



LUN Security Considerations for Storage Area Networks

by Hu Yoshida

Contents

LUNs and the Data Transfer Process: a Short History	1
SCSI-based Data Transfer and the LUN Security Problem	2
UNIX and Windows NT LUN Security Issues	2
LUN Security with SANs	3
Host Software	3
Host Bus Adapter Utilities	4
Switch Zoning	5
Mapping Within a Storage Controller	6
LUN Security is Essential	7

LUN Security Considerations for Storage Area Networks

by Hu Yoshida

Storage area network (SAN) technology offers IT departments the ability to connect multiple server hosts on a network and provide shared access to any storage resource connected to that network. While SANs reduce costs by pooling storage resources, they also expose the data on shared storage to access by unauthorized users and overwrites by multiple hosts. This data, which is addressed by logical unit (LUN), must be protected from unauthorized or indiscriminate host access.

Currently, there are four approaches to providing LUN security: host software, host bus adapter utilities, switch zoning, or mapping within a storage controller. This paper describes each approach, identifies where each might be used, and details the process of securing LUNs from unauthorized host access.

LUNs and the Data Transfer Process: a Short History

Small Computer Systems Interface (SCSI) is the primary protocol for transferring data between host processor and storage devices in UNIX[®] and Windows NT[®] operating systems. SCSI is based on a master/slave paradigm. The connection between an open systems (UNIX or Windows NT) server and storage is accomplished through a host bus adapter (HBA). The HBA plugs into a slot on the internal bus of the server and connects a SCSI cable to the storage device or subsystem. Known as the initiator, the HBA begins the transfer of data to and from a target device. Although the SCSI protocol supports a combination of initiators and targets on a common bus, configurations have been limited to one initiator since most operating systems cannot manage shared devices.

Initially, HBAs were called controllers and talked directly to disks, which originally served as targets. When controller-based storage subsystems came onto the scene, another level of identification was required to address the controller port on the storage subsystem as well as the disks behind the controller. With the new labeling, the initiator would see the storage port that was connected on the SCSI bus as the target, and the target would identify the disks behind it as LUNs. Today, storage units are addressed based on a slightly different controller, target, and LUN combination. “Controller” now refers to the identification of the host bus adapter on the system bus, not to an external storage control unit.

SCSI-2 specifications support eight LUNs per target. The newer SCSI-3 specifications define an encoded 64-bit identifier, which allows for many more LUNs. However, the vendor defines the maximum number of LUNs available in any device. In the Hitachi Freedom Storage[™] 5800 subsystem, a total of 64 LUNs can be defined across the total number of storage ports. In the Freedom Storage 7700E,

each port can support up to 120 LUNs. Nevertheless, the number of LUNs that the host sees depends on the capabilities of the host bus adapter driver. If the driver is written to recognize only eight LUNs per target, it will report only eight LUNs, even if the storage array is capable of reporting 120 LUNs.

Some operating systems use a volume manager to provide another level of abstraction between the SCSI driver and the host file system. Volume managers build virtual volumes or devices on top of storage LUNs. These virtual volumes are accessed by the file system or database in the same way that LUNs would be accessed. Volume managers can gather a number of LUNs into a volume group, which can be carved up and presented to the file system as a number of virtual volumes (sometimes called logical devices). Volumes are composed of other virtual objects, which can be manipulated to change the volume's configuration.

Volume managers originally were used to provide software RAID protection for non-intelligent disks. Today, with intelligent storage arrays like the 7700E, software RAID protection is unnecessary. However, volume managers can enhance operations by manipulating volume manager objects to add capacity, optimize performance, and enable nondisruptive backup and other administrative tasks. They can also draw from two or more volumes on multiple physical devices and present the data to the file system as one volume. This is called mirroring. With fibre channel, the mirror could reside on a physical volume located up to 6.2mi./10km away.

Some volume managers have other functions, such as alternate path management. Windows NT will not have a volume manager until the release of Windows 2000, when Microsoft® will OEM a volume manager from VERITAS®. Windows NT 4 provides a fault-tolerant file system for disk mirroring. Currently, the Disk Administrator fills some of the functions of the volume manager. For a complete discussion on the features and merits of volume managers, please refer to www.veritas.com.

SCSI-based Data Transfer and the LUN Security Problem

When SCSI bus systems are initialized, the SCSI driver in the host bus adapter will “walk” the bus to discover what targets the bus holds. It asks each target to report the LUNs to which it connects. Each initiator that shares the same bus will see the same targets and LUNs. The drawback, since there is no reserve/release capability at the target/LUN level, is that multiple initiators can overwrite each other's data. (Although there is a definition for SCSI reserve/release in the SCSI-3 specifications, this has not been implemented by vendors.) This creates what is known as the LUN security problem.

UNIX and Windows NT LUN Security Issues

In UNIX, the LUNs on a SCSI bus can be partitioned between systems. UNIX systems will see all the LUNs on the bus but will enable the operator to mount, or select, only the LUNs that the operator has permission to use. As long as each operator on each UNIX system plays by the same rules, LUN security can be preserved. However, this will not work if Windows NT systems are involved. A Windows NT system assumes that it owns all the LUNs it discovers and writes a signature on each LUN during the discovery process. This signature is used to ensure that each LUN it sees is unique and not duplicated.

LUN Security with SANs

The problem of LUN security becomes much larger in a SAN environment. A SAN allows many more SCSI host bus adapters to connect to storage subsystems over a network (rather than a bus) and to access many more LUNs. A SAN Fibre Channel Arbitrated Loop (FC-AL) can address up to 126 nodes, while a fibre fabric can access up to 16 million nodes. A node can be a SCSI initiator or a target/LUN with any-to-any connectivity to any node in the loop or fabric. To realize the full capability of a storage area network, a solution for LUN security is imperative. The following sections explore the four approaches to providing LUN security: host software, host bus adapter utilities, switch zoning, and mapping within a storage controller.

Host Software

A number of software vendors, such as Data Direct, Mercury Computer Systems™, CrosStor, Retrieve, and others, provide middleware that can help provide LUN security. This software intercepts I/O requests and routes them over a LAN to a file server to gain access to a pool of storage that is managed by the file server. The file server manages files and grants access to the requesting system.

Initially, most of this middleware was developed and sold in the video and prep-press markets, which were early adopters of fibre. Transoft® Networks, Inc., which was recently acquired by Hewlett-Packard®, currently offers a configuration database and a GUI interface that allows the user to drag and drop LUNs between Windows NT systems without a reboot. Transoft's approach does not require a file server, but it does require a Qlogic host bus adapter (HBA) and a modified HBA driver. Some products, such as SANergy™, also provide heterogeneous file sharing by writing raw I/O to a common NTFS file system owned by the SANergy file server. (See Figure 1, Software Approach to LUN Security.)

The drawbacks to these systems is the need for another software layer, which may be sensitive to release levels and single points of failure in the file server. Another consideration is the type of file system that is used. Data Direct uses a proprietary file system, while others use more standard file systems like NFS, NTFS, or CIFS.

The advantage of host software is that it can provide centrally managed security and lock management beyond the LUN to the block level. This level of security will be required for future data sharing. Implementation of this level in the storage hardware is not feasible until file level—not just block level—information becomes available outside of the host.




Figure 1: Software Approach to LUN Security. The file server controls the file system metadata, as well as allocations, authorizations, authentications, and locks. Metadata is communicated across the LAN with standard NFS. The client agent intercepts NFS read/write and initiates block I/O to the standard NFS file system.

Host Bus Adapter Utilities

Emulex Corporation and JNI™ offer LUN security measures via drivers that allow the user to mask off the LUNs that a particular host system should not see. These drivers initialize with all the fibre channel LUNs that they discover on the network, and the LUNs are reported for each of the fibre channel nodes in the loop or fabric. The nodes are identified by a unique WorldWide Name (WWN) that is stamped on the fibre channel node chip set by the manufacturer. (A WWN is assigned by the IEEE and is similar to a MAC address in IP or a URL on the Internet.) A node is a fibre port on a storage subsystem or device. The drivers provide a masking utility that can be initiated off the desktop in Windows NT or via the command line with UNIX. The masking utility allows the operator to edit the list of WWNs and LUNs down to the set that is authorized for that host. (See Figure 2, Emulex Persistent Bind Feature.) The host is then rebooted and will only see the specified, or unmasked, LUNs. Emulex calls this feature “Persistent Bind” since it is based on the unique WWN of the storage node, which always remains the same even if its physical addressing is changed. If the storage node were to be moved to another location in the SAN, the HBA would still see the same LUNs.

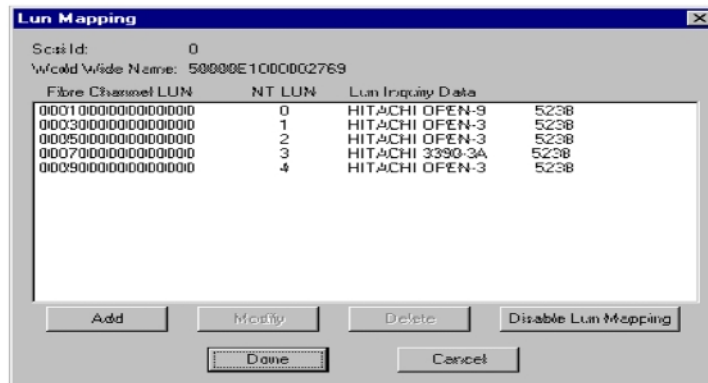


Figure 2: Emulex Persistent Bind Feature. This host-based utility is used to add or delete LUNs that are reported for each fibre channel storage node that is identified by a WWN. Physical addressing may change, but the binding is persistent (or unchanging) since it is based on the unique WWN of the storage port.

This security solution works well in small installations. In large installations, a great deal of coordination is required to ensure that the LUNs are partitioned properly across all host consoles. Consider the challenges of mapping 20 or more hosts with multiple HBAs across 100 or more LUNs, when each host must be configured individually. Add to that the disruption that occurs when a storage port adapter with a unique WWN must be swapped out and replaced with another unique WWN for the same set of LUNs. There is also the danger of a “rogue” server that does not coordinate its security with the other hosts.

The advantage for security at the HBA level is that this solution provides LUN masking for all the fibre devices on the SAN. It is not specific to the fibre infrastructure of hubs, switches, and routers, to a vendor’s storage subsystem, or to middleware on the hosts.

Switch Zoning

If a fibre channel network is based on a switch like ANCOR[®], Brocade[®], or Vixel, switch zoning can be used to provide masking down to the port level for all nodes that are known to the switch. All LUNs that are attached to a port can be masked from hosts that do not access that port via switch zoning. Switch zoning cannot mask individual LUNs that sit behind a port. All the hosts connected to the same port will see all the LUNs addressed through that port. (See Figure 3, Switch Zoning Example.) Fabric switches require that any node that attaches to the switch must log in to the switch and register its WWN in the Simple Name Server (SNS) function of the switch. An internal 24-bit address is assigned to each WWN in the SNS. This allows the host drivers to discover targets through a lookup in the SNS instead of “walking” the entire network of a possible 16 million nodes. The SNS can be zoned by WWN, which has the advantage of flexibility, or by hardware port. A configuration based on WWN can be dynamically changed to serve changing business requirements. For instance, a tape library could be moved to different zones for different periods for backup. A node can be moved to a different port address without a change to the zone, since it is based on WWN.

The disadvantage of WWN zoning is that it can be spoofed if someone knows the WWN. Zoning by hardware port is less flexible, but has the advantage of being secure. In the Hitachi Freedom Storage 7700E, a port can be zoned by the switch utility program, and the LUN manager utility in the 7700E Remote Console can be used to assign LUNs to that port. Future implementations of WWN zoning will check the source WWN to eliminate spoofing by unauthorized WWNs.

Switch zoning is easily implemented through out-of-band management user interfaces like Ethernet. It can be centrally managed through a LAN or Web interface. The drawback is that it only supports initiators and targets that attach to a switch. It does not provide security beyond the port level of a storage subsystem and it cannot mask LUNs from initiators that access the same storage port.




Figure 3: Switch Zoning Example.

Mapping Within a Storage Controller

Some storage subsystems, such as the Freedom Storage 7700E, have the ability to do LUN masking within their storage controllers. Through a remote console attached to multiple 7700E storage subsystems on a private LAN, the WWN of all fibre-channel-attached host bus adapters can be mapped against the LUNs contained in the 7700E. This allows multiple host bus adapters to access different LUNs through the same storage port, independent of any intervening SAN infrastructure, such as hubs or switches (see Figure 4, Hitachi Freedom Storage 7700E LUN Security).

The LUN security utility on the remote console can be used to add or delete LUNs from the mask for each WWN of a host bus adapter. After this mapping is set, the initiators will only see the LUNs that they have permission to view during IOSCAN/boot time. The remote console can be attached to another LAN, which would enable the LUN masking utility to be launched from a central management console using tools like Remotely Possible® or Symantec's pcAnywhere.

The advantage of LUN masking in the storage controller is that it allows many more hosts like Windows NT to attach to a given 7700E through a common fibre channel port and still maintain LUN security. It can work in point-to-point mode or through hubs as well as switches, in loop or fabric mode. Since it is based on the WWN of the host bus adapter, it is independent of the physical loop or switch address. This LUN masking is implemented or checked during IOSCAN/boot time—not with every I/O—in order to sustain the high performance of the 7700E. However, it is possible for a user with the proper systems authorization to change configurations after IOSCAN time and bypass the mask. It has to be remapped if a host bus adapter fails and needs to be replaced. Ease of access to the WWN of the host bus adapter is dependent on the operating system and driver.

Hitachi Data Systems intends to integrate the 7700E LUN Security Utility with SAN management programs like Computer Associates' SANITI, which can discover and map all the components of a SAN. SANITI would supply the WWN of the HBAs and provide notification when the LUN configuration has changed after IOSCAN/boot time.

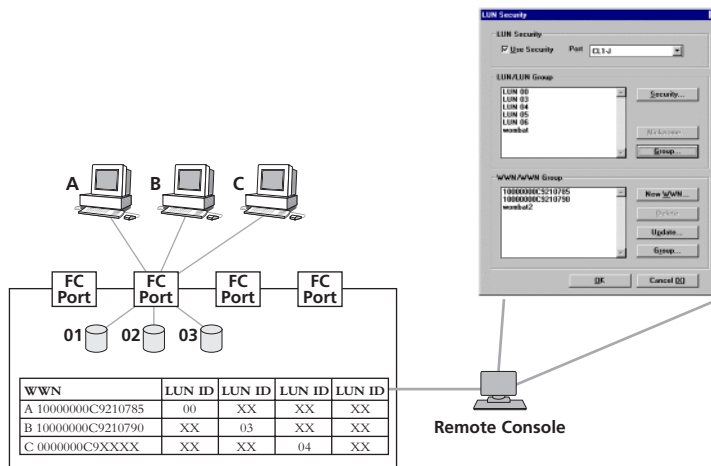


Figure 4: Hitachi Freedom Storage 7700E LUN Security. Users enter HBA WorldWide Names through the 7700E's Remote Console. The 7700E assigns LUNs to WWNs and allows multiple-host access through the same 7700E fibre port.

LUN Security is Essential

The advantages of any-to-any connectivity afforded by SAN can be a liability unless LUN security is implemented. We have reviewed four ways in which vendors provide a measure of LUN security today. In a mixed environment with many different vendor storage devices, there may well be a need to implement some or all of these methods.

Although this topic is generally referred to as LUN security, it should be thought of as LUN masking since that is really the extent of the protection that is provided. There is no security in the sense of authorization and authentication. These LUN masking techniques work adequately in a well-behaved environment, but they can be easily bypassed. The user should be aware of the shortcomings of each implementation and take additional measures to ensure security. Such measures might include an authorization process for access to these features and implementation of multiple masking techniques to ensure a cross-check between the HBAs and the storage subsystem, for example. The coordination and management of these different methods is not integrated at this time. Although the vendors provide configuration utilities that can be launched from a common desktop, the user will need to map LUN and server configurations on a spreadsheet in order to coordinate mapping across these different implementations. Hitachi Data Systems recognizes the need for an integrated management approach and is working with various standards bodies and vendors to that end.

Hitachi Data Systems

www.hds.com

Corporate Headquarters

750 Central Expressway
Santa Clara, California 95050-2627
U.S.A.
(408) 970-1000
info@hds.com

Asia Headquarters

39-09 Tower One
Lippo Centre
89 Queensway
Hong Kong
2525-2385
infoasia@hds.com

Australia/New Zealand Headquarters

11-17 Khartoum Road
North Ryde NSW 2113
Australia
02-9325-3300
info@hds.com.au

Canada Headquarters

380 Saint-Antoine Street West
Suite 7000
Montreal, Quebec H2Y 3X7
Canada
(514) 982-0707
info@hdscanada.com

Europe Headquarters

Sefton Park
Stoke Poges
Buckinghamshire SL2 4HD
United Kingdom
01753-61-8000
info@hds.co.uk

Latin America Headquarters

750 Central Expressway, MS 3468
Santa Clara, California 95050-2627
U.S.A.
(408) 970-7447
infolatin@hds.com

U.S. Headquarters

750 Central Expressway
Santa Clara, California 95050-2627
U.S.A.
(408) 970-1066
ussalesinfo@hds.com

Hitachi Data Systems is registered with the U.S. Patent and Trademark Office as a trademark and service mark of Hitachi, Ltd. The Hitachi Data Systems logotype is a trademark and service mark of Hitachi, Ltd.

Freedom Storage is a trademark of Hitachi Data Systems Corporation.

UNIX is a registered trademark, licensed exclusively through X/Open Company Limited.

Windows NT is a trademark of Microsoft Corporation.

VERITAS is a trademark of VERITAS Software Corporation.

Mercury Computer and SANergy are trademarks of Mercury Computer Systems, Inc.

Hewlett-Packard is a trademark of Hewlett-Packard Company.

JNI is a trademark of Jaycor Networks, Inc.

Ancor is a trademark and property of Ancor Communications, Inc.

Brocade is a trademark of Brocade Communications Systems, Inc.

Remotely Possible is a trademark of Computer Associates, Inc.

Transoft is a trademark of Transoft Networks, Inc., a Hewlett-Packard Company.

Macintosh is a trademark of Apple Computer, Inc.

Notice: This document is for informational purposes only, and does not set forth any warranty, express or implied, concerning any equipment or service offered or to be offered by Hitachi Data Systems. This document describes some capabilities that may be configuration-dependent, and features that may not be currently available. Contact your local Hitachi Data Systems sales office for information on feature and product availability.

©1999, Hitachi Data Systems Corporation.
All Rights Reserved/LD2.5M099/PERI-091