SAN segmenting prior to physically connecting multiple fabrics together.

- *The name of a zone object in one fabric should not be used for a different type of zone object in the other fabric (Zone type mismatch).* In other words, if you create a zone name on Fabric A, that same name should not be an alias or configuration name in Fabric B; otherwise the fabrics will not merge.

- *The definition of a zone object in one fabric is different from its definition in the other fabric (Zone content mismatch).* If an alias, zone or configuration name is the same on both Fabric A and B but the content or definition of that object is different between the fabrics the fabrics will not merge.
- *Zoning is enabled in both fabrics and the zone configurations that are enabled are different (Zone configuration mismatch).* Because of this mismatch the switches within each fabric are not going to assume one fabric has the correct zone configuration enabled. The fabrics will not merge until one of the merging fabrics has its zone configuration disabled.
- *Not only must each switch within a fabric have a unique domain ID but each switch within the multiple fabrics of the enterprise should have a unique id as well.* For example, If Fabric A has five switches with domain IDs 1 through 5 and Fabric B has five switches with the same domain IDs these two fabrics will not merge until all switches within both fabrics have a unique domain ID.

NOTE: If you use hard zoning, changing the domain ID's may affect access to devices. Hard zones are based on the domain ID and the port number.

When enabling a new configuration is highly recommended that the fabric be quiesced. Once the new zoning is enabled, a state change notification is sent to all the nodes that have registered to receive the state change. This will cause the host bus adapters to re-query and in turn might cause IO disruption based on the driver implementation.

Merging fabric together can be accomplished by simply disabling the effective configuration on one fabric, then plugging both fabrics together. The problem with this method is that once you disable the effective configuration, you open up that fabric so all servers will see all storage. Also once you plug the fabrics together, devices from the second fabric will not be accessible until you add them into the effective configuration.

To merge these two fabrics without having to disable the effective configuration for the entire fabric, it is necessary to disable at least one switch in each fabric or have a spare switch available. This will be the switch used for merging the zones and creating the new configuration. Keep in mind that there can be multiple defined configurations, but only one can be the effective or enabled configuration. The example in this document will allow merging without having to disable an active configuration.

These are some of the switch commands needed when merging fabrics.

- *`Cfgclear`*, clears all the zoning information for the ENTIRE FABRIC. Use with caution!
- *`Cfgsave`*, saves zone configurations to non-volatile memory
- *`Cfgenable`*, enables a zone configuration
- *`Cfgdisable`*, disables a zone configuration.
- *`Switchenable`*, enables a switch.
- *`Switchdisable`*, disables a switch.
- *`ConfigUpload`*, uploads the switch configuration to a host file
- *`ConfigDownload`*, downloads the switch configuration from a host file

Example Fabric Setup

FabricA has 4 switches with 3 zones. FabricB has 4 switches with 3 zones. We want to merge these two fabrics together. A good naming convention for the aliases and zones is to start them with a unique identifier for each fabric. For this example all the aliases in FabricA_CFG will start with the letter A. This way there can be a Server1 in each fabric, and they will have unique names when they merge.

FabricA_CFG

- A_Zone1
  - A_Server1
  - A_Server2
  - A_HSG80_1
- A_Zone2
  - A_Server3
  - A_Server5
  - A_HSG80_2
- A_Zone3
  - A_Server1
  - A_Server8
  - A_HSG80_3

FabricB_CFG

- B_Zone1
  - B_Server1
  - B_Server4
  - B_HSG80_3
- B_Zone2
  - B_Server6
  - B_Server7
  - B_HSG80_4
- B_Zone3
  - B_Server10
  - B_Server1
  - B_HSG80_6

A `cfgshow` on any FabricB switch will show:

```
FabricB_SW1:admin> cfgshow

Defined configuration:
cfg:   FabricB_CFG
B_Zone1; B_Zone2; B_Zone3

zone:  B_Zone1 B_Server1; B_Server4; B_HSG80_3
zone:  B_Zone2 B_HSG80_4; B_Server7; B_Server6
zone:  B_Zone3 B_Server10; B_Server1; B_HSG80_6

 alias: B_HSG80_3
     50:00:1f:e1:00:07:ae:94
 alias: B_HSG80_4
     50:00:1f:e1:00:07:bb:32
 alias: B_HSG80_6
     50:00:1f:e1:00:07:bb:91
 alias: B_Server4
     10:00:00:00:c9:22:58:77
 alias: B_Server1
     10:00:00:00:c9:22:6c:7b
 alias: B_Server7
     10:00:00:00:c9:21:84:fc
 alias: B_Server6
     10:00:00:00:c9:23:ba:b0
 alias: B_Server10
     50:00:1f:e1:00:00:08:22
Effective configuration:
cfg:   FabricB_CFG
 zone:  B_Zone1 10:00:00:00:c9:22:6c:7b
```

```
      10:00:00:00:c9:22:58:77
      50:00:1f:e1:00:07:ae:94
 zone:  B_Zone2 50:00:1f:e1:00:07:bb:32
      10:00:00:00:c9:21:84:fc
      10:00:00:00:c9:23:ba:b0
 zone:  B_Zone4 50:00:1f:e1:00:00:08:22
      10:00:00:00:c9:22:6c:7b
      50:00:1f:e1:00:07:bb:91
```

A `cfgshow` from FabricA will show:

```
FabricA_SW1:admin> cfgshow

Defined configuration:
cfg:   FabricA_CFG
A_Zone1; A_Zone2; A_Zone3

zone:  A_Zone1 A_Server1; A_Server2; A_HSG80_1
zone:  A_Zone2 A_HSG80_2; A_Server3; A_Server5
zone:  A_Zone3 A_Server8; A_Server1; A_HSG80_3

 alias: A_HSG80_1
      50:00:1f:e1:00:00:6d:11

 alias: A_HSG80_2
      50:00:1f:e1:00:00:02:f2

 alias: A_HSG80_3
       50:05:08:b2:00:1b:9b:d0

 alias: A_Server1
       10:00:00:00:c9:21:84:25

 alias: A_Server2
       10:00:00:00:c9:23:18:36

 alias: A_Server3
       10:00:00:00:c9:22:6d:68

 alias: A_Server5
       50:00:1f:e1:00:09:1d:94

 alias: A_Server8
       50:00:1f:e1:00:01:6a:61

Effective configuration:
cfg:   FabricA_CFG

zone:  A_Zone1 10:00:00:00:c9:21:84:25
      10:00:00:00:c9:23:18:36
       50:00:1f:e1:00:00:6d:11

 zone:  A_Zone2 50:00:1f:e1:00:00:02:f2
       10:00:00:00:c9:22:6d:68
       50:00:1f:e1:00:09:1d:94

 zone:  A_Zone3 50:00:1f:e1:00:01:6a:61
       10:00:00:00:c9:21:84:25
       50:05:08:b2:00:1b:9b:d0
```

This example starts with FabricA. To merge the fabrics, follow these steps:

1. Save both configurations using the `configupload` command. Saved configuration files will be FABRIC_A.TXT and FABRIC_B.TXT

   ```
   FabricA_SW1:admin> configupload
   Server Name or IP Address [host]: 10.6.6.130
   User Name [user]: administrator
   File Name [config.txt]: FABRIC_A.TXT
   Protocol (RSHD or FTP) [rshd]: ftp
   Password:
   upload complete
   FabricB_SW1:admin> configupload
   Server Name or IP Address [host]: 10.6.6.130
   User Name [user]: administrator
   File Name [config.txt]: FABRIC_B.TXT
   Protocol (RSHD or FTP) [rshd]: ftp
   Password:
   upload complete
   ```

2. If you use a spare switch make sure there is no zoning information on it before adding it into the fabric. (`cfgclear`, `cfgsave`) When you plug the spare switch into FabricA, it will automatically acquire FabricA's zone information. After plugging in the spare switch on FabricA, disable the spare switch or the switch you chose to work with.

   ```
   SANSpare:admin> switchdisable
   ```

3. ON THE DISABLED SWITCH in FabricA, use the `configdownload` command to download FABRIC_B.TXT. If there is a naming conflict, the download will not complete.

```
SANSpare:admin> configdownload
Server Name or IP Address [host]: 10.6.6.130
User Name [user]: administrator
File Name [config.txt]: FABRIC_B.TXT
Protocol (RSHD or FTP) [rshd]: ftp
Password:
zone config "FabricB_CFG" is in effect
Updating flash ...
Committing configuration...
done.
download complete
SANSpare:admin>
```

4. On the disabled switch you will see both the FabricA_CFG and the FabricB_CFG information, and the name aliases and zones from both fabrics but the FabricB_CFG has become the effective configuration.

```
SANSpare:admin> cfgshow
Defined configuration:
 cfg:   FabricA_CFG
               A_Zone1; A_Zone2; A_Zone3
 cfg:   FabricB_CFG
               B_Zone1; B_Zone2; B_Zone4
 zone:  A_Zone1 A_Server1; A_Server2; A_HSG80_1
 zone:  A_Zone2 A_HSG80_2; A_Server3; A_Server5
 zone:  A_Zone3 A_Server8; A_Server1; A_HSG80_3
 zone:  B_Zone1 B_Server1; B_Server4; B_HSG80_2
 zone:  B_Zone2 B_HSG80_3; B_Server2; B_Server6
 zone:  B_Zone4 B_Server9; B_Server1; B_HSG80_6
 alias: A_HSG80_1
      50:00:1f:e1:00:00:6d:11
alias: A_HSG80_2
      50:00:1f:e1:00:00:02:f2
 alias: A_HSG80_3
      50:05:08:b2:00:1b:9b:d0
 alias: A_Server1
      10:00:00:00:c9:21:84:25
 alias: A_Server2
```

```
      10:00:00:00:c9:23:18:36

 alias: A_Server3
      10:00:00:00:c9:22:6d:68

 alias: A_Server5
       50:00:1f:e1:00:09:1d:94

 alias: A_Server8
       50:00:1f:e1:00:01:6a:61

 alias: B_HSG80_2
       50:00:1f:e1:00:07:ae:94

 alias: B_HSG80_3
       50:00:1f:e1:00:07:bb:32

 alias: B_HSG80_6
       50:00:1f:e1:00:07:bb:91

 alias: B_SErver4
       10:00:00:00:c9:22:58:77

 alias: B_Server1
       10:00:00:00:c9:22:6c:7b

 alias: B_Server2
       10:00:00:00:c9:21:84:fc

 alias: B_Server6
       10:00:00:00:c9:23:ba:b0

 alias: B_Server9
       50:00:1f:e1:00:00:08:22

Effective configuration:
cfg:   FabricB_CFG
zone:  B_Zone1 10:00:00:00:c9:22:6c:7b
      10:00:00:00:c9:22:58:77
      50:00:1f:e1:00:07:ae:94
zone:  B_Zone2 50:00:1f:e1:00:07:bb:32
      10:00:00:00:c9:21:84:fc
      10:00:00:00:c9:23:ba:b0
zone:  B_Zone4 50:00:1f:e1:00:00:08:22
      10:00:00:00:c9:22:6c:7b
      50:00:1f:e1:00:07:bb:91
```

5. ON THE DISABLED SWITCH, disable the effective configuration.

   ```
   SANSpare:admin> cfgdisable "FabricB_CFG"
   ```

6. Bring up the web based GUI on the disable switch or use the telnet session and create your new configuration and add in the merged zones you want. For this example we will call the new configuration NEW_CFG. *DO NOT ENABLE* the NEW_CFG yet. Your NEW_CFG should look like this:

```
SANSpare:admin> cfgshow
Defined configuration:
cfg:   FabricA_CFG
A_Zone1; A_Zone2; A_Zone3

 cfg:   FabricB_CFG
B_Zone1; B_Zone2; B_Zone4

 cfg:   NEW_CFG B_Zone1; B_Zone2; A_Zone1; B_Zone4; A_Zone2; A_Zone3
zone:  A_Zone1 A_Server1; A_Server2; A_HSG80_1
zone:  A_Zone2 A_HSG80_2; A_Server3; A_Server5
zone:  A_Zone3 A_Server8; A_Server1; A_HSG80_3
zone:  B_Zone1 B_Server1; B_Server4; B_HSG80_2
zone:  B_Zone2 B_HSG80_3; B_Server2; B_Server6
zone:  B_Zone4 B_Server9; B_Server1; B_HSG80_6

alias: A_HSG80_1
       50:00:1f:e1:00:00:6d:11

 alias: A_HSG80_2
        50:00:1f:e1:00:00:02:f2

 alias: A_HSG80_3
        50:05:08:b2:00:1b:9b:d0

 alias: A_Server1
        10:00:00:00:c9:21:84:25

 alias: A_Server2
        10:00:00:00:c9:23:18:36

 alias: A_Server3
        10:00:00:00:c9:22:6d:68

 alias: A_Server5
        50:00:1f:e1:00:09:1d:94

 alias: A_Server8
        50:00:1f:e1:00:01:6a:61

 alias: B_HSG80_2
        50:00:1f:e1:00:07:ae:94

 alias: B_HSG80_3
        50:00:1f:e1:00:07:bb:32

 alias: B_HSG80_6
        50:00:1f:e1:00:07:bb:91

 alias: B_SErver4
        10:00:00:00:c9:22:58:77

 alias: B_Server1
        10:00:00:00:c9:22:6c:7b

 alias: B_Server2
        10:00:00:00:c9:21:84:fc

 alias: B_Server6
        10:00:00:00:c9:23:ba:b0

 alias: B_Server9
        50:00:1f:e1:00:00:08:22

Effective configuration:
no configuration in effect
```

7. You have the original configurations, plus the NEW_CFG that you will need to enable once you join the disabled switch back into FabricA. When you enable the switch the NEW_CFG will become part of FabricA, but will not yet be the effective configuration. Make sure that the disabled switch has a unique domain ID before you enable it. It may have changed when you did the configdownload.
8. Enable the disabled switch, and if everything is correct the fabric will not segment. In the worst case, the spare switch will segment from FabricA and give you a chance to fix any problems. All switches in FabricA will now have the new zoning information.

   ```
   SANSpare:admin> switchenable
   ```

9. On FabricA use the `cgfenable "NEW_CFG"` command to enable the NEW_CFG.

   ```
   SANSpare:admin> cfgenable "NEW_CFG"
   ```

10. Save the new configuration using the `configupload` command. (We will call this NEW_CONFIG.TXT)

Now your effective configuration will look like this:

```
SANSpare:admin> cfgshow
Defined configuration:
cfg:   FabricA_CFG

                A_Zone1; A_Zone2; A_Zone3

 cfg:   FabricB_CFG
      B_Zone1; B_Zone2; B_Zone4

 cfg:   NEW_CFG B_Zone1; B_Zone2; A_Zone1; B_Zone4; A_Zone2; A_Zone3
zone:  A_Zone1 A_Server1; A_Server2; A_HSG80_1
zone:  A_Zone2 A_HSG80_2; A_Server3; A_Server5
zone:  A_Zone3 A_Server8; A_Server1; A_HSG80_3
zone:  B_Zone1 B_Server1; B_Server4; B_HSG80_2
zone:  B_Zone2 B_HSG80_3; B_Server2; B_Server6
zone:  B_Zone4 B_Server9; B_Server1; B_HSG80_6

alias: A_HSG80_1
      50:00:1f:e1:00:00:6d:11

 alias: A_HSG80_21
       50:00:1f:e1:00:00:02:f2

 alias: A_HSG80_31
       50:05:08:b2:00:1b:9b:d0

 alias: A_Server11
       10:00:00:00:c9:21:84:25

 alias: A_Server21
       10:00:00:00:c9:23:18:36

 alias: A_Server31
       10:00:00:00:c9:22:6d:68

 alias: A_Server51
       50:00:1f:e1:00:09:1d:94

 alias: A_Server81
       50:00:1f:e1:00:01:6a:61

 alias: B_HSG80_21
       50:00:1f:e1:00:07:ae:94
```

```
 alias: B_HSG80_31
        50:00:1f:e1:00:07:bb:32
 alias: B_HSG80_61
        50:00:1f:e1:00:07:bb:91
 alias: B_SErver41
        10:00:00:00:c9:22:58:77
 alias: B_Server11
        10:00:00:00:c9:22:6c:7b
 alias: B_Server21
        10:00:00:00:c9:21:84:fc
 alias: B_Server61
        10:00:00:00:c9:23:ba:b0
 alias: B_Server91
        50:00:1f:e1:00:00:08:22
Effective configuration:
cfg:   NEW_CFG
zone:  A_Zone1 10:00:00:00:c9:21:84:25
       10:00:00:00:c9:23:18:36
       50:00:1f:e1:00:00:6d:11
 zone:  A_Zone2 50:00:1f:e1:00:00:02:f2
       10:00:00:00:c9:22:6d:68
       50:00:1f:e1:00:09:1d:94
 zone:  A_Zone3 50:00:1f:e1:00:01:6a:61
       10:00:00:00:c9:21:84:25
       50:05:08:b2:00:1b:9b:d0
 zone:  B_Zone1 10:00:00:00:c9:22:6c:7b
       10:00:00:00:c9:22:58:77
       50:00:1f:e1:00:07:ae:94
 zone:  B_Zone2 50:00:1f:e1:00:07:bb:32
       10:00:00:00:c9:21:84:fc
       10:00:00:00:c9:23:ba:b0
 zone:  B_Zone4 50:00:1f:e1:00:00:08:22
       10:00:00:00:c9:22:6c:7b
       50:00:1f:e1:00:07:bb:91
```

Matching FabricB to FabricA

Next, we must get FabricB to match FabricA. This is easier than the previous procedure. You can either use the same spare switch from FabricA or use a switch that is all ready in FabricB.

1. If you use a spare switch make sure there is no zoning information on it before adding it into the fabric. (cfgclear, cfgsave) When you plug the spare switch into FabricB, it will automatically acquire FabricB's zone information. After plugging in the spare switch on FabricB, disable the spare switch or the switch you chose to work with.

   ```
   SANSpare:admin> switchdisable
   ```

2. Use the `configdownload` command to download the NEW_CONFIG.TXT on the DISABLED SWITCH.

   ```
   SANSpare:admin> configdownload

   Server Name or IP Address [host]: 10.6.6.130

   User Name [user]: administrator

   File Name [config.txt]: NEW_CONFIG.TXT

   Protocol (RSHD or FTP) [rshd]: ftp

   Password:

   zone config "NEW_CFG" is in effect

   Updating flash ...

   Committing configuration...

   done.

   download complete

   SANSpare:admin>
   ```

3. On the DISABLED switch the NEW_CFG will be the enabled configuration, so make sure you `cfgdisable` NEW_CFG, before enabling it on FabricB. Make sure that the disable switch has a unique domain ID before you enabled it. It may have changed when you did the `configdownload`.

   ```
   SANSpare:admin> cfgdisable "NEW_CFG"
   ```

4. Enable the disabled switch. All zoning information on FabricA and FabricB should look the same except for the enable configuration. The best way to look at both fabrics is to enable logging on the telnet screen and capture the output of the `cfgshow` command and compare the files from both fabrics.

   ```
   SANSpare:admin> switchdisable
   ```

5. To complete FabricB, issue the `cfgenable "NEW_CFG"` command. Now both fabrics will be the same.

   ```
   SANSpare:admin> cfgenable "NEW_CFG"
   ```

6. Plug an ISL from FabricA into FabricB, and both fabrics will merge. If the fabrics segment then refer to the troubleshooting section.

Troubleshooting

The following section describes troubleshooting steps for isolating problems related to storage access. When initially building a SAN, lack of access either to individual storagesets or entire storage systems is not uncommon. This can usually be traced to an incorrect device setting or an inadvertent cabling or configuration setup error in the initial hardware configuration. The steps listed will assist you in isolating access problems.

1. On the server:

   a. From the server, determine if lack of access is to all of the storage (the entire storage system) or only to a portion of the storage (specific storagesets). If there is no access to only a portion of the storage system, refer to step 3.

   b. If access is not available to the entire storage system, verify from the server that the correct driver versions are loaded and that all parameters for the driver are correct. For multi-path applications, verify that the multi-path software is set up correctly.

   c. Verify that all Fibre Channel cables are plugged in and that all green indicator LEDs are on.

   d. Examine the event or error logs on the system.

2. On the Fibre Channel switch to which the server is connected:

   a. Verify the appropriate cable connection and that the port green indicator LED is on.

   b. Execute a “switchshow” command on the switch and verify that the server HBA is logged into the fabric correctly. Verify the correct port connection: F-Port, L-Port public, or L-Port private (refer to the specific HBA for more information on the correct login port types).

   F-Port: Tru64 UNIX, Compaq OpenVMS, HP-UX Fabric, Linux, Microsoft Windows NT, Windows 2000, SGI IRIX, and Sun Solaris.

   L-Port, 1 public: Novell NetWare.

   L-Port, x private, x phantom: HP-UX FC-AL.

   c. Verify all switch configuration and parameter settings.

   d. Verify that the switch is in the fabric and not segmented, “fabricshow.”

   e. Verify that all E-Ports are online.

3. On the Fibre Channel switch to which the storage is connected:

   a. Verify the appropriate cable connection and that the port green indicator LED is on.

   b. Execute a “switchshow” command on the switch and verify that the server HBA is logged into the fabric correctly. Verify the correct port connection: F-Port or L-Port private.

   F-Port: MA6000, MA/RA8000, EMA/ESA12000, EMA16000 set to FABRIC Topology.

   L-Port, x private, x phantom: MA6000, MA/RA8000, EMA/ESA12000, EMA16000 set to LOOP_HARD Topology, RA4000.

c. For MA6000, MA/RA8000, EMA/ESA12000, EMA16000:

Verify the connections to the storage system. Execute a “show connections” command at the CLI and verify that the server connection is “online.” Verify the connections are named correctly.

4. On the storage system:

a. Verify correct controller settings and configuration, “show this” and “show other.”

b. Verify that the controller ports are online and configured for the correct topology setting.

c. Verify that the storage sets are online to the appropriate controller without errors.

d. Verify that the storage sets are correctly configured and enabled for access, “show unit dn.”

e. Verify that unit offset parameters are correct. Also verify that the appropriate storage controller port is indicated in the connection name that will be accessed by the unit you have enabled.

f. Verify that the connection OS parameter type is set correctly for the operating system that is using the connection.

5. General Fibre Channel switch verification:

a. If zoning is in effect, verify that the effective zone matches the enabled zone, “cfgshow.”

b. Verify that all zone definitions are correct.

c. Verify that zoning alias names are assigned to the correct WWIDs.

d. Verify that the servers and storage being accessed are in the same zone. If zoning is in effect, the WWID must be in a zone that is in the enabled configuration or it will not have access to the fabric.

e. From the switch Web Tools GUI, examine the name server table. Verify that the appropriate WWIDs are listed and what zones they are in. Verify that the zones required are in the enabled configuration.

f. Fabric segmentation occurs when you connect together two switches or two fabrics and one of the following mismatch conditions exists between them:

- Zoning configuration mismatch
- Zoning type mismatch
- Zoning content mismatch
- Switch configuration parameter mismatches

NOTE: All switches in a fabric must have the same switch parameter settings with the exception of the following parameters:
switch name
IP address
domain ID

If you are experiencing fabric segmentation, carefully review and compare these settings in each of the two switches or fabrics.

6. QuickLoop verification:

NOTE: QuickLoop is only required for HP-UX private loop attachment.

   a. Verify that the QuickLoop license is installed.

   b. Verify that the switch ports are set to QuickLoop mode.

   c. If using QuickLoop with two Fibre Channel switches, verify that the switches are in a QuickLoop partnership.

A

Supported Products

Heterogeneous Open SAN Products

All StorageWorks by Compaq SAN topologies support a heterogeneous mix of Compaq and multi-vendor hardware platforms and operating systems, and a mix of Compaq storage system product types. This appendix provides the list of products supported in a Heterogeneous SAN. Refer to Chapter 4, "Heterogeneous SAN Platform and Storage System Rules", for configuration information and rules specific to each hardware platform and operating system version, storage products, interconnects, and interoperability for all products in the SAN.

NOTE: For the latest product support list, refer to http://www.compaq.com/storage.

Supported Operating Systems

The operating systems supported in the Compaq Heterogeneous SAN are:

- Compaq OpenVMS
- HP-UX
- IBM AIX
- Microsoft Windows NT/Windows 2000
- Novell NetWare
- Red Hat Linux and SuSE Linux
- SGI IRIX
- Sun Solaris
- Tru64 UNIX

Supported Cluster Products

- TruCluster Software Products Version 1.6 for Tru64 UNIX Version 4.0F, and TruCluster Server Version 5.0A/5.1 for Tru64UNIX Version 5.0A/5.1)
- Compaq OpenVMS
- HP MC/ServiceGuard
- Microsoft Cluster Server (MSCS)
- Novell NetWare Clusters
- Sun Clusters and the MA6000, MA/RA8000, EMA/ESA12000, EMA16000 Storage Systems
- VERITAS Cluster 1.3 with the Enterprise Virtual Array

Storage Products

The storage products supported in the Compaq Heterogeneous SAN include the Compaq entry level, mid-range, and enterprise level RAID Arrays.

The storage products supported in the Compaq Heterogeneous SAN are:

- Compaq StorageWorks Enterprise Virtual Array
- Compaq RAID Array 4000 (RA4000)
- Compaq RAID Array 4100 (RA4100)
- Compaq StorageWorks Modular SAN Array 1000 (MSA1000)
- Compaq StorageWorks Modular Array 6000 (MA6000)
- Compaq StorageWorks Modular Array 8000 (MA8000)
- Compaq StorageWorks RAID Array 8000 (RA8000)
- Compaq StorageWorks Enterprise Modular Array 12000 (EMA12000)
- Compaq StorageWorks Enterprise Modular Array 16000 (EMA16000)
- Compaq StorageWorks Enterprise Storage Array 12000 (ESA12000)
- Compaq SANworks Secure Path software
- Compaq SANworks Management Appliance
- Compaq SANworks Open SAN Manager
- Compaq SANworks Element Manager for HSV
- Compaq SANworks Element Manager for HSG
- Compaq SANworks Network View
- Compaq SANworks Resource Monitor
- Compaq SANworks Storage Allocation Reporter
- Compaq SANworks Data Replication Manager (DRM)
- Compaq SANworks Enterprise Volume Manager (EVM)
- Compaq SANworks Virtual Replicator (VR)
- Compaq SANworks Command Scripter
- Compaq SANworks SANscript
- Compaq Enterprise Backup Solutions (EBS)
- HP PV Links multi-path software (HP-UX 10.20 only)

Interconnects and Components

Host Bus Adapters

The Fibre Channel Host Bus Adapters (HBAs) supported in the Compaq heterogeneous Open SAN are:

- Compaq 380574-001 (DS-KGPSA-BC) (Emulex LP7000) Windows NT/Windows 2000, Tru64 UNIX, OpenVMS
- Compaq 168794-B21 (DS-KGPSA-CA) (Emulex LP8000) Tru64 UNIX, OpenVMS
- Compaq 176479-B21 (DS-KGPSA-CB) (Emulex LP8000) Windows NT/Windows 2000
- Compaq 380574-001 Windows NT/Windows 2000 and MSA1000
- Compaq 120186-B21 (64-bit), 223180-B21 (32-bit) Novell NetWare
- Compaq 254457-B21 (64-bit cPCI) (ACS 8.6) cPCI Sun Solaris
- Compaq 380575-001 (DS-SWSA4-SB) (JNI FC-1063) (32-bit) Sbus Sun Solaris
- Compaq 123503-001 (DS-SWSA4-SC) (JNI FC64-1063) (64-bit) Sbus Sun Solaris
- Compaq 380576-001 (DS-SWSA4-PC) (JNI FCI-1063) (32-bit) PCI Sun Solaris
- Compaq 167433-B21 (QLogic QLA2200F/66) Red Hat, SuSE Linux x86/Alpha
- Compaq 218409-B21 (QLogic QLA220F/66) HP-UX
- Compaq 197819-B21 (Cambex PC1000F) Fabric IBM AIX
- HP A3404A FC-AL HSC Bus, K Class HP-UX
- HP A3591A/B FC-AL HSC Bus, D Class HP-UX
- HP A3636A FC-AL HSC Bus, T Class HP-UX
- HP A3740A FC-AL PCI, L Class HP-UX
- HP A5158A FC-AL PCI A/L/V/N Class HP-UX
- QLogic 2200F/66 SGI PCI- FC-1POPT PCI IRIX
- QLogic 2200F/66 SGI XT-FC-1POPT XIO IRIX

NOTE: The Compaq branded HBAs listed above are supplied with Compaq specific drivers and firmware. These are the only HBAs supported for the specified operating systems.

Fibre Channel Switches

The Compaq Fibre Channel switches supported in the Compaq Heterogeneous SAN are:

- Fibre Channel SAN Switch 8
- Fibre Channel SAN Switch 16
- Fibre Channel SAN Switch 8-EL
- Fibre Channel SAN Switch 16-EL

NOTE: The Fibre Channel SAN Switch 8-EL comes with a single E-Port license. It is possible to upgrade the 8-EL to multiple E-Port licenses.

- FC-AL Switch 8

Gigabit Interface Converters (GBICs)

- Short-wave (850nm)
- Long-wave (1310nm)
- Very Long Distance (1550nm)

Gigabit Link Module (GLM)

- Short-wavelength (850nm)

SAN Interfaces/Fiber Optical Cables

- Fibre Channel connection via 50/125 or 62.5/125 multi-mode, and 9/125 single-mode fiber optic cable.

Data Replication Manager (DRM) Interfaces/Transports

- Fibre Channel connection via 50, 62.5 micron multi-mode or 9 micron single-mode fiber optic cable
- Fibre Channel connection via Wave Division Multiplexing (WDM)
- ATM over a single T1/E1 Wide Area Network (WAN)
- ATM over multiple T1/E1 WAN (Inverse Multiplexing)
- ATM over T3/E3 WAN
- ATM over fractional and/or shared T3/E3 and OC3 WAN
- IP over 10/100 copper based Ethernet
- IP over 1 Gbps Optical Ethernet

Heterogeneous SAN Product Support

Tables that list the currently available storage products supported in a Heterogeneous SAN are provided in the document “Supplemental Tables for the Heterogeneous Open SAN Design Reference Guide.” Access this table through this link:

Supplemental Tables

(http://www.compaq.com/products/storageworks/san/documentation.html)

Glossary

This glossary defines terms used in this guide or related to this product and is not a comprehensive glossary of computer terms.

arbitrated loop

See FC-AL

Asynchronous Transfer Mode (ATM)

Communications networking technology for LANs and WANs that carries information in fixed-size cells of 53 bytes (5 protocol and 48 data)

controller pair

Two interconnected controller modules which together control a disk array

Corporate Fabric

A SAN fabric using Compaq SAN Switch 8, 16, 8EL, and 16EL model switches. A Corporate Fabric can also include the SAN Switch Integrated 32/64 model switch.

Director Fabric

A SAN Fabric using Compaq Director 64 model switches

Enterprise Virtual Array

The StorageWorks Enterprise Virtual Array is a high performance, high capacity, and high availability storage solution for the high-end enterprise class marketplace. Each Enterprise Virtual Array storage system consists of a pair of HSV virtualizing storage controllers and the disk drives they manage.

Enterprise/Modular RAID Array

Storage system based on an HSG60 or HSG80 controller. These systems include MA6000, MA8000, RA8000, EMA12000, EMA16000, and ESA12000 storage systems

Entry-Level Fabric

A SAN fabric using Compaq C8 model switches

Fibre Channel Arbitrated Loop (FC-AL)

A Fibre Channel topology that links multiple ports (up to 126) together on a single shared simplex media

gigabit interface converter (GBIC)

The hardware devices inserted into the ports of the Fibre Channel switch that hold the Fibre Channel cables. GBIC devices are available for short-range applications (0.5 to 500 meters), long-range applications (up to 10 km), and very long distances (up to 100 km)

hop

One or more interswitch links between a pair of Fibre Channel switches

Host Bus Adapter (HBA)

An adapter used to connect the host server to the fabric

inter-switch Link (ISL)

A fibre cable connecting a port on one switch to a port on another switch

Selective Storage Presentation (SSP)

This feature provides the ability to restrict access to a given Fibre Channel LUN.

Storage Area Network (SAN)

A high-speed network that interconnects different kinds of data storage devices with associated data servers on behalf of a larger network of users, typically using Fibre Channel technology

Wavelength Division Multiplexing (WDM)

The technique of placing multiple optical signals on a single optical cable simultaneously. Dense wave division multiplexing (DWDM) places many signals on a cable. Coarse wave division multiplexing (CWDM) places only a few signals on a cable.

Index