

SAN Fundamentals

How Fibre Channel SANs are Built, Secured and Managed

Tim Thomas

Systems Engineering Team

Storage Group

Sun Microsystems

May 18, 2006

© 2005 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, CA 95054 USA

All rights reserved.

This product or document is protected by copyright and distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this product or document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any. Third-party software, including font technology, is copyrighted and licensed from Sun suppliers.

Parts of the product may be derived from Berkeley BSD appliances, licensed from the University of California.

Sun, Sun Microsystems, Sun StorEdge, the Sun logo, are trademarks, registered trademarks, or service marks of Sun Microsystems, Inc. in the U.S. and other countries.

UNIX is a registered trademark in the United States and other countries, exclusively licensed through X/Open Company, Ltd.

Windows is a registered trademark of Microsoft Corporation in the United States and other countries.

All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

The OPEN LOOK and Sun's Graphical User Interface was developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a non-exclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUIs and otherwise comply with Sun's written license agreements.

RESTRICTED RIGHTS: Use, duplication, or disclosure by the U.S. Government is subject to restrictions of FAR 52.227-14(g)(2)(6/87) and FAR 52.227-1987), or DFAR 252.227-7015(b)(6/95) and DFAR 227.7202-3(a). DOCUMENTATION IS PROVIDED AS IS AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS HELD TO BE LEGALLY

Table of Contents

Introduction.....	1
What is Fibre Channel ?.....	1
How Fast is Fibre Channel ?.....	1
What is a Storage Area Network ?.....	1
What About IP Based SANs ?.....	1
Fibre Channel Switches and Fabrics.....	2
What is a Fabric?.....	2
Server Host Bus Adapters.....	3
Fabric Name Server.....	3
World-Wide Names.....	3
Fabric Login.....	3
Principal Switch.....	4
Fibre Channel Registered State Change Notifications.....	4
Inter-Switch Links	5
SAN Access Control.....	6
Fibre Channel Switch Zones.....	6
Fibre Channel, SCSI LUNs and LUN Masking.....	6
Using Zoning and LUN Masking Together.....	6
Extending SANs with Inter-Switch Links.....	7
Flow Control and Buffer Credits.....	8
Joining SANs with Fibre Channel Routers.....	9
Server Operating Systems and SANs.....	11
Sun StorEdge SAN Foundation Software.....	11
Multipathing.....	12
Booting Servers Off SAN Attached Storage.....	13
SANs and Tape.....	14
Tape Library Sharing.....	16
Static Tape Library Sharing.....	16
Dynamic Drive Sharing.....	16
Managing SANs.....	17
Challenges of SAN Management.....	17
SAN Management Applications.....	17
Summary.....	20

References..... 21

Further Resources..... 22

Introduction

This document discusses the core technologies that are used to build and secure Fibre Channel based Storage Area Networks.

First, let us discuss a few basic topics:

What is Fibre Channel ?

Fibre Channel is a flexible standards based networking architecture that can be used as a transport mechanism for a number of Upper Level Protocols. The most common Upper Level Protocols are TCP/IP and SCSI.

Fibre Channel is a Serial Full Duplex protocol. It has sophisticated flow control allowing it to be extended over long distances.

How Fast is Fibre Channel ?

Fibre Channel first became available at 1 Gbit, but by the time SANs became popular 2 Gbit Fibre Channel was shipping and this is the speed most commonly used. 4 Gbit Fibre Channel products are now available and 8 Gbit products are under development. 4 Gbit equipment is, and 8 Gbit will be, backwardly compatible with 2 Gbit equipment.

10 Gbit Fibre Channel was a subject of wide discussion within the industry a few years ago and some 10 Gbit products were released, but it is not backwardly compatible with earlier Fibre Channel equipment so its usage is limited.

What is a Storage Area Network ?

Within this document the term Storage Area Network (SAN) is used to describe a Fibre Channel network over which SCSI-FCP (SCSI over Fibre Channel) protocols run.

What About IP Based SANs ?

The term SAN is now commonly extended to include iSCSI based solutions which run over a LAN/WAN infrastructure, not Fibre Channel. Reference (v) provides detailed information on iSCSI and IP SANs and positions them relative to Fibre Channel SANs.

Fibre Channel Switches and Fabrics

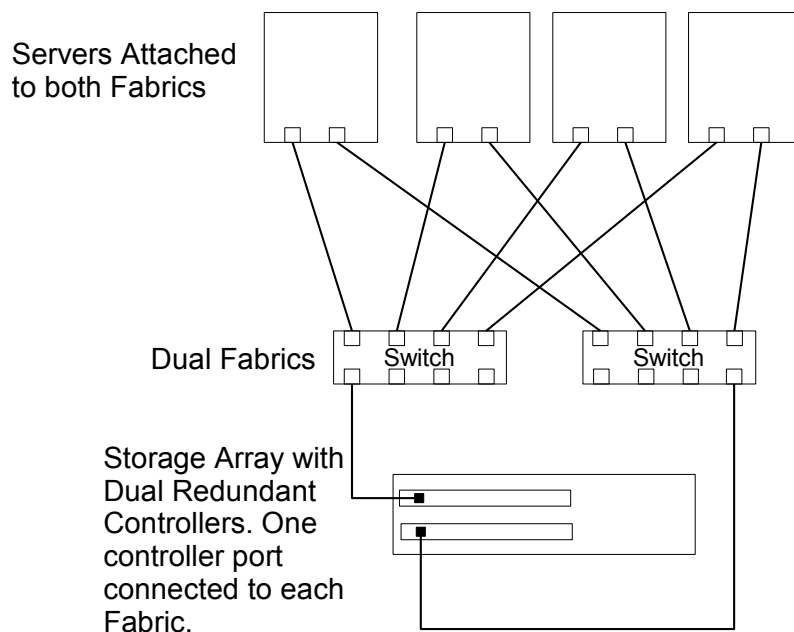
What is a Fabric?

A Fibre Channel switch on its own, or two or more switches connected together, forms a Fibre Channel Fabric. A switch can only be in one Fabric. Some large switches can now be partitioned into what are effectively multiple switches, but each partition is still in only one Fabric.

The primary function of a Fabric is to correctly route traffic from one port to another, where the source (initiator) and destination (target) ports may or may not be connected to the same switches but must be in the same Fabric (unless Fibre Channel Routers are used, see later section).

SANs usually comprise of at least two separate Fabrics for redundancy. Servers and storage have connections to both Fabrics for resilience and server side software multipaths between the paths.

The below figure shows four servers attached via redundant Fabrics to a storage array.



Server Host Bus Adapters

To be able to connect to a SAN a server requires one or more Fibre Channel Host Bus Adapters (HBAs) to be installed and configured. Fibre Channel HBAs handle all the processing of the Fibre Channel stack using onboard ASICs.

Fabric Name Server

The Fibre Channel standard defines a name service called the Fabric Simple Name Server. The "Simple" is commonly omitted. The Fabric Name Server stores information about all the devices in the Fabric. An instance of the name server runs on every Fibre Channel switch in a SAN.

World-Wide Names

All Fibre Channel devices have a unique IEEE issued 64 bit World-Wide Name (WWN). This is used to identify devices in a SAN.

Some devices allow their WWN to be changed (e.g. some Fibre Channel tape drives) and this can be useful in certain service situations but, as with changing MAC addresses on Ethernet Network Interface Cards, this needs to be done carefully.

Fabric Login

When a device (e.g. server HBA port, array controller port, Fibre Channel tape drive port) is connected to a Fibre Channel switch it goes through an initialization process called a Fabric Login.

1. The switch registers the device's WWN with the Fabric Name Server which dynamically generates a 24 bit Port Address for the device which is unique within the Fabric.
2. If the device is an initiator (typically a server HBA port) it will then attempt to login to every other device it can see in the Fabric to find out what it is. This process is required so that the server can discover the characteristics of the devices and so "connect" them correctly to its operating system.

The Port Address assigned to a device describes the location of the device in the SAN. It has information that identifies the switch and switch port the device is connected. The information provided by the Port Address enables traffic to be routed very efficiently across the Fabric.

Principal Switch

In a Fabric one of the switches is the Principal Switch. It is responsible for keeping all of the Name Server instances in the Fabric up to date. The decision about which switch in a Fabric is Principal can be forced by the administrator, or the switches can be left to decide for themselves using built-in logic. If the Principal Switch fails another will take over the role automatically.

If a new device is connected to a Fabric the Name Server on the Principal Switch is responsible for generating a Port Address and notifying all of the other switches. Similarly, if a device is removed from the Fabric then the Name Server on the Principal Switch is notified and it then sends notifications to all of the other switches.

If a new switch is connected to a Fabric it exchanges information with the Principal Switch and vice-versa, and then the Name Server on each of the other switches in the Fabric is updated. A switch joining a stable Fabric cannot become the Principal.

Fibre Channel Registered State Change Notifications

Registered State Change Notifications (RSCNs) are the mechanism by which changes in a Fabric are propagated.


RSCNs are generated by a switch when anything changes e.g. a device connects or disconnects. RSCNs are only sent to devices that have registered with the Fabric to receive them. Only Initiators register to receive RSCNs, target devices do not. Server HBA ports are initiators, ports on arrays (unless the array is replicating out of the port) and tape drives are targets.

Some events, such as enabling or disabling a switch, will cause RSCNs to be sent to all devices in a Fabric that have registered to receive them. Lesser events, such as a server rebooting, cause RSCNs to only be sent to devices in the same zone(s) as the device.

Historically, when a Fibre Channel device received an RSCN it logged out of the Fabric and then logged back in to the Fabric so that it could query the Fabric Name Service to find out what devices it could now see following the changes. This was a disruptive process and was known to cause problems with Fibre Channel tape drives. The latest switch, HBA and tape firmware complying with the latest Fibre Channel standards for error recover (FCP-2) minimise the disruption caused by RSCNs and can often handle processing an RSCN in a manner that is not disruptive to I/O.

Inter-Switch Links

Connections between switches are called Inter-Switch Links (ISLs). When switches are connected together with one or more ISLs management information is exchanged and the switches are then in the same Fabric.



Fabrics are expanded in this way to increase the number of ports in a Fabric or to extend the Fabric between buildings. Merging two live Fabrics together needs to be planned carefully. When Fabrics are merged one switch will become the Principal switch for the new larger Fabric and new Port Addresses may be given to some devices, this should not create any issues but is best done at a quiet time.

Switch ports are universal, so any port on a switch can be used for an ISL.

SAN Access Control

Fibre Channel Switch Zones

A zone is made up of a number of devices grouped by their World Wide Names (WWN), or is a group of switch ports.

Zoning information is managed and stored in the Fabric Name Server and is communicated to new switches when they join a fabric.

Devices can only see other devices in the same zone, so zones enable servers and the storage devices they use to be isolated from other servers and their storage devices.

Zones can overlap i.e. A switch port can be in more than one zone and a device WWN can be in more than one zone.

Zones can span multiple switches i.e. They can contain device WWNs or ports from anywhere in a Fabric.

For an overview of zones and zoning strategies see reference (i).

Fibre Channel, SCSI LUNs and LUN Masking

What is seen as a LUN today from a servers point of view is more often than not actually a slice, or partition, of a RAID set created within an array that is presented to the SAN by the array RAID controller.

An array presents a Fibre channel port or ports to a SAN and the slices are then presented on the ports using the SCSI-FCP protocol and the servers see them as LUNs. Fibre channel switches see the controller port and its WWN but do not “speak SCSI” and so cannot control access to the slices presented on a port on an individual basis.

As many servers may be sharing an array controller port in a SAN it is necessary to have a way of controlling access to the slices presented to the SAN and LUN Masking provides that.

LUN Masking is a function of the array controller. When LUN Masking is enabled each slice in an array has an access control list which is a list of the World Wide Numbers of the HBAs that are allowed to access it. This access control list is defined by the storage administrator using array management tools.

When a server scans the SAN for storage devices it will only see those slices that it has been granted access to.

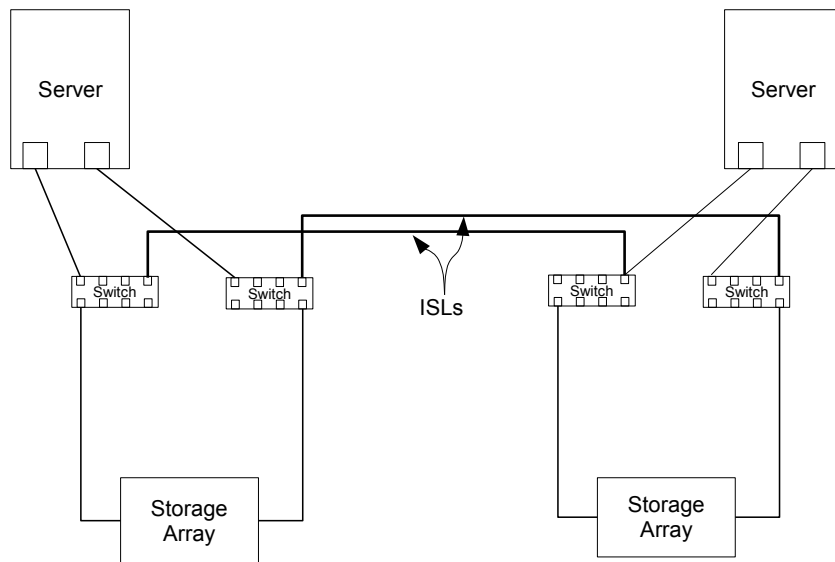
Since most servers will be connected to a SAN with at least two HBA ports via two separate Fabrics access to a slice/partition will have to be granted to both of the server HBA port WWNs.

Using Zoning and LUN Masking Together

Zoning and LUN masking are complementary. You Zone to a device port or WWN and then, where required, use LUN masking to control access to individual slices presented by the array port(s).

Extending SANs with Inter-Switch Links

The below figure shows an extended or stretched SAN. This comprises two Fabrics, each with a switch at both sites. A single ISL is shown between the switches in each fabric, but more can be used to improve availability and performance.



This infrastructure could be used to simply share devices between two buildings or sites, for a stretched Cluster, or with arrays that can replicate data over a SAN for Disaster Recovery, or for all three at the same time.

ISLs can be anything up to few hundred metres long with no special equipment required. By using specialist equipment ISLs can be extended over long distance fibre links or over WAN links by tunnelling Fibre Channel over IP, the second option being very popular when replicating between two storage arrays.

Reference (iv) discusses the distance limitations of Fibre Channel at various speeds along with details of cabling and connectors. It is also important to consult the documentation of the specific switches being used to find out what distances are supported between switches at what speeds.

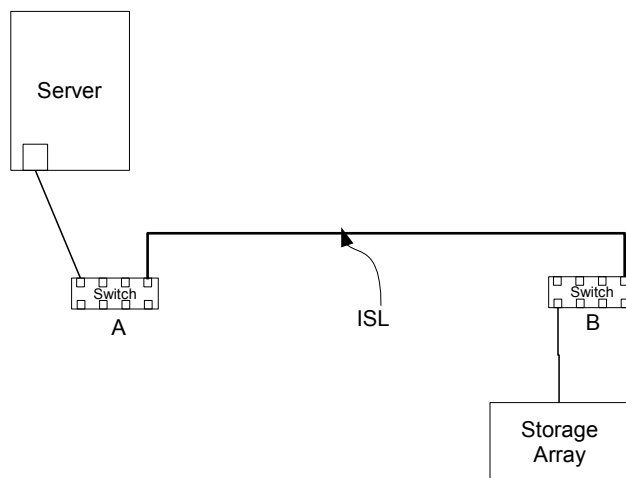
Extending SANs over WAN's is discussed briefly in a later section.

Flow Control and Buffer Credits

Fibre Channel switches have sophisticated flow control mechanisms to maximise throughput. This is based on a system of Buffer Credits. Depending on the vendor, a switch may have a pool of Buffer Credits for all of its ports, or a number of Buffer Credits assigned to each port.

In a multi-switch Fabric every switch will know how many buffers the switches it is connected to has available, as will the server HBAs. Exchanging this information is all part of the Fabric initialization process.

Consider an example where two switches (A & B) are connected together with an ISL and a server connected to switch A is writing to a storage array connected to switch B.



Switch A sends frames (similar to packets in ethernet) across the ISL to Switch B. Switch B receives each frame and immediately routes it on to its destination sending acknowledgements back down the fibre (Fibre Channel is Full Duplex) to Switch A.

Once Switch B has acknowledged receipt of a frame Switch A can clear the buffer that was holding that frame locally, which enables it to receive more frames from the server.

The longer the ISL, the more unacknowledged frames that will be in flight down the fibre at any given time. If all of a Switch A's buffer space is full of unacknowledged frames latencies will build up in the SAN: The ISL becomes a bottleneck. Your switch manual will tell you how many buffer credits are available per port and the maximum distances supported between ports at various speeds before you are in danger of running out of Buffer Credits.

SANs are often termed non-blocking. This is not because they have infinite amounts of bandwidth but because data is never sent somewhere where it does not have a place to go. Fibre Channel networks will not block, but they can become congested.

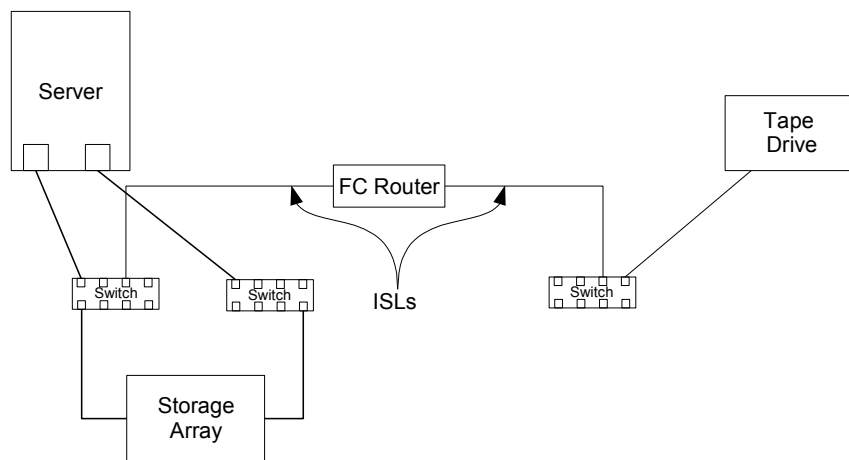
Joining SANs with Fibre Channel Routers

When extending a SAN with ISLs the Fabrics span sites. As previously discussed, if two existing Fabrics are joined a major Fabric re-configuration takes place and no I/O is done whilst it completes. As a one-off event this can be managed, but once the Fabrics are joined any event generating an RSCN which all elements in the Fabric need to be notified of may generate significant management traffic across the ISLs. More significantly, with the elements being potentially a long way apart, the time required for the Fabric re-configuration to complete could lead to server I/O requests timing out.

By careful SAN design and planning these problems can be minimized, however Fibre Channel Routers offer an elegant way of sharing devices between Fabrics without merging them.

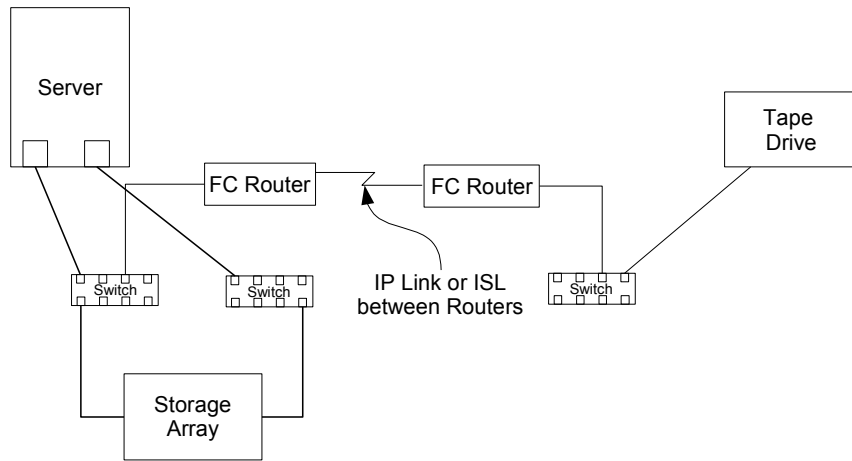
Fibre Channel Routers are used to map only the specific devices that need to be shared between Fabrics. If a device mapped into one Fabric from another has a Port Address already used within that Fabric the Router takes care of any address translation required. The Fabrics are not merged and RSCNs will remain within the local Fabrics.

In the example below a server in one Fabric needs to back up to a Fibre Channel tape drive in another. Rather than connect the Fabrics together we use a Fibre Channel Router to map the drive from one Fabric to the other.



The Router looks like just another Fibre Channel Switch to the Fabrics it is connected to, and the devices it is routing for appear as if connected directly to the local Fabrics.

Most Fibre Channel Routers have the capability of routing over Fibre from Router to Router and also of tunnelling Fibre Channel from Router to Router over IP, so we can extend a SAN as shown below.



In summary, with SAN Routers you can share devices between Fabrics and get the benefits of an extended SAN without the potential vulnerabilities and management issues that simply extending a Fabric with ISLs might cause.

Server Operating Systems and SANs

At a low level within a server's operating system, the driver for the Fibre Channel Host Bus Adapter(s) will present the LUNs and tape drives it discovers in the SAN as SCSI devices by decoding the SCSI commands from the Fibre Channel frames.

In most Operating Systems the devices are then managed via the standard SCSI driver stack and are handled no differently than direct attached SCSI devices, but in the case of Solaris an enhanced driver stack has been developed.

Sun StorEdge SAN Foundation Software

Since Solaris 8 (4/01), Sun has offered a very high level of integration with Storage Area Networks.

The Sun StorageTek™ SAN Foundation Software (SFS) is used by Solaris to discover SAN attached devices and includes commands to create and delete device nodes for those devices. It is also known as Leadville.

The SFS is available as patches for Solaris 8 and Solaris 9 on the SPARC platform and forms an integral part of Solaris 10 for SPARC and x86/x64 platforms.

The SFS is only supported with Sun HBAs.

See reference (ii) for more on the SFS.

Multipathing

Because the use of multipathing is now so common, server Operating Systems increasingly have their own native multipathing functionality. Examples are Microsoft's MPIO and the Sun StorageTek™ Traffic Manager (SSTM) software for Solaris, also known as MPxIO. Storage vendors now qualify with these native multipathing solutions.

The benefit to customers is that they get multipathing software with the server's operating system at no additional charge, so they no longer need to buy and maintain 3rd party multipathing software. This reduces costs and simplifies support.

SSTM is provided as patches for Solaris 8 (4/01) and later and Solaris 9 on the SPARC platform. It is integrated into Solaris 10.

SSTM works with Fibre Channel attached storage (Direct Attached or SAN attached) and in Solaris 10 (1/06) support for iSCSI was added.

Booting Servers Off SAN Attached Storage

Booting a server off SAN attached storage requires that the server be able to see its boot devices across the SAN even before the operating system is loaded. This is achieved by storing the boot LUN's WWN in the Boot PROM of the server or in the HBA.

If you wish to boot a server off a SAN and you have purchased your server and storage from the same company then speak to them about your requirements.

If you have purchased your server and storage from different companies you should initially contact the storage vendor about how to approach this as it is usually they who test SAN booting of servers from various suppliers off of their storage and only certain combinations of HBAs and switches will have been tested.

SANs and Tape

Storage Area Networks and tape drives with Fibre Channel interfaces allow flexible high performance backup solutions to be built.

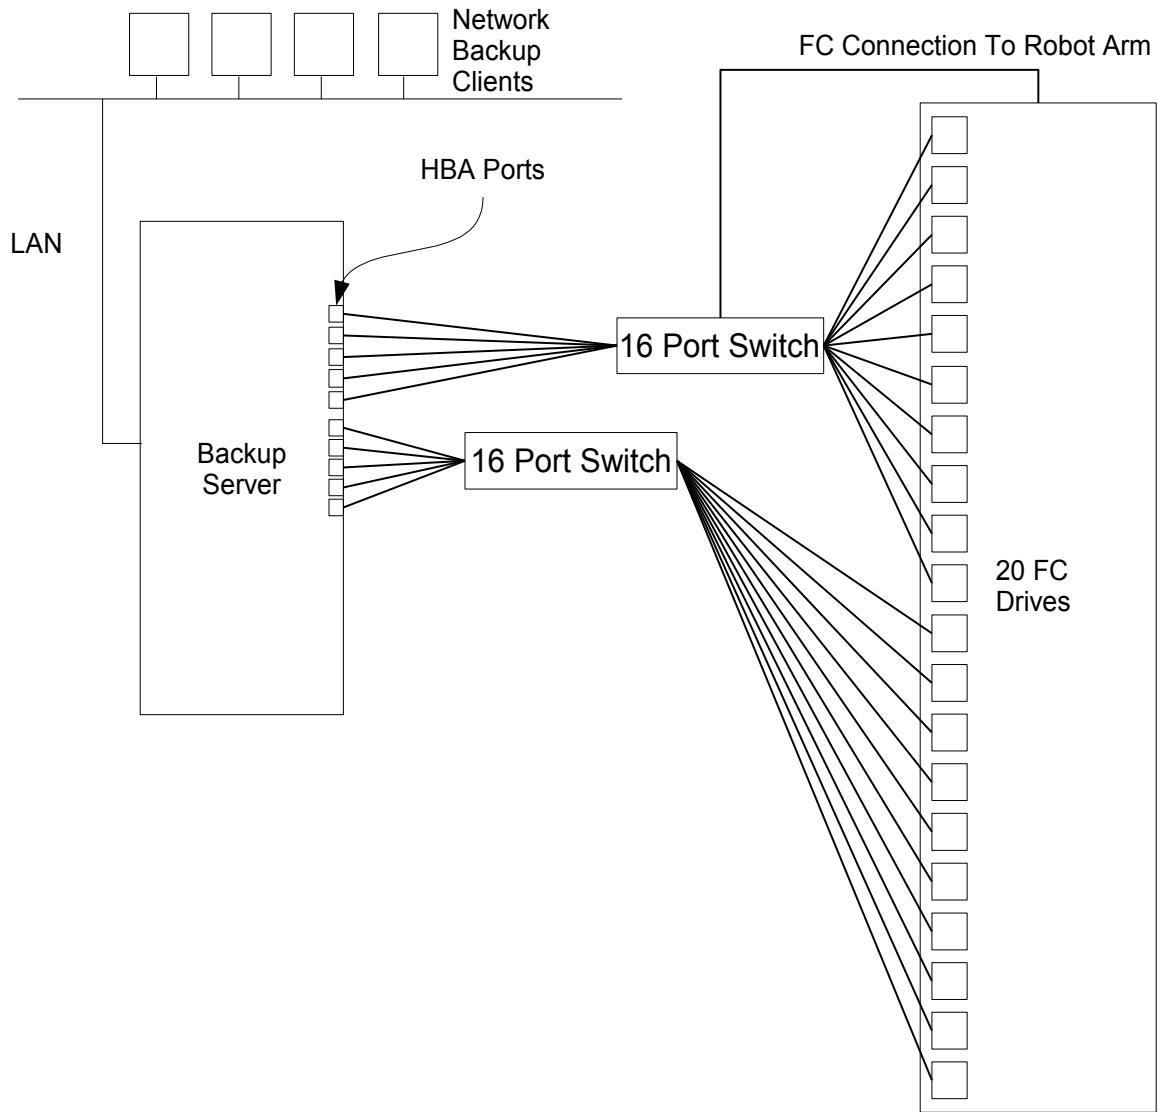
Zoning is used to secure tape drives to server HBA ports in the same way as it can be used with ports on array controllers. LUN masking is not relevant to tape drives and there is no multipathing for tape drives.

Array controller ports and tape drive ports can be accessed from the same HBA port on a server and can be in the same zone. That being said, it is an industry best practice to keep disk and tape in separate zones. This practice evolved to protect the tape drives from spurious RSCNs and also to ensure best performance.

Whilst RSCNs are now less of a concern the sense of the second point becomes apparent when you consider that three LTO-2 Fibre Channel drives can consume all the data that a 2 Gbit HBA port can send to them.

SAN based backup solutions can very quickly use up a lot of switch ports as each drive must be individually connected to the SAN as must the HBA driving them. In many cases the tape library robot arm requires its own switch port also.

The following example shows a single backup server connecting to 20 Fibre Channel tape drives with 2 drive ports per HBA port. This uses up 31 SAN ports.



Tape Library Sharing

Backup software allows tape libraries to be shared between multiple servers. This sharing of tape libraries may be done in two ways, the second of which is only possible using a SAN.

Static Tape Library Sharing

Multiple servers share a tape library and the number of drives available to each server can be static.

To do this the SAN is zoned so that each tape drive in the library is visible to only one server.

A master server controls the library robot. This master server may or may not have access to tape drives. The other servers are slaves so far as controlling the robot goes. The slave servers send tape cartridge mount and dismount requests to the library via the master using the LAN. The slave servers transfer data directly to the tape drives once the cartridges are loaded. This can be implemented with copper SCSI and Fibre Channel based libraries.

This approach allows a tape library to be shared but could lead to some drives being under or over utilized.

Dynamic Drive Sharing

This is used where multiple servers share a tape library, and where some servers require very large amounts of tape drive resources at certain times, but where it would be useful to share those resources differently at other times. e.g. large application on an enterprise server requires ten drives for overnight backup to complete in time, but it is desirable that those drives can be allocated two each to five smaller servers once that backup is complete.

To do this the SAN is zoned so that each tape drive in the library is visible to multiple servers.

Master and slave servers work much as before however the master server now additionally determines which servers can use which drives at what time based on backup schedules and policies.

This allows a tape library to be shared very flexibly between multiple servers but the SAN design for this can be complex and can consume many switch ports and server HBA ports.

Managing SANs

Challenges of SAN Management

A SAN must be managed at multiple points:

- Server OS
- HBA
- Switch
- Storage Array/Tape Library

The biggest challenge is understanding the dependencies of the various elements upon each other, as well as having to learn multiple tools. For example, a HBA has to be swapped on a server. The new HBA has a different WWN than the old one and so zoning information in the SAN and LUN masking information on the arrays has to be updated or a switch needs to be replaced – who should be told ? Which applications will be affected ?

Small SANs can be managed effectively using the individual element managers of the switches, storage arrays and tape libraries along with server OS features however as SANs grow the amount of knowledge required to make a change can become significant.

In a large SAN, relatively simple operations require a detailed knowledge of the environment. Skilled operators, good documentation and rigorous change control can mitigate the potential issues that might arise when managing large SANs, however such processes can be expensive to implement in people time and can reduce the agility of an organization. A new breed of management tools have been created to address this.

SAN Management Applications

Ideally a SAN Management Application should be able to perform management functions such as Switch Zoning, Array LUN Management and Full Path Management (e.g. creating a LUN and mapping it to a server HBA port taking care of any re-zoning and LUN masking on the way) via a standard interface. Features like this are collectively known as Provisioning.

The SAN Management Application should normalize the management of the underlying elements so that administrators can use the same GUI (or CLI) regardless of the brand of storage, switch, HBA and server. The benefits of this are huge as operators need only learn one tool for day-to-day work and in understanding the relationships between the various elements the SAN Management Application can reduce problems due to “driver error”.

Until the recently SAN Management Applications relied on multiple vendors proprietary management APIs to manage the various elements. It soon became clear that API swapping and licensing between various vendors was not a long term solution and the Storage Networking Industry Association (SNIA) started the Storage Management Initiative and developed the Storage Management Initiative Specification, known as SMI-S, which is still evolving today.

SMI-S defines a set of standard management APIs that allow a SAN Management Application to manage HBAs, Switches and Storage Arrays in a normalised fashion. The Storage Management Application uses SMI-S management interfaces to discover and manage elements.

The software component that interfaces between the SAN Management Application and the managed element is called a Provider. It is the Provider that is SMI-S compliant. The vendor of the element being managed usually writes the Provider and they map SMI-S operations to their proprietary API in the Provider.

Providers can be implemented in a number of places:

1. The Provider for an element can be integrated into the Storage Management application.
2. The Provider is in the device e.g. On the service processor of an array or inside a switch.
3. The Provider takes the form of middleware (known as Proxies) that runs on an intermediate host and manages the dialogue between Storage Management application and the element.

An elements IP address or DNS name, along with security information (login information), are required to be able to discover it. If a proxy is used then this information is required for the proxy, not the element.

The picture would be incomplete without being able to discover servers. Discovering a server usually requires that an application specific agent be first installed on the server to get all of the required information.

Without an agent installed the Storage Management Application cannot find out information about filesystems, files, applications or the server itself. An exception is Microsoft Windows. Microsoft Windows exposes some information through the Windows Management Instrumentation (WMI) interface however the information is not as complete as that which an Storage Management Application's agent would produce especially around mapping HBA's, volumes and filesystems together.

So SAN Management Applications can be run agentless, but the information available from servers will be limited. Other than with Microsoft Windows, all that will usually be discovered is the HBA ports via the switches and they would need to be manually grouped to form a "virtual server". All SAN Management tools have a grouping function since without a host agent the Storage Management application has no idea that HBAs are in the same server. Once all the necessary agents and Providers are in place the SAN Management Application discovers details of all the elements on the SAN and stores them in a database. The application can then determine the relationships between the various elements from the information it has extracted and can validate changes against the database before implementing them on the actual elements.

Revisiting the earlier example of a HBA change, we now have the capability to lookup every point in the SAN where the old WWN was used and change it to the new one or have the tool do it for us via a wizard type interface.

For more on the SNIA and the Storage Management Initiative see Appendix B.

The Sun StorageTek™ Operations Manager software is a suite of SMI-S compliant SAN Management products. See reference (iii) for more on this software. Software such as this does not totally remove the need to use of individual element managers as the management standards do not yet extend to the management of features like array-to-array replication or snapshots (though work on APIs for these kinds of features is in progress) however they will perform the bulk of day-to-day SAN management tasks quickly and efficiently as well as handling SAN health and event management from a single application.

Summary

Fibre Channel is a well-established, widely deployed technology with a proven track record. This paper has introduced the reader to the the basics of Fibre Channel and has explained how Fibre Channel SANs are built, secured and managed.

In the Further Resources section of this document is a list of web sites and books that are a starting point for additional research into Fibre Channel SANs. Reference (v) provides an introduction to iSCSI and IP based SANs.

References

- i. Zoning Implementation Strategies for Brocade SAN Fabrics
http://www.brocade.com/san/white_papers/pdf/Zoning_Imp_WP_00.pdf
- ii. Increasing Storage Area Network Productivity
<http://www.sun.com/blueprints/0704/817-7544.pdf>
- iii. Sun StorageTek Operations Manager
http://www.sun.com/storage/software/storage_mgmt/esm/
- iv. SNIA Europe Technology Guide: Fibre Channel Fibres, Speeds and Distances
http://www.snia-europe.com/admin/cmfiles/miscellaneous_ifiles/TGFibreChannel.pdf
- v. IP SAN Fundamentals
SunWIN token #472865

Further Resources

- i. Storage Networking Industry Association
<http://www.snia.org>
- ii. SNIA Storage Management Initiative
<http://www.snia.org/smi/home/>
- iii. Fibre Channel by Alan F. Benner. Published by McGraw-Hill. ISBN: 0-07-005669-2