An Update on Distributed I/O at LLNL

Lansing Sloan

Presented to SCSI Working Group

July 11-12, 1995

Lawrence Livermore National Laboratory

7000 East Avenue, M/S L-60

Livermore, CA 94550-9900 USA

e-mail: ljsloan@llnl.gov

phone: (510) 422-4356

fax: (510) 423-8715

Overview

Goals

We are de-emphasizing SCSI: why?

Emphasis now: third-party data over TCP/IP

Alternative: "network attached secure discs"

Comparison of approaches

Security remarks

Scalable I/O Facility (SIOF) Goals

- Hierarchical storage management (HSM) and high, scalable bandwidth
- Practical solution to the I/O bottleneck
  - Scalable parallel HSM I/O for MPPs and clusters
- Reduce storage cost
  - Mass-market components
  - Network attached peripherals (NAPs)
- Parallel WAN access to stored data

SIOF - Key Technologies

- Switched Fabric (Fibre Channel, ATM, Serial HIPPI, or FDDI)
  - Parallelism and scalability
  - Fast Date Rates
  - Small form factor interconnect
- Network Attached Peripherals (NAP)
- Security for NAPs (and for WAN)

Network Attached Peripheral (NAP)

- Peripheral with network interface, security mechanism and control protocol to allow it to be fully controlled by a remote File Server.
- Advantages
  - eliminates workstation server - reduces cost and system administration
  - avoids unnecessary data copies - global file system
  - scalable I/O performance if using switched fabric

Network Attached Peripherals reduce costs

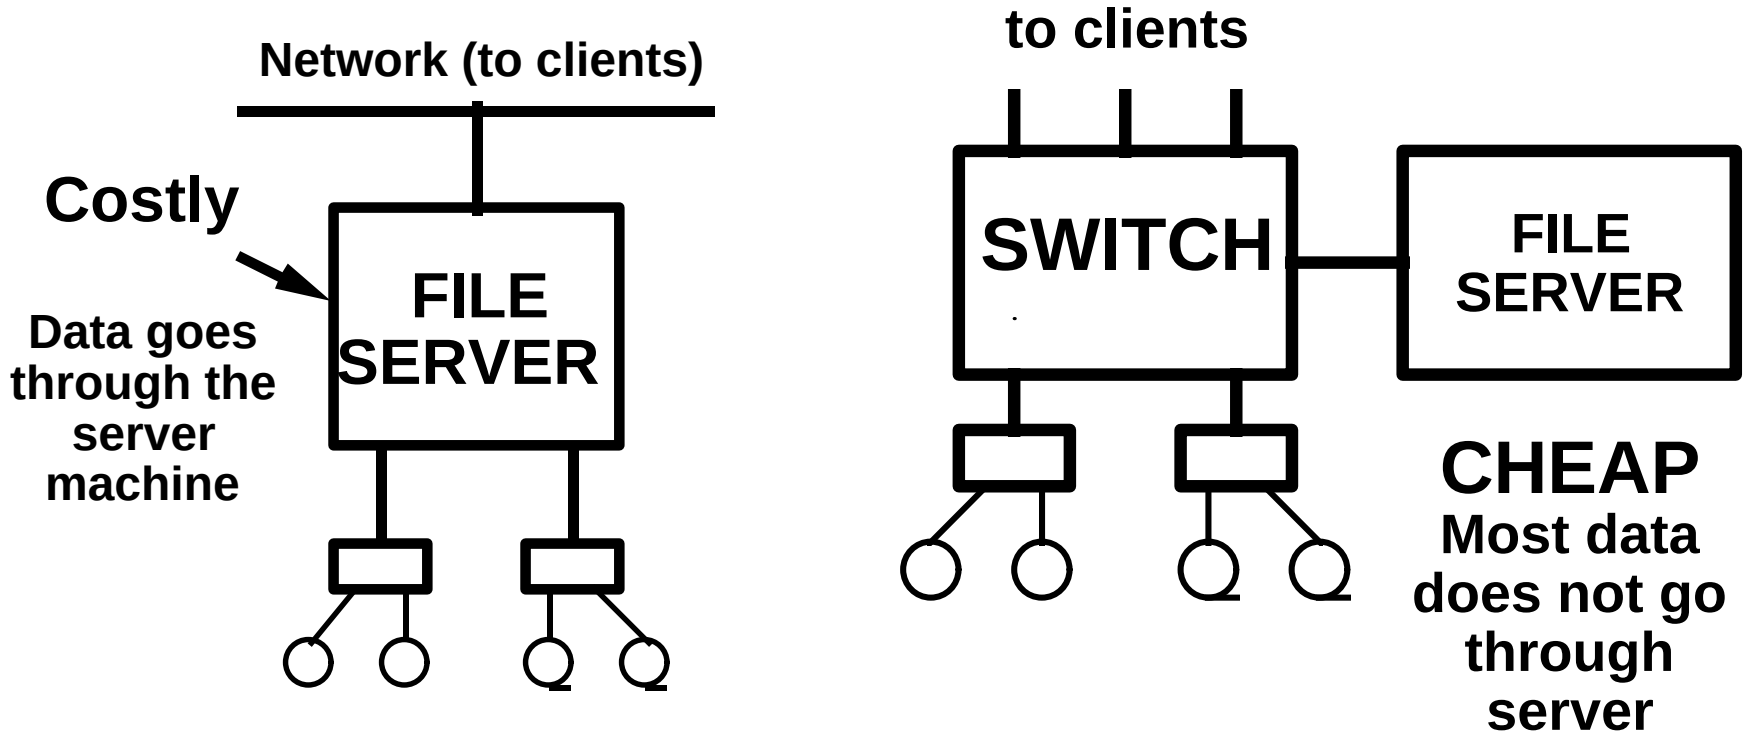

Deliverables for 1995

- Demonstrate at SuperComputing 1995 scalable parallel network-centric peripheral systems with a Meiko CS-2 MPP accessing
  - tapes in parallel (front-ended by workstations)
  - discs in parallel (front-ended by workstations emulating network attached peripherals)
- Plan deliverable for HPSS Release 3
  - Specify NAP behavior, start coding

Problems with SCSI and Fibre Channel

- Cannot use SCSI for SC95.
  - Class-2 FC switch not available in time
  - Class-1 FC SCSI support not available in time
  - Switched-FC SCSI support not ready in time
- Long-term opportunity
  - Industry support switched-FC SCSI with multiple vendor interoperability

We are emphasizing TCP/IP now

- HPSS already supports data transport over TCP/IP.
- Meiko supports TCP/IP with switched Class 1 FC (tested) but needs obsolete FC version.
- We might have to use an IBM RS6000, which might not support the obsolete FC version.
- Fall-backs are to use TCP/IP on ATM (tested with Meiko) or other media.

First Generation Implementation

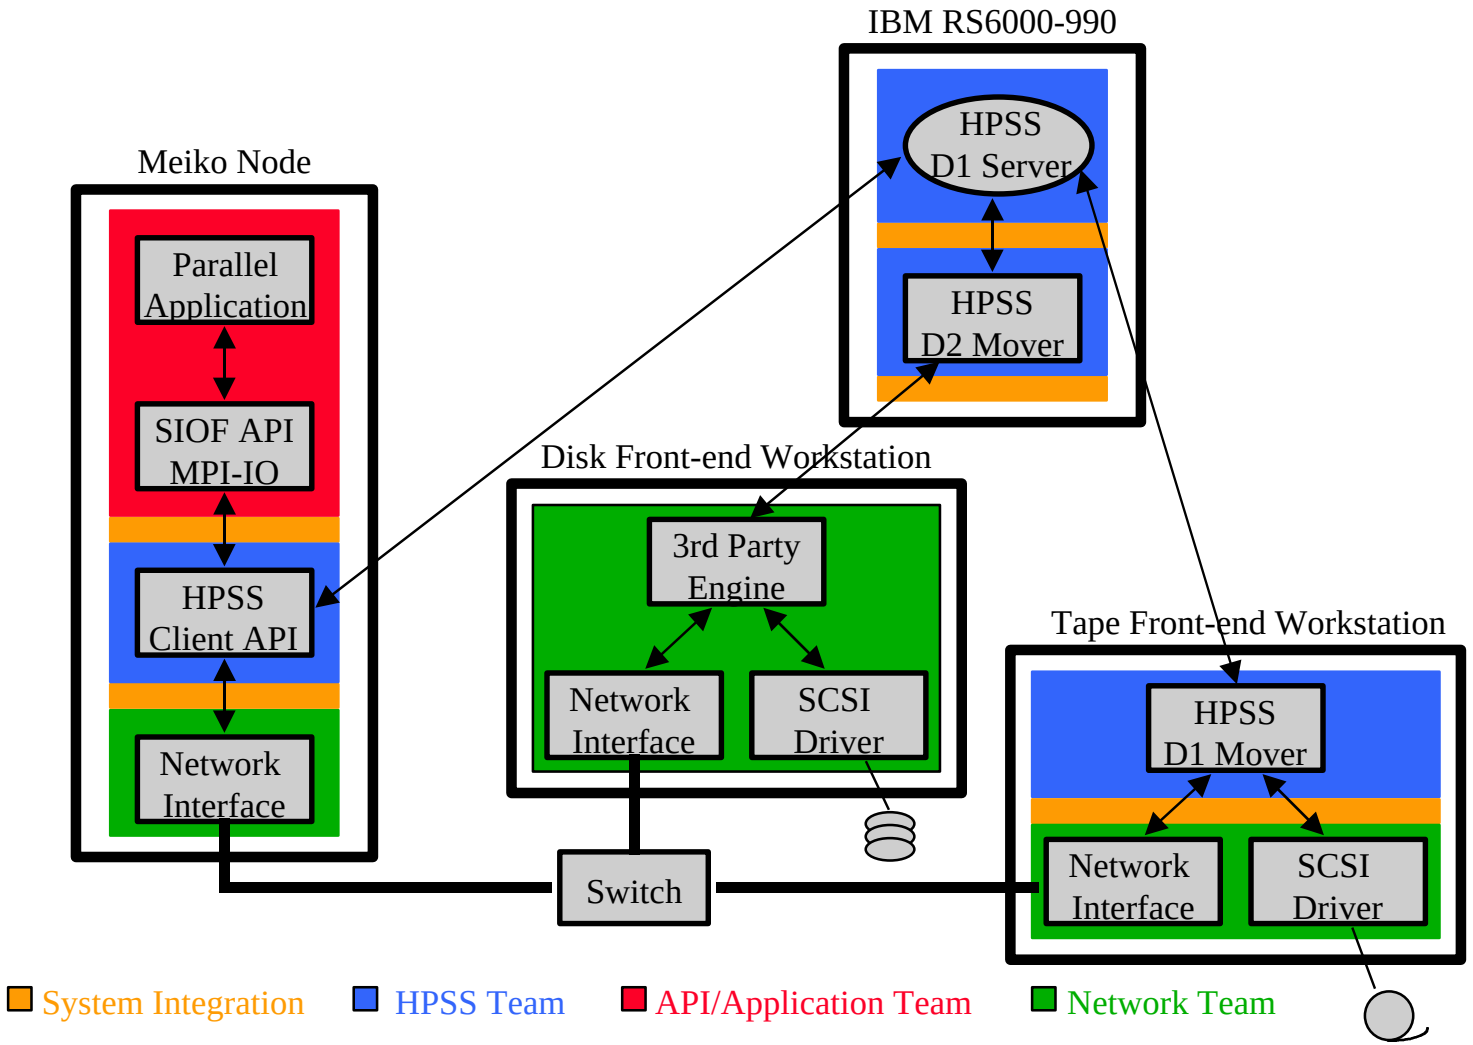

NAP Development Plans

- Phase 1 - SC ‘95
  - Workstation front-end NAP converter
  - Investigate command protocols
  - Begin prototyping security mechanisms
- Phase 2 - FY95
  - NAP converter using off the shelf PCI board
- Phase 3 - FY96
  - Custom PCI board for maximum performance
  - Prototype hardware security mechanisms

Embedded NAP Converter

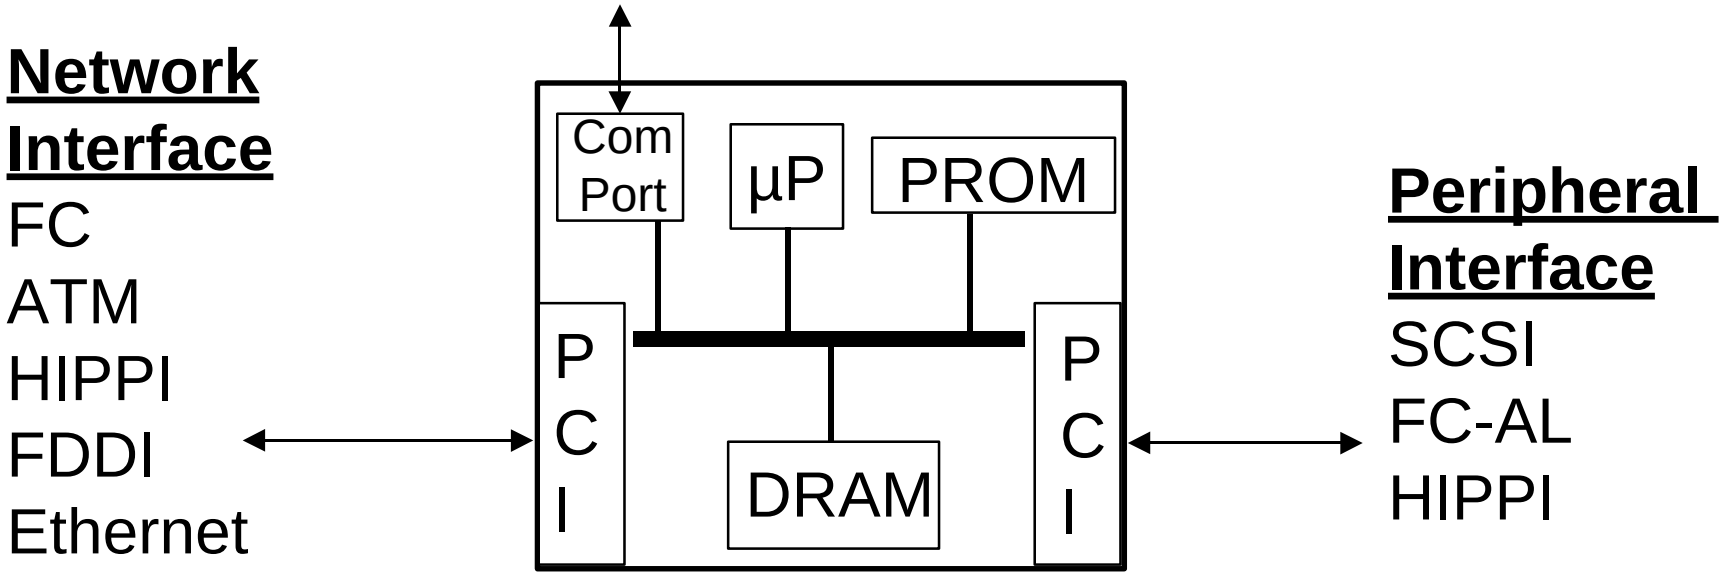

Other approach: Network Attached Secure Discs (NASD)

- Proposed by Garth Gibson at National Storage Industry Consortium, May 1995
- Discs are directly on switched networks.
- A disc has its own file system.
- Discs include security for file systems.
- Applications access files directly.

LLNL is interested in NASD

+ Real potential for secure NAPs
+ Should allow efficient application I/O
- Early to know what will result
- Early to know about industry support
- Direct access by applications interferes with hierarchical storage management
- Each disc must support site security policies
- Cost and performance for security?

We remain interested in SCSI

+ Current direction of peripheral mass-market
+ High performance
+ Cost effective
- Current market supports no switches
- Need some security enhancements
- Limits physical layer - not on ATM, FDDI, ...
- Probably need gateway to WAN environment

For now, LLNL will continue with TCP/IP NAP

+ Supports almost all physical media
+ Provides WAN connectivity
+ Reduce current HPSS costs
+ Already supported by industry in NAP like devices - NFS & FTP servers
+ Easy to prototype, software implementations
+ Industry is working hard on TCP/IP
  - security (big electronic commerce market)
  - performance

TCP/IP has some disadvantages

- Possibly slower, less cost-effective than SCSI
- Might need WAN and/or security gateways to lower costs or optimize WAN use
- Cost and performance for security?
- If costs more than SCSI, TCP/IP may be a stopgap solution

NAP security

- Goal: Ensure storage server fully controls NAPs (secure reliable control path).
- Objectives
  - Performance - minimize latency
  - Minimize modifications/additions to peripherals - mass-market
  - Minimize configuration management
  - Portable - not media dependent (desired)
  - Extendible to WAN environment (desired)

NAP security: control path

- Physically separate control net (current HPSS)
  - Simplest but costly - not mass-market
  - WAN environment?
- Use security features of fabrics (switch)
  - Requires the least amount of change to peripheral
  - Configuration management difficult
- Encryption
  - Costly addition to peripheral - not mass-market?
  - Performance?

NAP security update

- Snooping may disclose parameters that can be used to mount attacks
  - SCSI is vulnerable to forgery
  - Solution: confine SCSI domain in safe area
- Snooping facilitates "replay" attacks
- Rogue clients may pass address parameters
  - Hard to detect
- Site-specific security policies

NAP security analysis is hard

- System architecture not complete - HPSS WAN
- Security requirements not clearly defined
  - Exact needs vary from site to site
- Need to prototype to effectively analyze solutions

NAP security: some conclusions

- Peripheral must participate - placing more responsibility in peripheral simplifies security
  - Minimum: NAP must be fully controlled by server
  - Do not want site-specific policies in all peripherals: mass-market and cost to configure
- Distributed systems allow for more attacks

Summary: LLNL's current efforts

- Focusing on TCP/IP NAP at present
- Investigating security issues
- Plan to participate in Network Attached Secure Disc effort
- Seeking industry partners for secure NAP