

iSCSI Security

ELEN 689 – Network Security

John Price

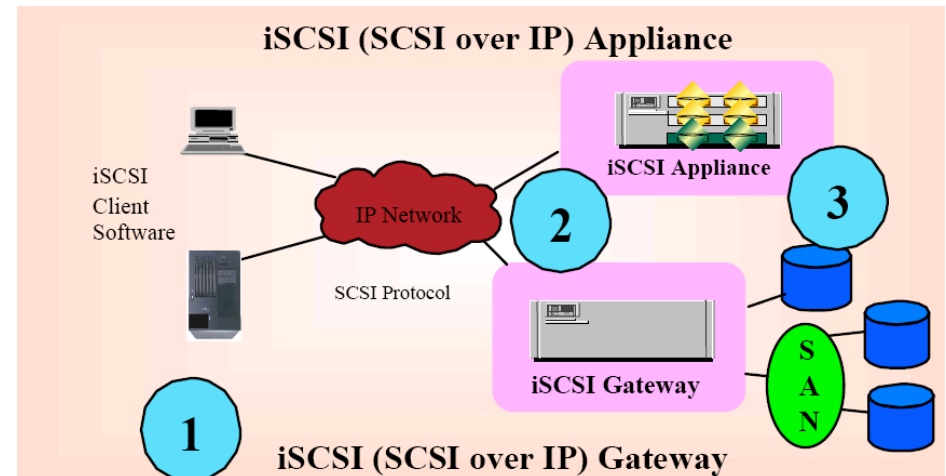
Peter Rega

Outline

- iSCSI Basics
- iSCSI and NAS Differences
- iSCSI Security
- Current Vulnerabilities
- References

iSCSI Basics

- Internet **S**mall **C**omputer **S**torage **I**nterconnect
- The iSCSI protocol is a mapping of the SCSI remote procedure invocation model over the TCP protocol
- Two approaches to iSCSI implementations: iSCSI Appliance and iSCSI Gateway
- Enables the transfer of Block I/O data over an IP network
- Operates on top of TCP through encapsulation of SCSI commands in the TCP/IP datastream



NAS and iSCSI Technology Overview, Storage Networking Industry Association, 2006

iSCSI Packets (PDU)



IP Storage Protocols: iSCSI, Storage Networking Industry Association, 2006

iSCSI Protocol Data Unit (PDU): Provides ordering and control information. Contains iSCSI control info, with optional SCSI Commands &/or Data

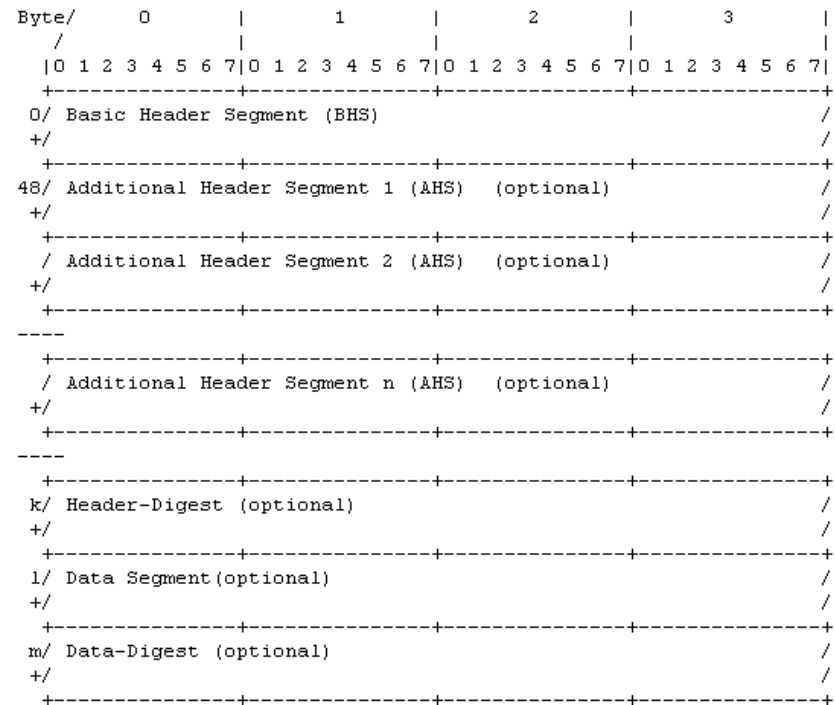
Provides Reliable data transport and delivery (TCP Windows, ACKs, ordering, etc.) Also demux within node (port numbers)

Provides IP "routing" capability so that packet can find its way through the network

Provides physical network capability (Cat 5, MAC, etc.)

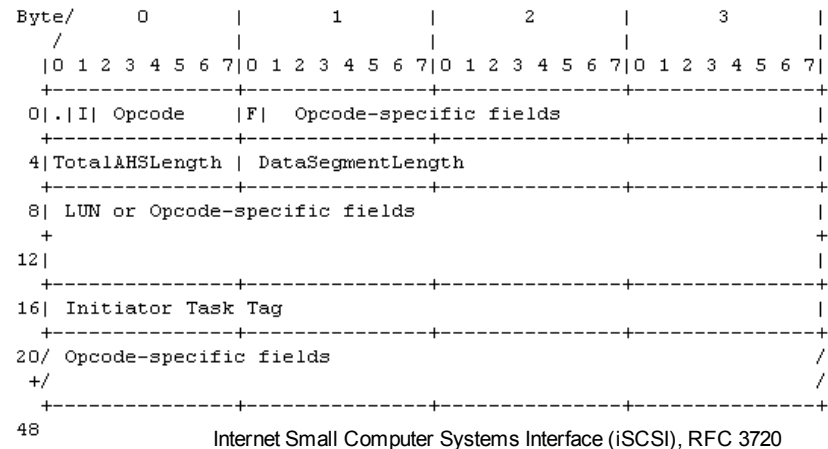
PDU Structure

- BHS is the only required segment in the PDU
- All segments and digests are padded to the closest integer number of four byte words
- All multi-byte integers are big-endian

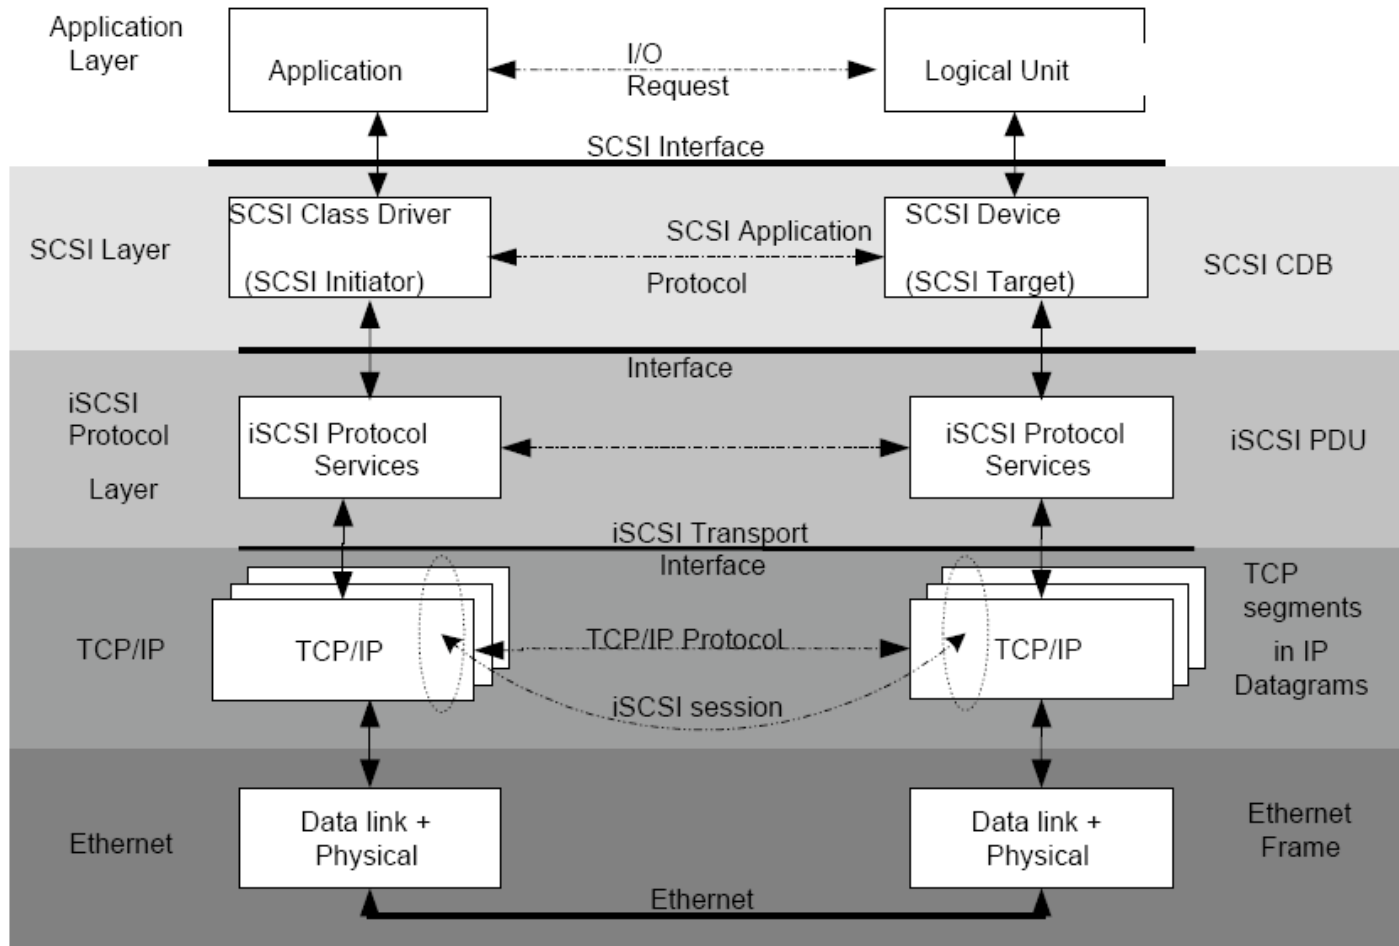


Basic Header Segment

- BHS is always 48 bytes long
- The opcode and the DataSegmentLength are ALWAYS in the BHS
- 2 categories of opcodes: initiator and target
- When used, the Initiator Task Tag and LUN must always be in the same location




iSCSI Layer Model



iSCSI Sessions

- Sessions can be thought of as a scsi port attached to one or more TCP/IP connection
- Login phase always begins a session
- Sessions guarantee SCSI commands delivered in order
- Can recover from lost connections



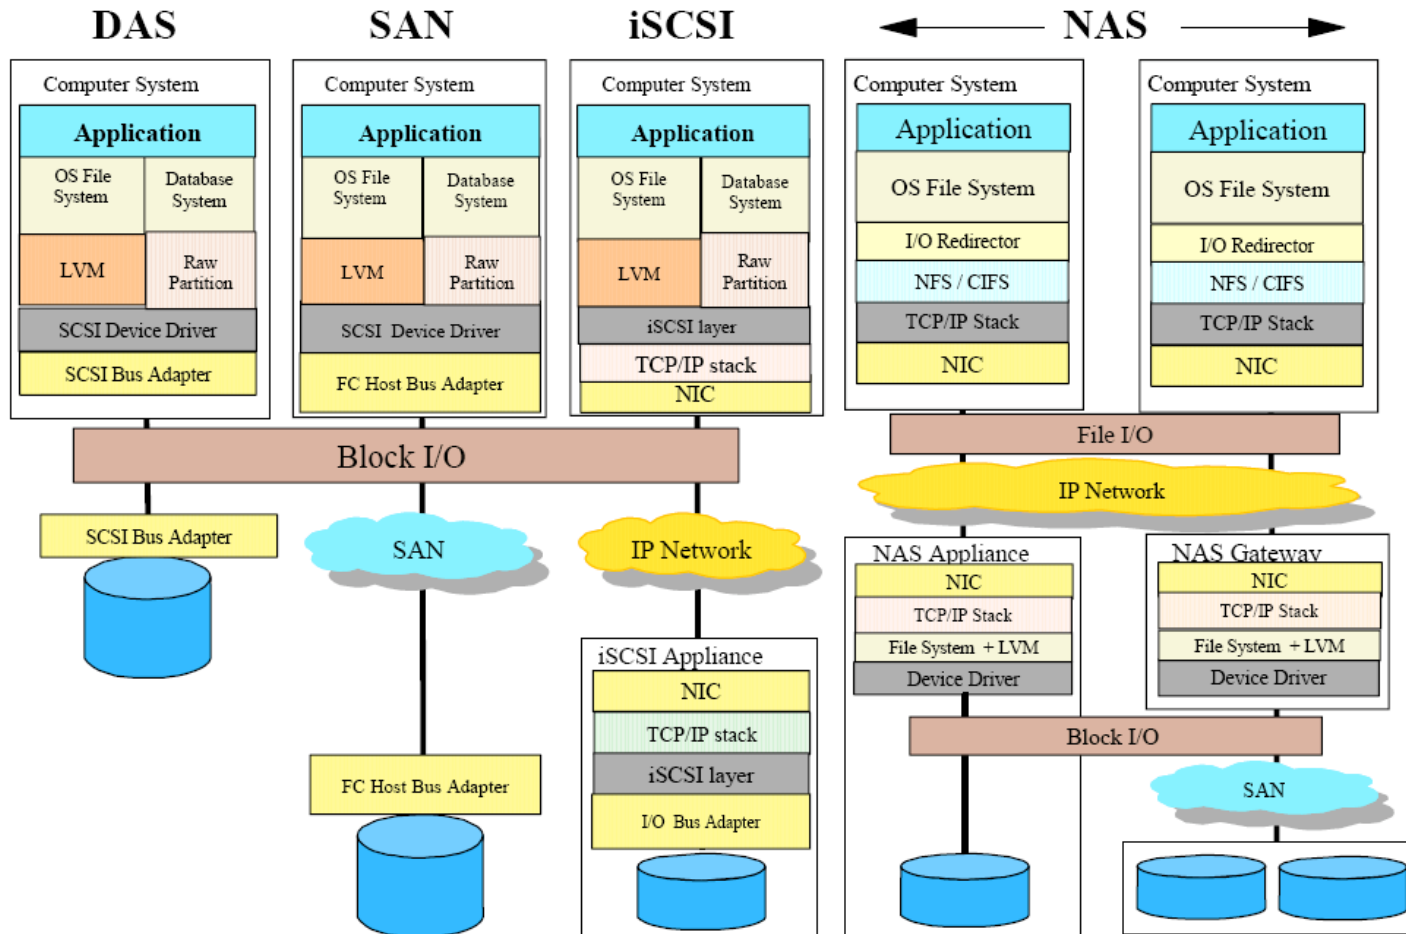
iSCSI Integrity

- iSCSI adds a cyclic redundancy check (CRC-32C) along with TCP/IP checksum and Ethernet level CRCs
- CRC “check word” is called a “digest”
- Header digest and data digest are optional to use in the PDU (MUST be used to fully comply with spec)
- The Digests insure correct operation, data placement, and data is unmodified

TCP Overhead

- Every TCP connection has processing overhead potential:
 - Connection setup/teardown
 - Window management
 - Congestion control
 - TCP Segmentation
 - IP Fragmentation
 - Checksum calculations
- TOEs are helpful but not required for iSCSI
- Therefore, iSCSI is currently used for entry level to mid size servers
- High end implementations would require Gigabit Ethernet and TOE


iSCSI vs. NAS



NAS and iSCSI Technology Overview, Storage Networking Industry Association, 2006

iSCSI vs. NAS

Percent CPU Overhead



With unmodified TCP/IP, iSCSI is 1/3 the overhead of NFS
With all TCP/IP copy overhead offloaded (0 Copy) iSCSI was 1/12 the overhead of NFS

- NFS
- SCSI over GE-TCP/IP
- % of NFS cpu used by SCSI over GE-TCP/IP
- SCSI over GE-TCP/IP 1 copy
- SCSI over GE-TCP/IP 0 copy
- NFS (with 0 TCP/IP data copies) est.
- % of NFS cpu (with 0 TCP/IP data copies) used by SCSI over GE-TCP/IP (0 copies)

Goal: Use NAS for Sharing Files, and iSCSI for everything else

iSCSI Security Considerations

- Must prevent against active (man in the middle, masquerading) and passive (snooping) attacks.
- iSCSI protocol allows for configuration without any security measures
 - Use only in extreme cases/closed environments

Phases of iSCSI security

- In-band Initiator-Target Authentication
 - Initial series of connection login packets exchanged between the target and initiator.
 - Security here exists at the iSCSI connection level.
 - Targets and Initiators are not required by spec to authenticate each other.
 - Unable to assure secure IP level communications, so care must be taken.

In-band Initiator-Target Authentication

- In-band authentication only assures users of correct endpoints.
 - If an underlying security mechanism is not used (Ipsec) then all packets are sent in plain text.
 - Only assures initiators that they have connected to a trusted target, and vice-versa.

Methods of In-band Authentication

Name	Description
KRB5	Kerberos V5 - defined in [RFC1510]
SPKM1	Simple Public-Key GSS-API Mechanism defined in [RFC2025]
SPKM2	Simple Public-Key GSS-API Mechanism defined in [RFC2025]
SRP	Secure Remote Password defined in [RFC2945]
CHAP	Challenge Handshake Authentication Protocol defined in [RFC1994]
None	No authentication

IP Layer Security

- iSCSI uses IPsec for packet level (IP level) protection. IPsec provides:
 - Cryptographic integrity
 - Authentication
 - Confidentiality
- If multiple devices are used, then the combined device is considered the iSCSI device

iSCSI and IPsec

- Data authentication and integrity is provided by a cryptographic Message Authentication Code (MAC)
- Types of MAC's
 - Unconditionally secure (one time pad)
 - Hash-function based
 - Stream cipher-based
 - Block cipher-based

iSCSI and IPsec

- Any iSCSI compliant device must impliment IPsec with ESP in tunnel mode
 - May also provide IPsec in transport mode
- With the speeds that iSCSI can operate at, a single 32 bit number (used as the SA in IPsec) could cycle very fast.

Network Appliance Vulnerability

- Unauthenticated iSCSI Initiators can bypass iSCSI authentication on NetApp Filers by manipulating the iSCSI Login Negotiation protocol.
- Login negotiation is in 3 phases:
 - Security Mode: client and server verify identity
 - Operational Mode: client and server negotiate session
 - Full Feature Mode: client and server exchange SCSI commands
- Clients can force server to go to operational mode without proving identity
- Data stored on iSCSI-mapped LUNs can be read and altered by the attacker
- Has been fixed by Network Appliances

Linux-iSCSI Vulnerability

- Linux iSCSI implementation installs the iscsi.conf file with world-readable permissions on some OS's, which allow local users to gain access to the cleartext CHAP password
- Only exploitable locally
- Would allow a local attacker to corrupt the password, obtain user passwords, or manipulate data

References

- Hibbard, Eric, Hitachi Data Systems, *Security 1: Storage Security - Beyond the Introduction*, 2006 Storage Network Industry Association
- Huffard, John L., and Zamar, Ahmed, *IP Storage Protocols: iSCSI*, 2006 Storage Network Industry Association
- Singer, Wolfgang, IBM Austria, *NAS and iSCSI Technology Overview*, 2006 Storage Network Industry Association
- Internet Small Computer Systems Interface (iSCSI), RFC 3720, <http://www.ietf.org/rfc/rfc3720.txt>
- SecurityFocus, Network Appliance iSCSI Authentication Bypass, <http://www.securityfocus.com/archive/1/414558>
- Linux-iSCSI Vulnerability, <http://nvd.nist.gov/nvd.cfm?cvename=CVE-2002-0849>
- <http://www.rsasecurity.com/rsalabs/node.asp?id=2177>