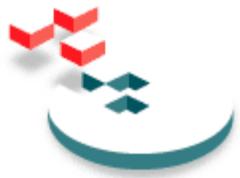


SNIA IP Storage Forum

IP Storage Case Study: SANDIA National Laboratories

Using Internet Fibre Channel Protocol (iFCP) and
Internet SCSI (iSCSI) for IP Storage solutions





Agenda

- The ASCI Tri-Lab
- ASCI Project Overview
- The ASCI HSI Research and Development Objective
- Storage Requirements
- Storage over IP Technologies
- VERITAS ServPoint SAN Software
- SuperComputing 2001
- Lessons Learned
- What's Next ...



About ASCI and Sandia



Nuclear Weapon Stewardship via software simulation requires accelerated advances in compute, storage, visualization, and network technologies

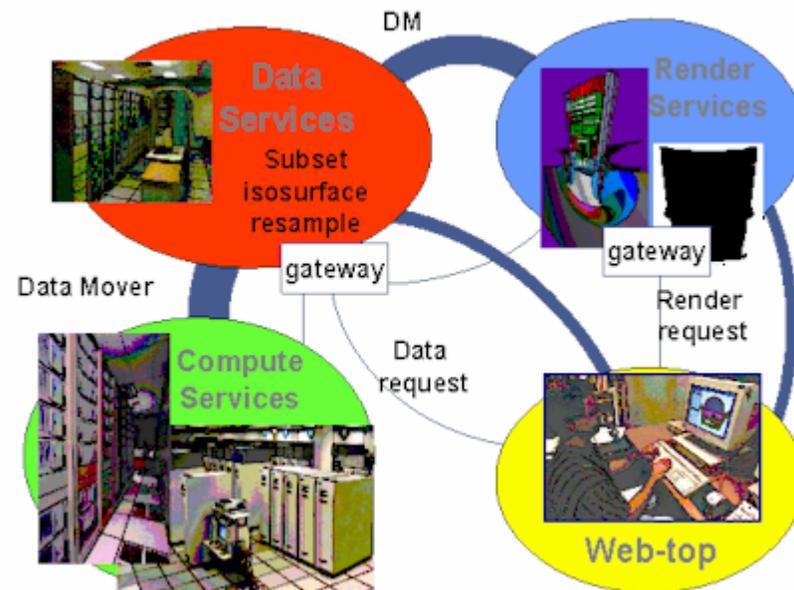


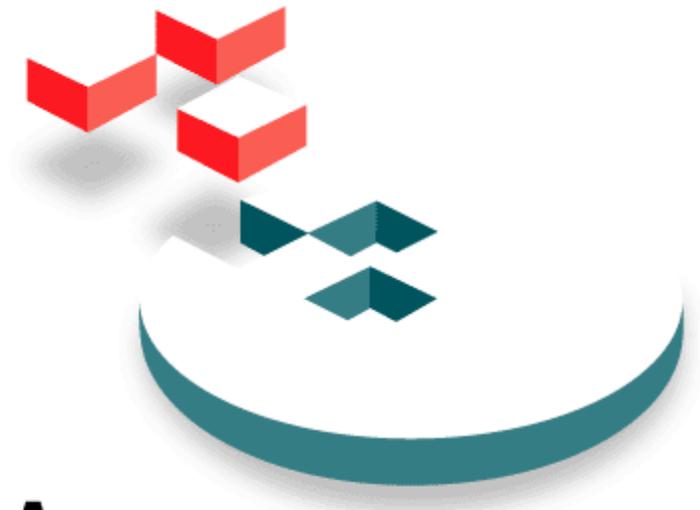
Sandia is a multi-program laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy under contract DE-AC04-94AL85000.



Challenges

- Large scientific calculations generate hundreds of terabytes of data
- Tri-Lab resources are geographically distributed
 - Compute platforms
 - Graphics engines
 - Storage
 - Scientists





SNIA IP Storage Forum

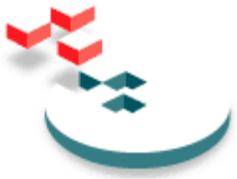
The Tri-Lab
Advanced Scalable
Computing Initiative (ASCI)





The ASCI High Speed Interconnect Project

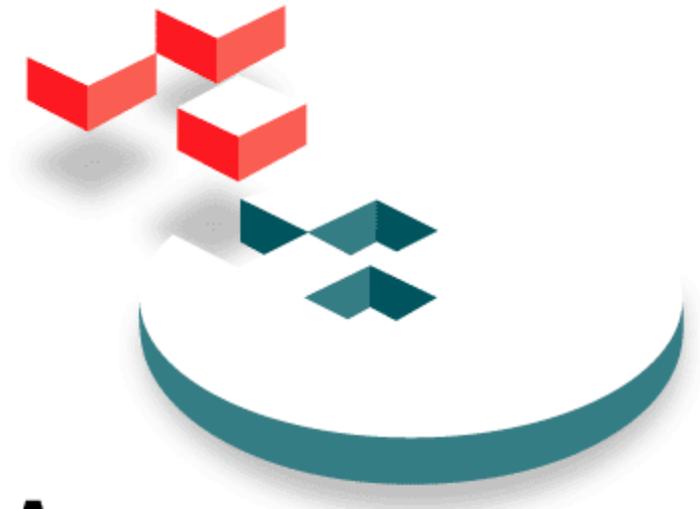
- Delivery of high throughput, low latency network and storage performance to the ASCI tools and applications
 - ASCI Tri-Lab experiments indicate existing applications and protocols utilize 50% of the available bandwidth across gigabit networks
 - The disparity between physical bandwidth and “deliverable” end-to-end bandwidth continues to be one of the most challenging problems faced in building networks of ASCI components
 - Collaboration between HSI, DISCOM, VIEW and SIO teams to develop and promote standards efforts for delivering COTS technology that meets the bandwidth, latency and scalability requirement of ASCI applications



Tri-Lab Industry Participants



* SNIA IP Storage Forum Members



SNIA IP Storage Forum

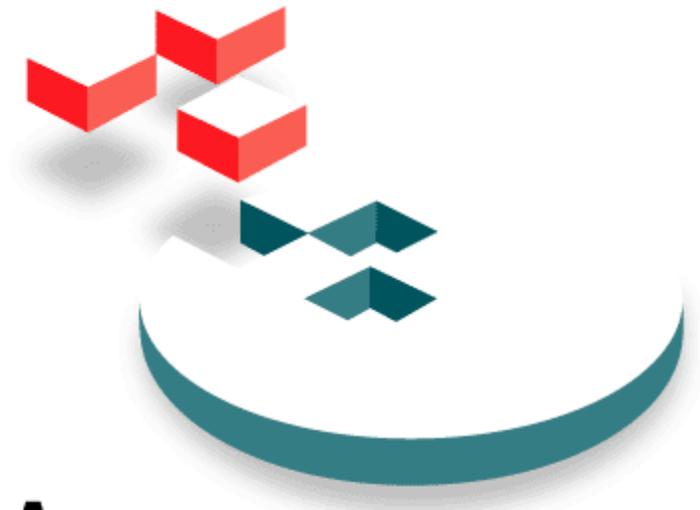
Tri-Lab Test bed
Research and Development Objectives





Tri-Lab Research & Development Goals

- Efficiently move terabytes of data over the SAN, LAN and WAN for interactive scientific visualization
- Demonstrate a robust network architecture
 - Automatic failover and recovery
 - Dynamic bandwidth provisioning
 - Quality of Service (QoS)
 - Interoperability between storage and network services
- Demonstrate next generation storage technologies
 - Emerging storage standards (iSCSI & iFCP)
 - Storage virtualization over the SAN, LAN & WAN
 - Centralized storage management
 - Quality of Storage Services (QoSS)



SNIA IP Storage Forum

Storage Requirements





Selecting Storage Partners

- Interoperability with a variety of platforms
 - SGI
 - Linux
 - Sun
 - Data Direct Networks
- Ability to meet technical requirements of project
- Industrial partnerships can help facilitate the design and implementation of a storage system



Storage Virtualization Solution – ServPoint SAN

- **BLOCK Level Storage**
 - Independence from server OS
 - Independence from the File System
 - Allowed for interoperability with SGI & DDN
- **Quality of Service – Server Cache**
 - Supports Caching for Reads
 - Adaptive Read-Ahead per Virtual Disk
- **Dynamic Storage Provisioning**
 - Striping across multiple disks and/or multiple arrays maximizes I/O throughput
 - Online relocation eliminates performance bottlenecks
- **Storage over IP technology supported**



Robust WAN Architecture

- AVICI SSR & TSR
 - Composite Link
 - One logical link can consist of up to 64 physical interfaces (OC3c-OC192c)
 - One IP address - small routing table
 - Load balance while maintaining packet ordering
 - Failover restoration within 45 ms
 - Dynamic bandwidth adaptation through addition or removal
 - of link member
- Spirent Adtech Broadband WAN Emulator
 - Emulates error and delay
 - CA to NM (~22 ms RTT)



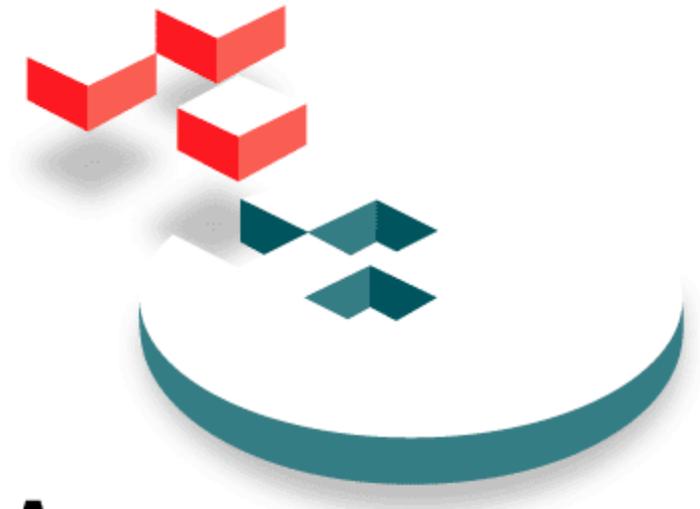
Robust WAN & LAN Architecture

- Riverstone Switch/Router
 - Load balancing and self healing through OSPF Equal Cost Multi-Path or 802.3AD Ethernet link aggregation
 - Ordered packet delivery within each TCP micro-flow
 - QoS using 802.1p, MPLS
- Nishan 4300 IP Storage Switch
 - IP load balancing and self healing through designated backup port
 - Ordered IP packet delivery within each TCP micro-flow
 - Conversion of FC end systems to iFCP or iSCSI
- Cisco SN 5420
 - iSCSI gateway
 - Gigabit Ethernet Host access remote FC storage in block mode



Robust SAN Architecture

- Brocade Silkworm 6400 FC Switch
 - Automatic data routing and re-routing of traffic
 - Heterogeneous device connectivity
- VERITAS ServPoint Appliance Software for SAN's
 - Quality of Service for storage devices
 - Local and remote virtualization of the storage devices
 - Block level I/O
 - Dynamic storage provisioning
 - Software implementation of iSCSI
- VERITAS SANPoint Control
 - Configuration, management and monitoring of all SAN devices



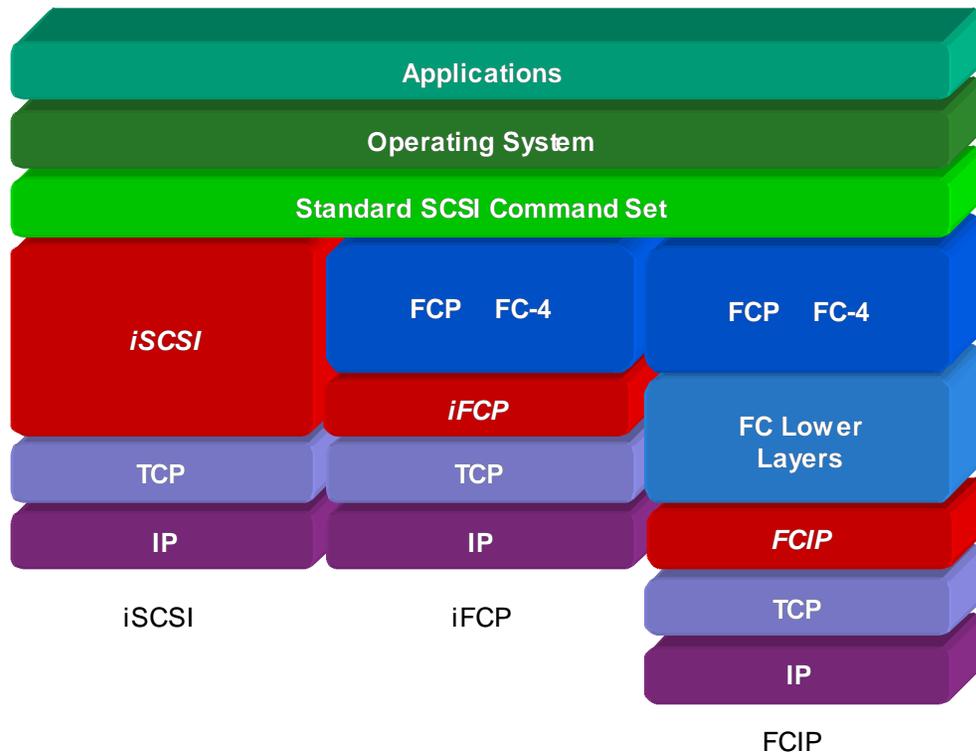
SNIA IP Storage Forum

Storage over IP

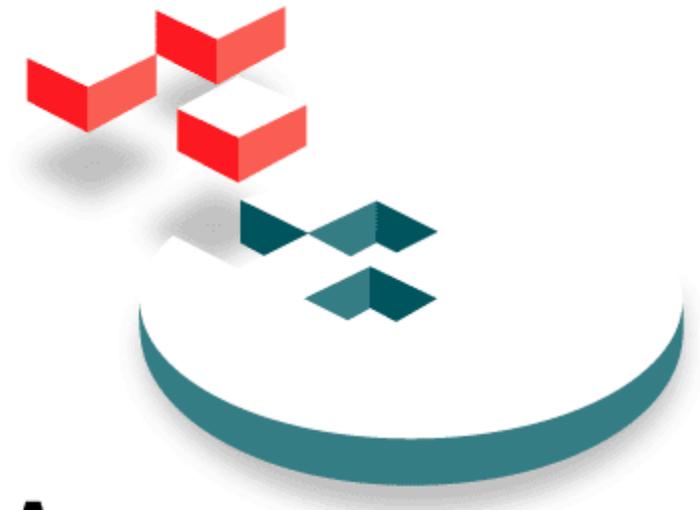




Emerging IP Storage



- iSCSI maps SCSI to TCP
- iFCP maps SCSI and FCP to TCP
- FCIP tunnels entire FC stack over TCP



SNIA IP Storage Forum

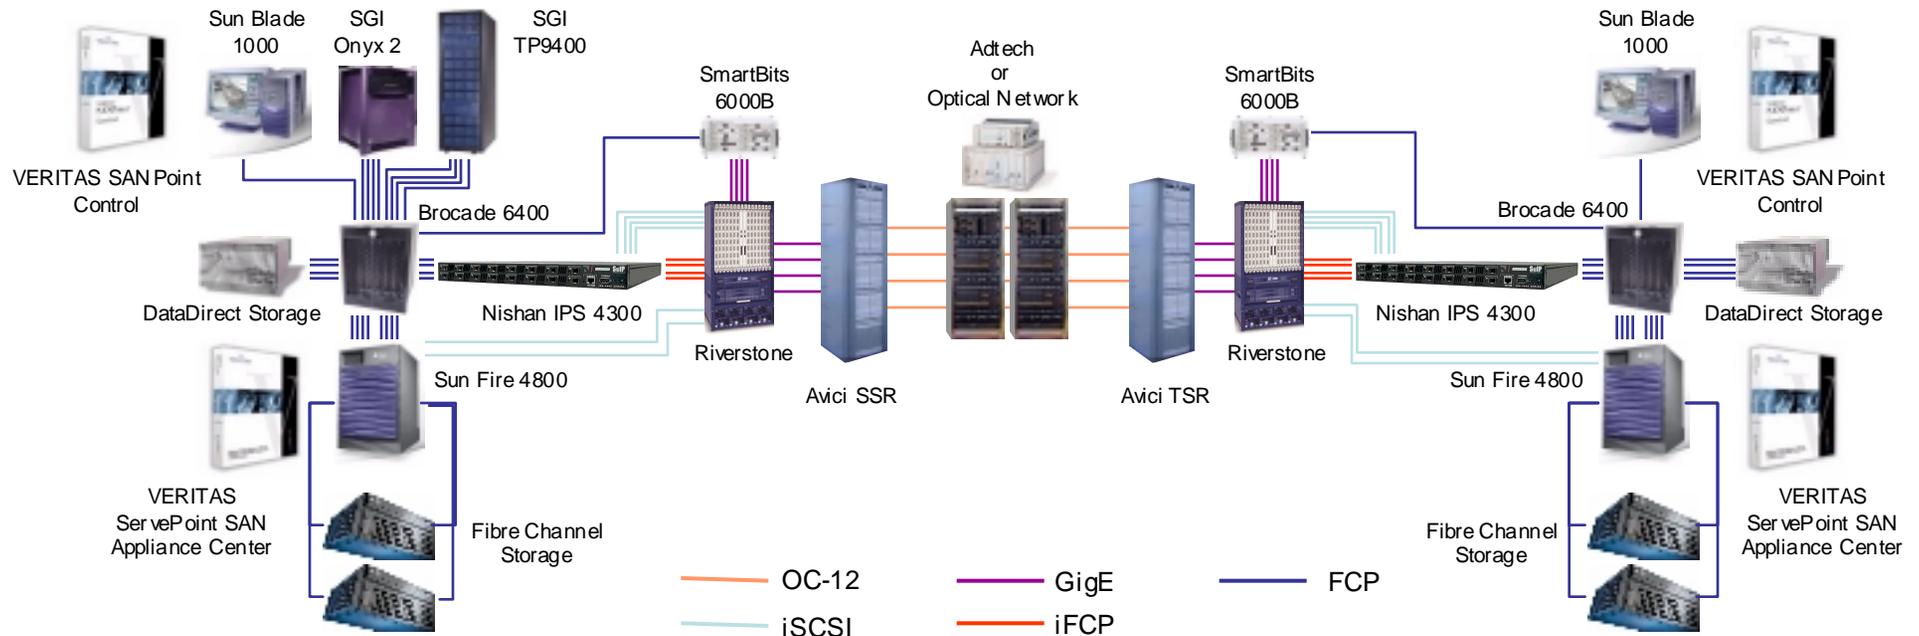
SuperComputing 2001





The Architecture

Sandia National Labs Next-Generation Storage and Networking Project

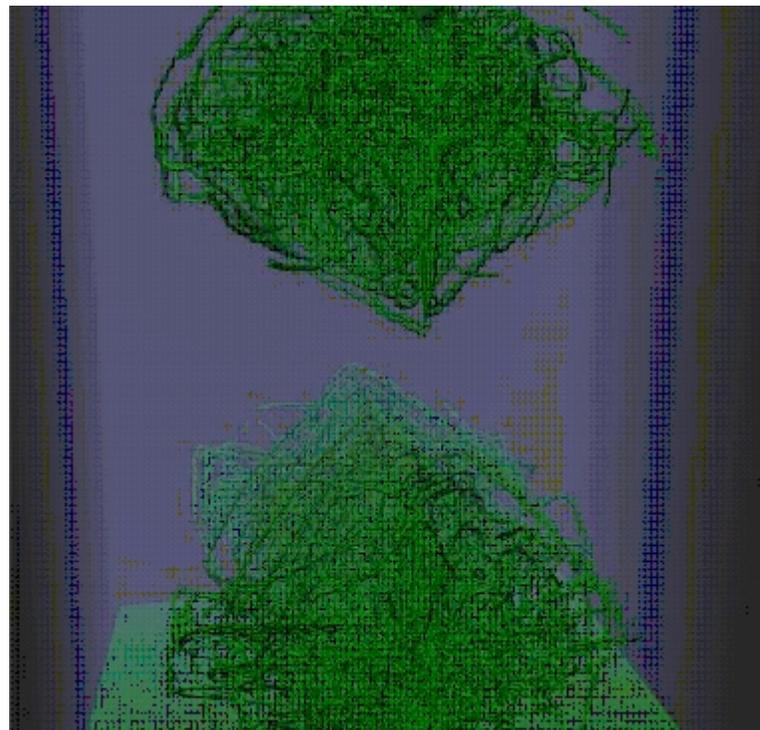


- Demonstrated high-throughput network storage solution
- Implemented a scalable distributed computing environment over TCP/IP architecture
- A collaboration between Sandia, VERITAS and 10 industry partners



Visualization Demonstration

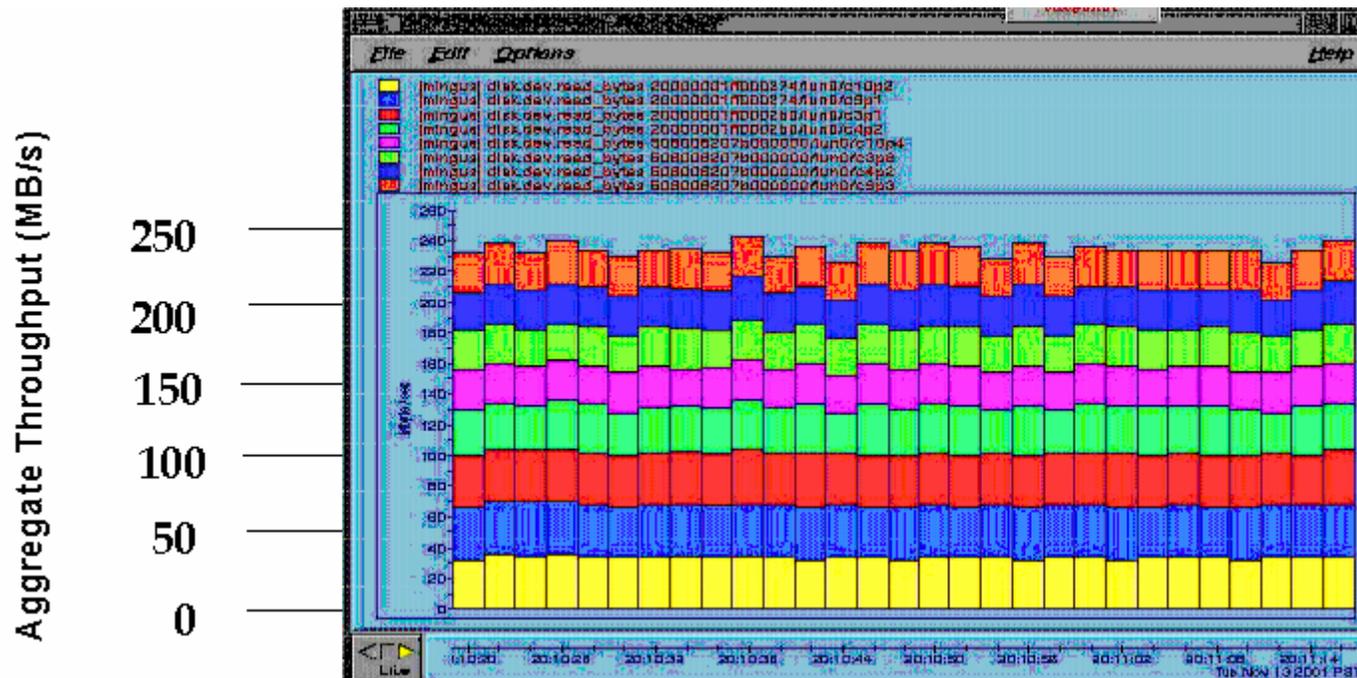
- The animation of a crack propagation in matrix of 2 billion Cu atoms
 - 21 GB of data representing 1399 frames at 15.7 MB per frame
 - Planes of dislocation are Cu atoms of highest energy level





SuperComputing 2001 Results

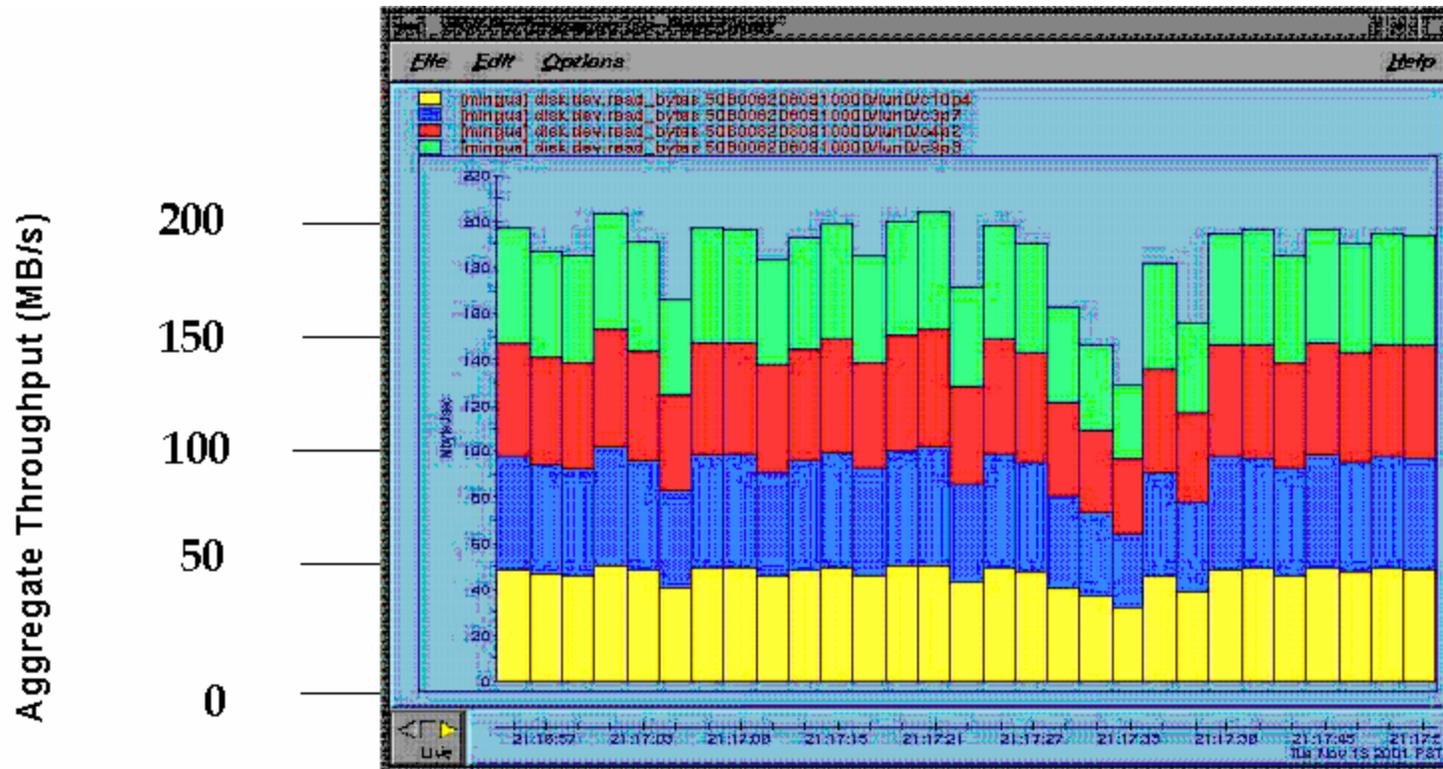
- SGI 64 processor ONYX server
- 8 VERITAS ServPoint SAN Virtual Disks
- 4 OC-12's links
 - Rate limited at 60 MB/s per OC-12
- Aggregate throughput = 240 MB/s





SAN Performance Results

- Four VERITAS ServPoint SAN Virtual Disks
- SGI File System striped across all four virtual disk
- Aggregate throughput = 180 MB/s





LAN Performance Results

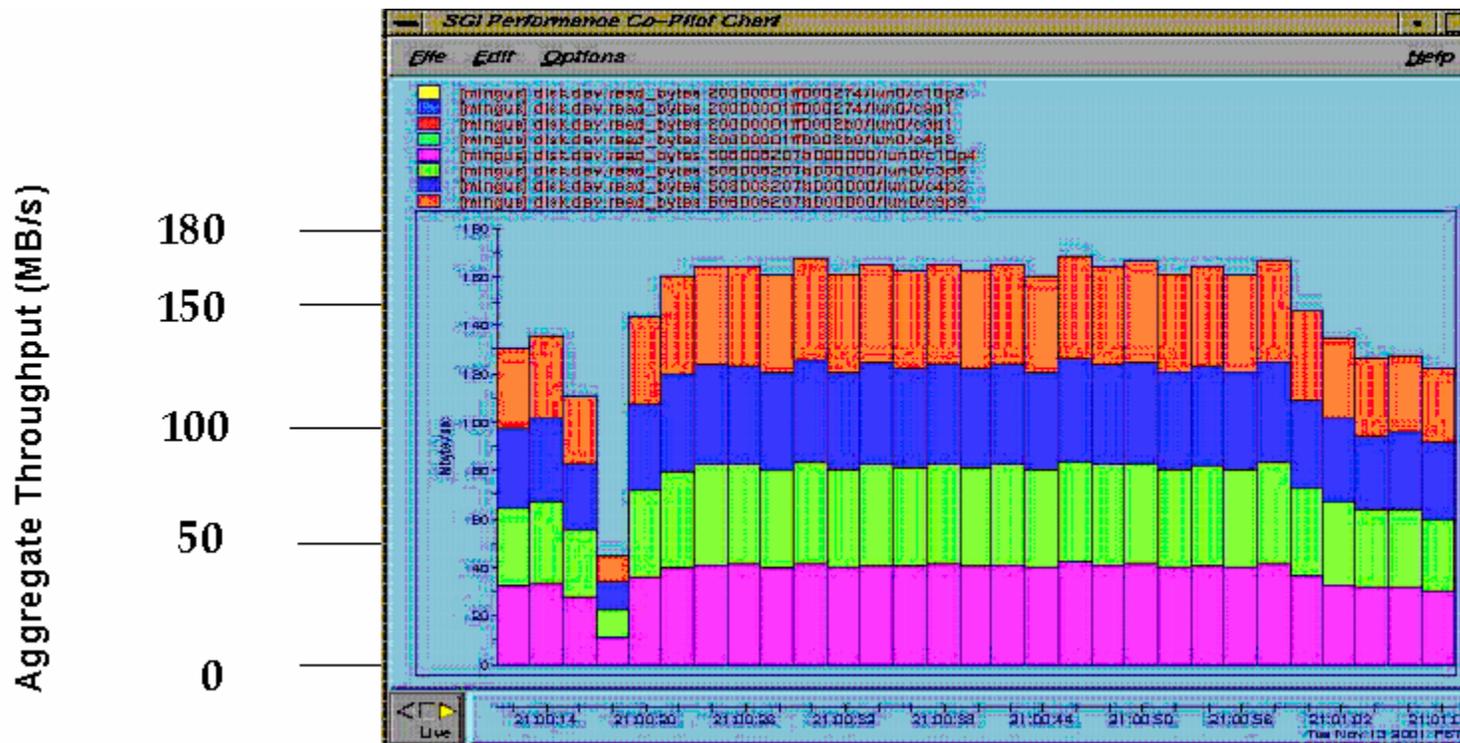
- Four VERITAS ServPoint SAN Virtual Disks
- SGI File System striped across all four virtual disk
- 4 – OC-12's with a 0ms delay
- Aggregate throughput = 160 MB/s

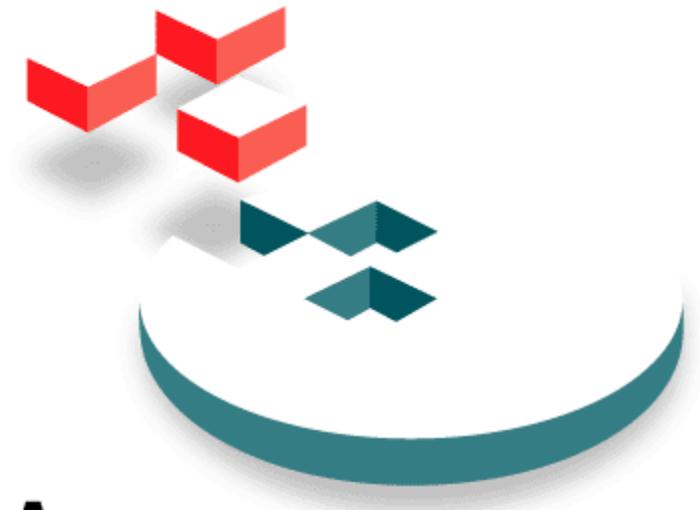




WAN Performance Results

- Four VERITAS ServPoint SAN Virtual Disks
- SGI File System striped across all four virtual disk
- 4 – OC-12's with a 22 ms delay
- Aggregate throughput = 160 MB/s





SNIA IP Storage Forum

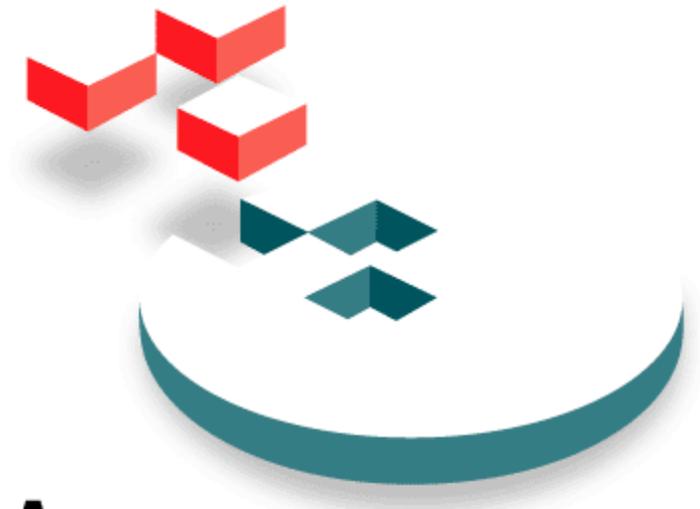
Lessons Learned





What do the results mean?

- Demonstrated a practical application using virtualized storage over an IP network
- iSCSI software showed reasonable performance
- Solutions available for interoperability between Fibre Channel and IP storage
- Networked storage has the potential to delivery comparable performance against direct attached storage
- Storage over IP technologies are rapidly improving



SNIA IP Storage Forum

What's Next...





Next Steps

- Define additional performance and test scenarios for ASCI projects
 - File-level performance tuning using VERITAS ServPoint SAN
 - Convergence of file-, storage-, and network-level Security
- Demonstrate Quality of Service from Storage through the Network
- Begin defining goals and requirements for SuperComputing 2002
- Possible adoption of this architecture into the ASCI Tri-Lab production network



Acknowledgements

- Special Thanks to the Team at Sandia National Laboratories for Contributing this Information
 - Helen Chen
 - Dov Cohen