



**Storage Area Networks
Data Security & Fabric Management**

White Paper

**Product Management
3/13/02**



Table of Contents

Introduction	3
Security and information storage systems	3
SAN design and management security	3
SAN data integrity and fabric management	3
Data integrity security concerns	4
Data access and integrity	4
Questions concerning data access and integrity security	4
Fabric management security concerns	5
Fabric-level security	5
Questions concerning fabric-level security	5
Data protection and fabric management	7
Security	7
Data integrity and security	8
Fabric zoning	8
What is zoning?	8
Types of zoning	8
Zoning configurations and components	8
Zone construction	9
What is a zone member?	10
Zone member identification	10
What is a zone set?	11
Zoning limitation	11
LUN masking	12
Persistent binding	13
Fabric management and protection	14
Fabric-to-fabric security technologies	14
Host-to-fabric security technologies	14
Management-to-fabric technologies	14
Configuration integrity technologies	14
Conclusion	15



Introduction

Security and information storage systems

When the word security is used in association with information storage systems in storage area networks (SANs), the first thought that comes to mind is of computer hackers infiltrating those information storage systems and causing problems. Unauthorized human access to information storage systems is a big concern, but there is yet another security issue that must also be addressed and that is technology containment. For example, in an environment where technologies are not contained, Windows NT server(s) would see every available LUN and try to take ownership of them all. In brief, technology containment keeps the servers from gaining unauthorized access or accidental access to certain resources in the SAN.

SAN design and management security

A basic advantage of SAN technology over storage systems that were previously captive or directly attached to their associated hosts, is that SANs make possible networked any-to-any connectivity. In a SAN, storage systems are de-coupled from their hosts and shared on a network where many hosts can potentially access them; making security an essential component in the design and management of storage systems based on SAN architectures.

SAN data integrity and fabric management

Open systems offer many different file systems, volume management and disk management formats and software, demanding that security be considered and implemented in the SAN for the following reasons:

- Data access and integrity
- Fabric management and protection from outside threats

The following discussion poses questions concerning data access and integrity, fabric management and protection; then presents technologies and methodologies that provide answers to those questions.



Data integrity security concerns

Data access and integrity

The first concern when implementing a SAN is to provide a higher level of data access and data integrity than direct-attached storage systems provide. This requires new technologies and security methods over those typically available for direct-attached storage systems. Frequently, multiple technologies and security methodologies are applied to any single SAN implementation to maximize data access and integrity. For example, a shared RAID-array is a system that presents logical disks or LUNs to hosts on the SAN for use. Without security measures, any host on the SAN fabric would be able to see those LUNs and potentially write data to them; in a scenario where data would be stored on those LUNs, data corruption or data loss would occur.

Questions concerning data access and integrity security

Concerning data access and integrity on a SAN, consider the following questions:

How can we segregate operating systems at the port level on the SAN fabric?

- It is undesirable to have Windows NT and Sun Solaris systems accessing the same RAID-array port on the fabric; due to Windows NT's practice of attempting to write disk signatures to new disks it finds attached to the fabric. That creates the need for a network-fabric enforced way of segregating ports into logical groups of visibility.

How can we segregate different application types on the fabric?

- For example, it may be necessary to ensure that finance systems on the SAN cannot access the data owned by engineering systems, or web systems. That creates the need for a fabric-enforced way of grouping ports on the fabric into zones of visibility, which could be based on application, function, or departmental rules.

How can we isolate any single logical unit (LUN) on an array, permitting only a certain host(s) access to that LUN and no others?

- A basic advantage of a SAN is that a large number of hosts can share expensive storage resources. As it concerns RAID, this demands that multiple hosts have access to the disk storage LUNs through a single-shared port on the array. Security methods must be employed to ensure that LUNs behind a port are accessible only to the intended hosts. Without special software and architectures to manage multi-host block-level read/write access (*when multiple systems access the same LUN concurrently*), data corruption or data loss would occur.

How can we, from the host side, ensure that hosts see their storage ports and LUNs consistently as new storage is added, and after each reboot?

- In the world of fibre-channel SANs, the assignment of SCSI target IDs is moved from the storage side to the host/fiber channel (FC) host bus adapter (HBA) side. Thus, target IDs can be dynamically reassigned as new storage is added to an individual host via the SAN. Since this feature is a fundamental advance of SAN, the assignment of target IDs must be managed to ensure their consistency across storage devices, fabrics, and after host configuration changes.
-



Fabric management security concerns

Fabric-level security

When implementing a SAN, a fundamental level of fabric security is required; especially when adding switches to the fabric and configuration management of that fabric. This level of security more closely resembles the common security concerns in traditional TCP/IP networks. This same level of security is also very important in co-location environments, where multiple SAN fabrics are used and owned by multiple organizations that are in close proximity to each other.

Questions concerning fabric-level security

Concerning fabric-level security on a SAN, consider the following questions:

How is switch-to-switch security managed on the fabric; also, how can we enforce policies that prohibit non-authorized switches or hosts from attaching to the fabric?

- In early SAN environments, additional switches (*configured with a default password and login*) could be easily attached to an existing operating fabric, and that new non-secure switch could be used as a single point of configuration administration for the entire fabric. Technologies are needed that enforce access control at the fabric level, and ensure that only authorized and authenticated switches can be added to the fabric.

How can security and configuration be centrally managed on a fabric?

- In the initial phases of SANs evolution and even today, large fibre-channel fabrics are frequently made up of many 8- or 16-port fibre-channel switch building blocks. Each switch features both in-band and out-of-band management components [*simple network management protocol (SNMP), telnet, etc.*], and a switch-centric security control model. As large SANs evolve, technologies are needed to centrally control security, regarding the access and management of the fabric; also, to minimize the number of administrative access and security control points on the fabric.

How can we ensure that only authorized hosts are allowed to connect to the fabric, and to a specific port designated by the administrator?

- Initially, in SAN configurations, a host fiber channel HBA could attach to any point in a fabric and if the HBA was capable of basic fabric login, that HBA became a participating member of the fabric. Technologies are needed that allow a fabric-centric method of access control to govern which hosts can attach to a specific port or switch on the fabric. This would prevent a rogue attacker with a NT system and a FC HBA from attaching to a non-secure SAN for the purpose of configuration changes, or data access.

Continued on next page



Fabric management security concerns, continued

Questions concerning fabric-level security, (continued)

How can we ensure that the management tools used to manage the SAN, and SAN-management requests are coming from an authorized source?

- Multiple in-band and out-of-band methods are used to manage SAN fabric configurations. A tunnel of communication must exist between SAN management consoles and frameworks, and the target fabric being managed. That tunnel of communication must be secure and confirmed as authentic to prevent an attacker from using a management tool to access a non-secure SAN.

How can we ensure that configuration changes on the fabric are valid when there are multiple points of configuration management?

- In early SAN configurations, multiple administrators could log into different switches on the same fabric and perform fabric-configuration changes concurrently. After enabling and propagating those configuration changes fabric-wide, the fabric configuration could become corrupt due to conflicts. Fabric corruption usually occurs when configuration changes are made through multiple points on a fabric. Technologies are needed to ensure that fabric configuration changes are performed through a centralized and secure point in the SAN, and that those configuration changes do not cause configuration conflicts.
-



Data protection and fabric management

Security

Answers to the previously posed data protection and fabric management security questions, needs, and concerns are available today in various technologies. Typically, multiple technologies and methodologies are used to provide the highest level of security for SANs, and some of these are the following:

For data access and security:

- Fabric zoning
- LUN masking
- Persistent binding

For fabric management and protection:

- Fabric-to-fabric security
- Host-to-fabric security
- Management-to-fabric security
- Configuration integrity

The following discussion is about data access and security, and fabric management and protection technologies and methodologies that provide security and management for storage area networks.



Data integrity and security

Fabric zoning SAN implementations make data highly accessible; as a result, heightened network security and processes optimized for data transfers are needed. Fabric zoning establishes the way devices in the SAN interact, establishing a certain level of management and security.

What is zoning? Zoning is a fabric-centric enforced method of creating barriers on the fabric to prevent set groups of devices from interacting with each other. SAN architectures provide port-to-port connections among servers and storage devices through bridges, switches, and hubs. Zoning sets up efficient methods of managing, partitioning, and controlling pathways to and from storage devices on the SAN fabric; as a result, storage resources are maximized, and data integrity and data security are maintained. Additionally, zoning enables heterogeneous devices to be grouped by operating system, and further demarcation based on application, function, or department.

Types of zoning There are two types of zoning: Soft zoning and Hard zoning.

- Soft zoning uses software to enforce zoning. The zoning process uses the name server database located in the fibre-channel switch. The name server database stores port numbers and World Wide Names (WWN) used to identify devices during the zoning process. When a zone change is made, the devices in the database receive Registered State Change Notification (RSCN). Each device must correctly address the RSCN to change related communication paths. Any device that does not correctly address the RSCN, yet continues to transfer data to a specific device after a zoning change, that device will be blocked from communicating with its targeted device.
- Hard zoning uses only WWNs to specify each device for a specific zone. Hard zoning requires each device to pass through the switch's route table so that the switch can regulate data transfers by verified zone. For example, if two ports are not authorized to communicate with each other, the route table for those ports is disabled, and the communication between those ports is blocked.

Zoning configurations and components Zone configurations are based on either the physical port that the device plugs into or the WWN of the device. Zoning components include:

- Zones
- Zone sets
- Zone members

Continued on next page

Data integrity and security, continued

Zone construction

A zone is made up of servers and storage devices on the SAN fabric that can access each other through managed port-to-port connections. Devices in the same zone can recognize and communicate with each other, but not necessarily with devices in other zones unless a device, in that zone, is configured for multiple zones. Figure 1 shows a three-zone SAN fabric with both zones sharing the tape library in zone 2.




Figure 1: Three-Zone SAN Fabric

Continued on next page

Data integrity and security, continued

What is a zone member?

Zone members are devices within the same assigned zone. See Figure 2. Zone-member devices are restricted to intra-zone communications, meaning that these devices can only interact with members within their assigned zone. A zone member interacting with devices outside its assigned zone is not permitted, unless that device is configured for multiple zones.

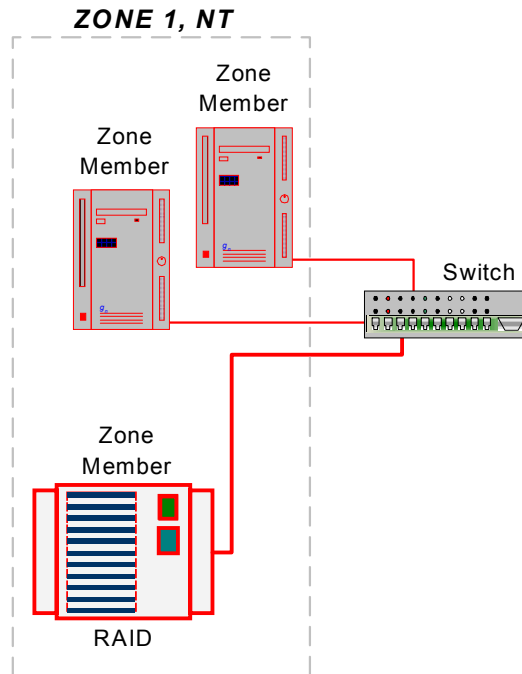


Figure 2: Zone Members

Zone member identification

Zone members are recognized by port number or World Wide Name (WWN). A WWN is a 64-bit number that uniquely identifies zone members.

Continued on next page

Data integrity and security, continued

What is a zone set?

A zone set is a group of zones that function together on the fabric. Each zone set can accommodate up to 256 zones. All devices in a zone see only devices assigned to their zone, but any device in that zone can be a member of other zones. In Figure 3, all 4 zones see member A.




Figure 3: Zone Set

Zoning limitation

Today, fabric zoning cannot mask individual tape or disk logical unit numbers (LUNs) that sit behind a device port.

Continued on next page

Data integrity and security, continued

LUN masking

LUN masking is a RAID system-centric enforced method of masking multiple LUNs behind a single port. LUN masking is configured at the RAID-array level, using World Wide Port Names (WWPNs) of server HBAs. See Figure 4. LUN masking allows disk storage resource sharing across multiple independent servers. With LUN masking, a single large RAID device can be sub-divided to serve a number of different hosts that are attached to the RAID through the SAN fabric. Each LUN (*disk slice, portion, unit*) inside the RAID device can be limited so that only one or a limited number of servers can see that LUN.

LUN masking can be done either at the server HBA or at the RAID device (*behind the RAID port*). It is more secure to mask LUNs at the RAID device, but not all RAID devices have LUN masking capability; therefore, some HBA vendors allow persistent binding at the driver-level to mask LUNs.

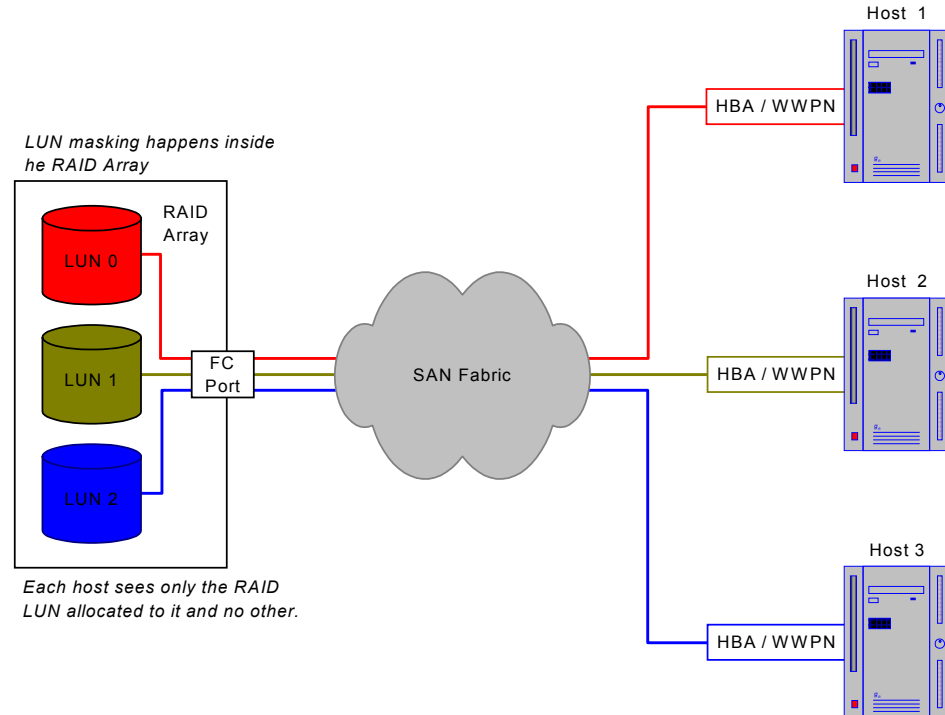


Figure 4: LUN Masking Example

Continued on next page

Data integrity and security, continued

Persistent binding

Persistent binding is a host-centric enforced way of directing an operating system to assign certain SCSI target IDs and LUNs. For example, where a specific host will always assign SCSI ID 3 to the first router it finds, and LUNs 0, 1, and 2 to the three-tape drives attached to the router. See Figure 5.

Operating systems and upper-level applications (*such as backup software*) typically require a static or predictable SCSI target ID for their storage reliability and persistent binding affords that happening.

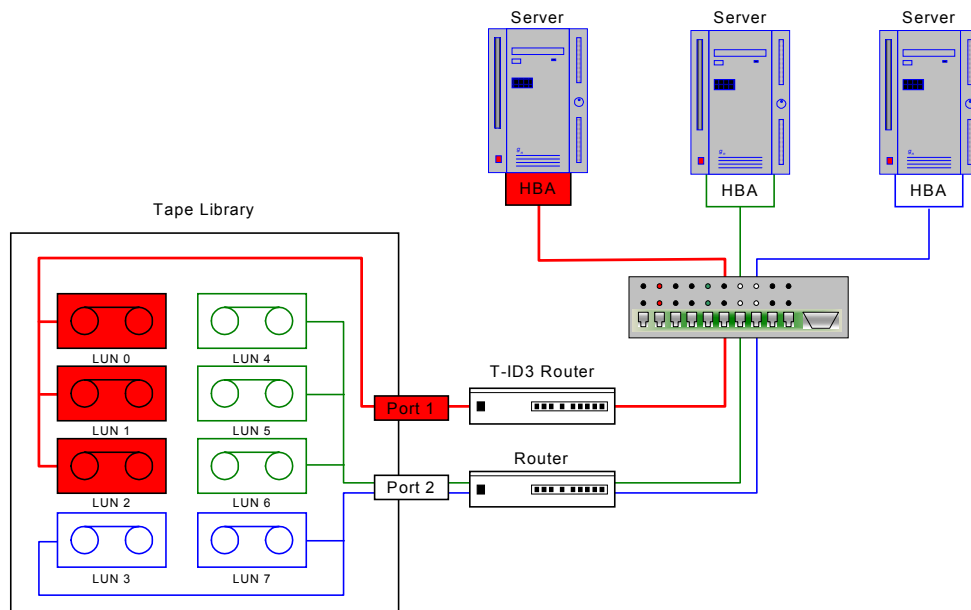


Figure 5: Persistent Binding



Fabric management and protection

Fabric-to-fabric security technologies

Fabric-to-fabric security technologies permit Access Control Lists (ACLs) to allow or deny the addition of a new switch to the fabric. Access lists filter network traffic by controlling whether routed packets are forwarded or blocked at the router interface. For example, access lists can *allow* one host the right to access a certain part of the network and *deny* another host that same access. Access control lists provide a basic level of security for accessing the network.

Public Key Infrastructures (PKI) technology can be applied as a mechanism for authenticating the identity of a new switch.

Additionally, fabric-wide security databases help to ensure that all new authorized switches added to the fabric inherit fabric-wide security policies, so that a new, out-of-the-box switch does not become a non-secured access point.

Host-to-fabric security technologies

Host-to-fabric security technologies can apply ACLs at the port-level of the fabric to allow or deny a particular host's FC HBA from attaching to a specific port. This would prevent an unauthorized intruder host from attaching to the fabric via any port. The host's ability to log into the fabric is explicitly defined and is allowed with this model.

Management-to-fabric technologies

Management-to-fabric technologies can use PKI and other encryption (*such as md5*) technologies to ensure that a trusted and secure management console-to-fabric communications layer exists. PKI and other encryption help ensure that the management console or framework used to control the fabric is authentic and authorized. In addition, encryption methodologies can restrict the number of switches on the fabric from which management and configuration changes are propagated to the rest of the fabric. That will create a SAN with a minimal number of security control points.

Configuration integrity technologies

Configuration integrity refers to technologies that ensure propagated fabric configuration changes only come from one location at a time, and are correctly propagated to all switches in the fabric with integrity. A distributed lock manager is one way of ensuring that a serial and valid configuration change is enabled on the fabric.



Conclusion

SAN implementations make data highly available by delivering shared storage in open, non-proprietary environments via any-to-any connectivity. The advantage of any-to-any connectivity can be a liability unless well thought out security policies are put into place to manage how devices interact within the SAN. Shared storage in a SAN environment requires safeguards to ensure data integrity, and to prevent unwanted access from unauthorized systems and users. This discussion briefly explored some of the technologies and their associated methodologies used to ensure data integrity, and to protect and manage the fabric. Each technology has advantages and disadvantages; and each must be considered based on a well thought out SAN security strategy, developed during the SAN design phase.

Moreover, to expect that the required level of security can be achieved from any one of the previously discussed technologies, independent of all others, is unwise. The astute information storage architect clearly understands that in a heterogeneous SAN environment, with diverse operating systems and vendor storage devices, that some combination or all of the aforementioned technologies could be required to ensure that the SAN is secure from unauthorized systems and users.

Finally, the SAN security strategy must be periodically addressed as the SAN infrastructure evolves and as new technologies emerge; this will ensure that the proper level of security is maintained and the SAN fabric is properly managed.

Datalink Corporation
8170 Upland Circle
Chanhassen, MN, 55317
www.datalink.com