TO: X3T10 SCSI Working Group
RE: On Protected Persistent Reservation
FR: Lansing Sloan
DT: November 7-8, 1995

Mail Stop L-60
Lawrence Livermore National Laboratory
7000 East Avenue
Livermore, CA 94550-9900

ljsloan@llnl.gov
Phone: 1-510-422-4356
FAX: 1-510-423-8715

Scalable I/O Facility Project
http://www.llnl.gov/liv_comp/siof.html

Outline

- Assumptions and Situation
- Proposal and analysis
- Perspective

Assumptions and Situation

NAP Goal

Competitive mass market secure Network Attached Peripherals (NAPs) that file servers can use.

Background of proposal

This proposal is based on discussions by Bob Snively (Sun Microsystems) and Lansing Sloan on July 13, 1995, after discussions at the ANSI X3T10 SCSI Working Group on July 12. Key ideas were the use of passwords to identify authorized initiators and rules to limit what other initiators can do. Bob made a presentation to the SCSI WG in September 1995.

SIOF architectural assumptions

- Data flows directly between clients and peripherals.
- Authorized storage systems should control peripherals. Clients and others should not.

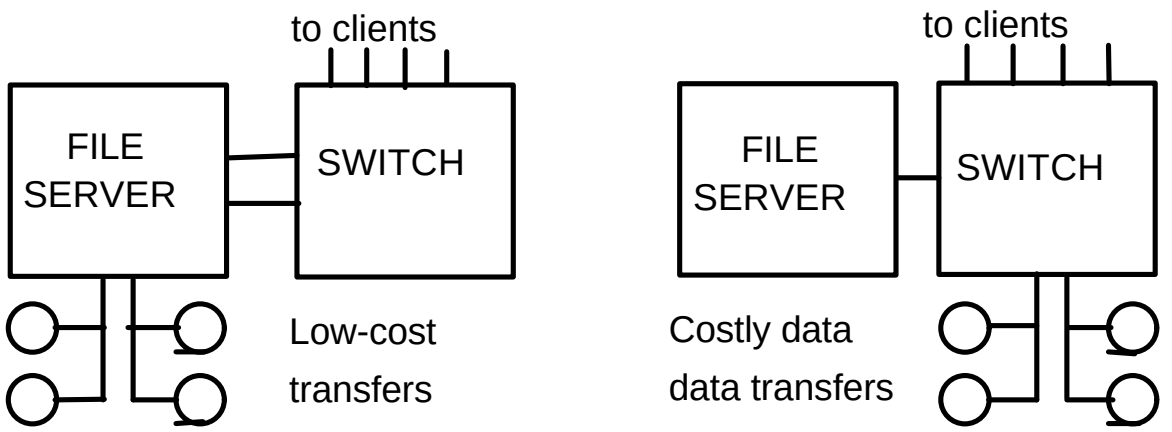

Technology assumptions

- A mostly-switched interconnect is used for scalability and high performance.
- Address reassignments can be detected and systems can be restarted after such reassignment.
    - Note: Undetected address reassignment could result in spoofing and in sending sensitive information (e.g., passwords) astray (e.g., to attackers).
    - LANs, busses, loops are OK in special contexts.

Security-related assumptions

- Physical security for storage: Storage system modules, storage system peripherals, the interconnections among these, and interconnect management system(s) are sanctioned, trusted, physically protected, and managed appropriately.
    - Cannot snoop and forge within a storage complex.
    - If clients and storage components share parts of a path, storage interconnections remain protected.
- The interconnect ensures correct source addresses adequately.
- Encryption is not used or assumed here.
- Security and authentication among clients and components of storage systems is outside the scope.

SCSI traditionally has no security mechanisms.

- Any "Initiator" can send commands to any reachable "Target." There are no provisions for peripherals to obey only authorized storage systems.

SIOF Needs for SCSI security in peripherals

- Peripherals able to recognize "their" controlling storage systems.
- Storage systems able to find and recognize "their" peripherals.
- Peripherals obey only their controlling storage systems.
- When specified by their controlling storage systems, peripherals exchange specified data with whatever is at specified addresses.

Proposal and remarks

Extended proposal for secure NAPs

Based on the Protected Persistent Reservation proposal.

- An SCSI peripheral accepts non-LUN commands from an Initiator that registered with a top-level password.
    - A top-level password survives power off/on.
    - A command changes the top-level password.
    - A command sets a one-LUN second-level password.
- An SCSI LUN accepts LUN commands from an Initiator that registered with the LUN using the LUN's second-level password.
- Initiators that provide passwords to SCSI peripherals are authorized to fully control the peripherals.
- Other "untrusted" Initiators are very limited in commands allowed (more details come later).
- A "third-party SCSI reservation" made by an authorized Initiator permits specific access by an untrusted Initiator.

Capabilities of Untrusted Initiators

- Some commands (REQUEST SENSE, INQUIRY, TEST I/O) seem OK without passwords.
- Some task management functions seem OK.
    - TERMINATE TASK and ABORT TASK affect only the Initiator's functions. (Corrupted media?)
    - CLEAR ACA seems OK.
- These are not OK because they can affect other Initiators (and fixing them does not seem needed):
    - TARGET RESET
    - CLEAR TASK SET

    - ABORT TASK SET

- No Asynchronous Event Notification?

- Lower-level events (FC examples such as FCP Login, FC Fabric logout) appear to be acceptable.

Does extended proposal satisfy our needs?

- Peripherals recognize "their" controlling storage systems when systems provide passwords (yes).

- If addresses change, storage systems cannot easily dynamically find and identify peripherals (concern).

    - Protecting passwords is a concern.

    - Static addresses allow static config tables (yes).

- Peripherals obey only Initiators with password, hence only controlling storage systems (yes).

- When specified by storage systems, peripherals exchange specified data with whatever is at specified addresses (qualified "yes").

    1 READ/WRITE with third-party RESERVE.

    2 COPY command (doesn't even need RESERVE).

Remarks on how to identify peripherals.

- Not a serious issue if addresses are static.

- Commands like INQUIRY can be allowed without passwords. They neither prevent nor facilitate attacks on security but help detect non-malicious changes.

- A storage system can write a unique unguessable data pattern on a storage peripheral or medium.

- The storage system can later retrieve the data pattern. If it is not correct, the peripheral may be misidentified (and passwords may be disclosed).

    - This assumes the medium is not removable.

Remarks on protecting the password

- Protect all paths between storage system and its peripherals from snooping.

- Ensure storage system never sends password to wrong destination. Hard with dynamic addresses.

    - Assume done somehow based on lower-layer mechanisms, e.g., static addresses or fabric provides configuration information to storage systems.

- We are looking at ways to ease installing new passwords (if lost or disclosed, or to ease re-sale).

- Three-state switch in peripheral: "disabled", "acquire one password", and "running".

Perspective

Alternative approaches are being pursued.

- LLNL SIOF NAP converter (our main current effort)
- CMU/NSIC Network Attached Storage Devices

Status of LLNL NAP effort

- Plan to demo (remotely) during Supercomputing 1995.
- Some data has been at least written.
- Initial implementation is to front-end ordinary peripherals with workstation (and switched Fibre Channel). Workstation provides NAP functions.
- This effort remains IP oriented.
- Have recently decided to allow byte-aligned transfers, not restrict to blocks.