LeftHand NETWORKS

Network Storage

When does it make sense?

LeftHand NETWORKS

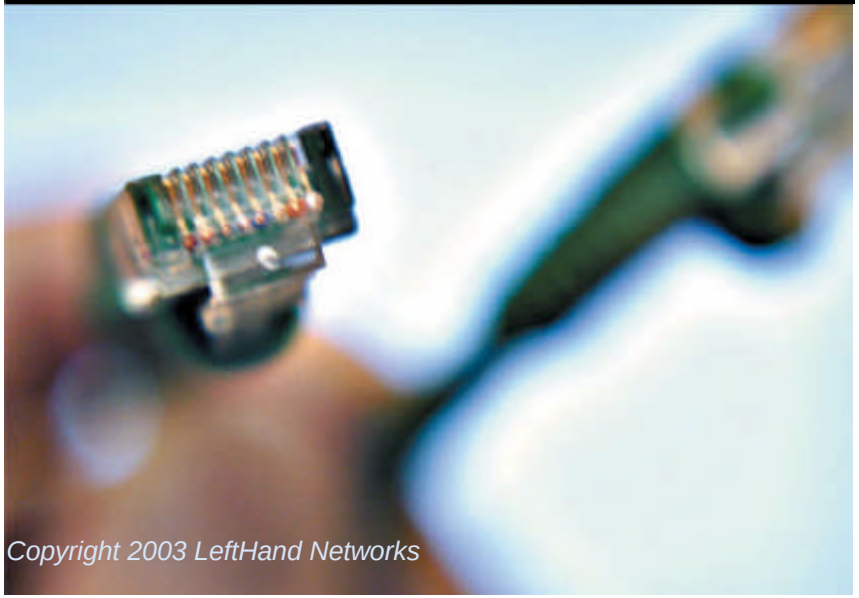

John Spiers - CTO

Storage as it should be

The Growth in Storage

- *More new information will be created over the next two years than over the entire history of humanity - more than 90 percent of it digital.*

Gartner-2002

Storage Management Challenge

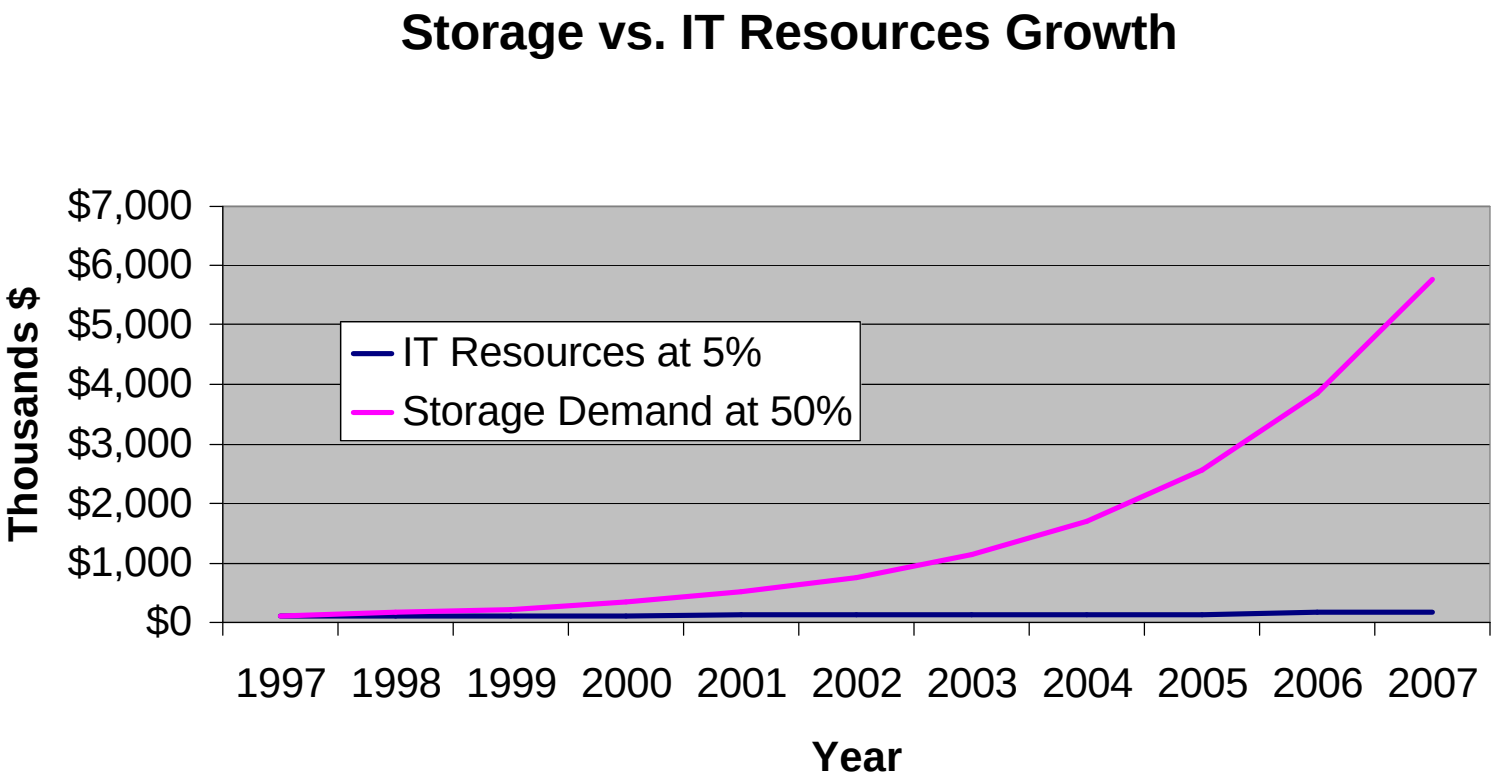

IT Resources must increase their efficiency to manage 35x more storage over the next 10 years!

The problem with Storage today?

- Problems with today's storage model
  - **Rapid decline in the cost of computing drove the decentralization of computing and the proliferation of Servers**
  - **Still isolated or task-specific computing**
  - **Increase in fragmentation of data and applications across the enterprise**
  - **Islands of storage where utilization is typically less than 50%**
  - **Storage scalability limited by server performance and design**
  - **Storage expansion often requires adding or upgrading servers**
  - This increases the complexity of managing the data storage

Today's Computing Architecture

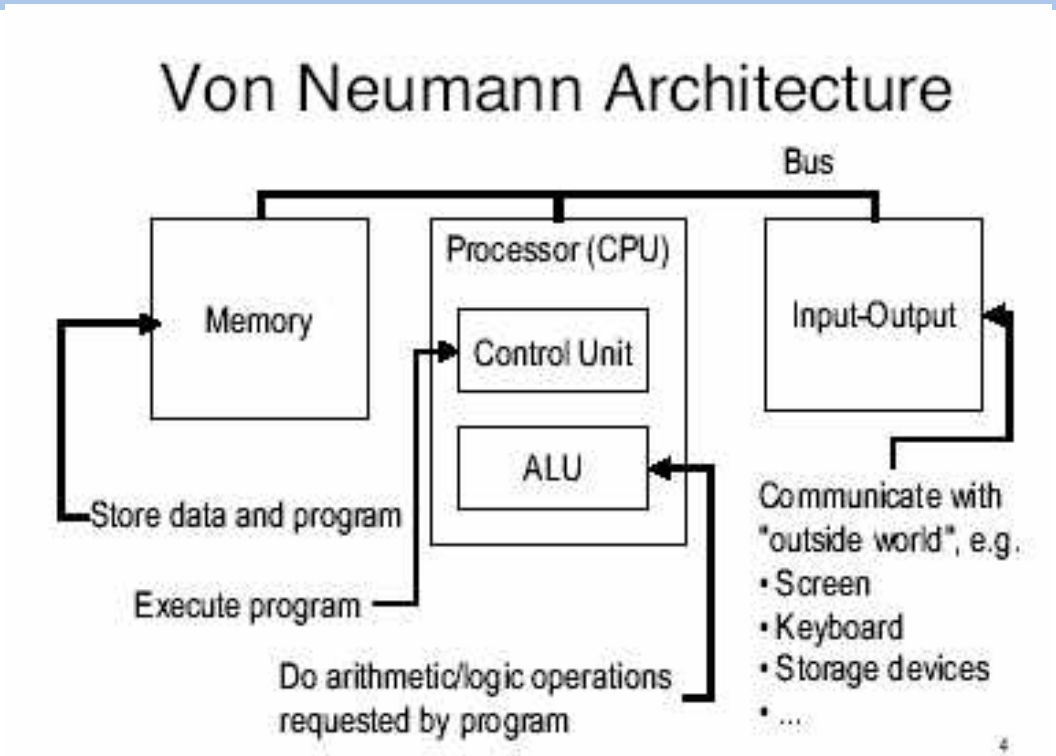

- First computer built using this model was the EDVAC computer in 1952
- Today's computer and storage architecture is fundamentally the same

Storage & Computing Trends

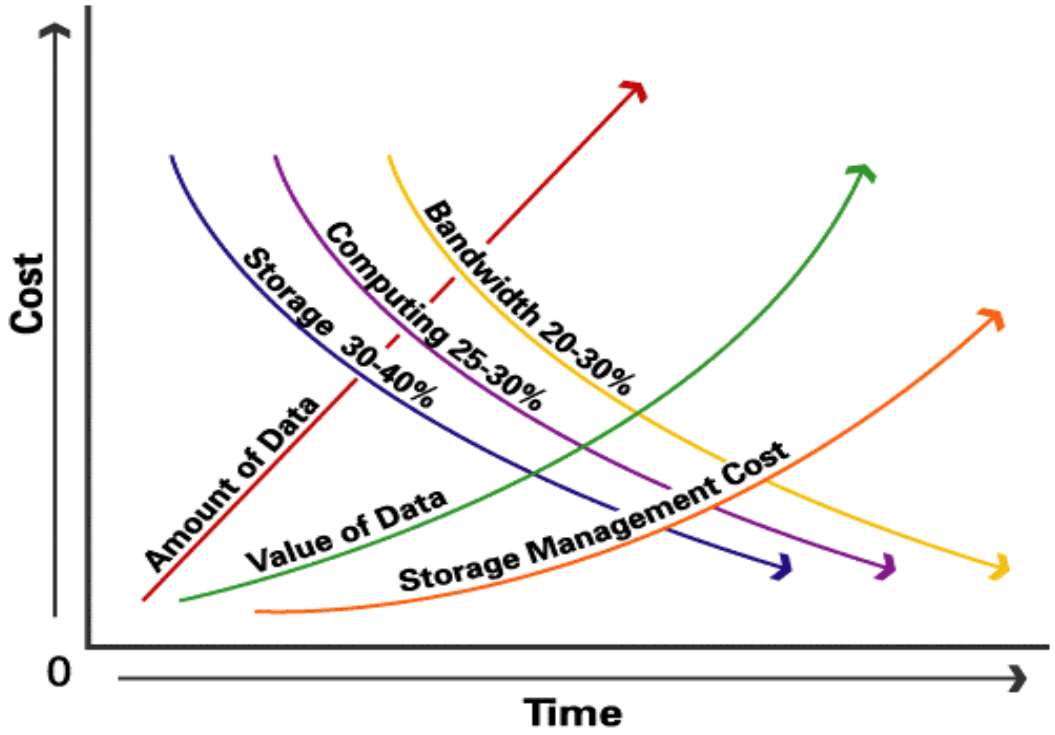

The Perfect Storm

Storage computing and bandwidth are all abundant

Digital data growing 70-80% CAGR

Value of data growing exponentially

Technology has changed….

- In the 90's the internal buss of the computer ran at 100MB/sec, the network ran at 1MB/sec.
- There has been an explosion in network speed that has created a technological "discontinuity".
  - Today's Networking technology (10Gbit) is faster than the internal computer buss
- The future: Seamless bridges between system and network providing:
  - Dynamically configurable connections to storage and other I/O
  - Consolidation of chassis backplanes to a single fabric type
  - Simplification of multi-node computing architectures

IP/Ethernet As Disruptive Technology

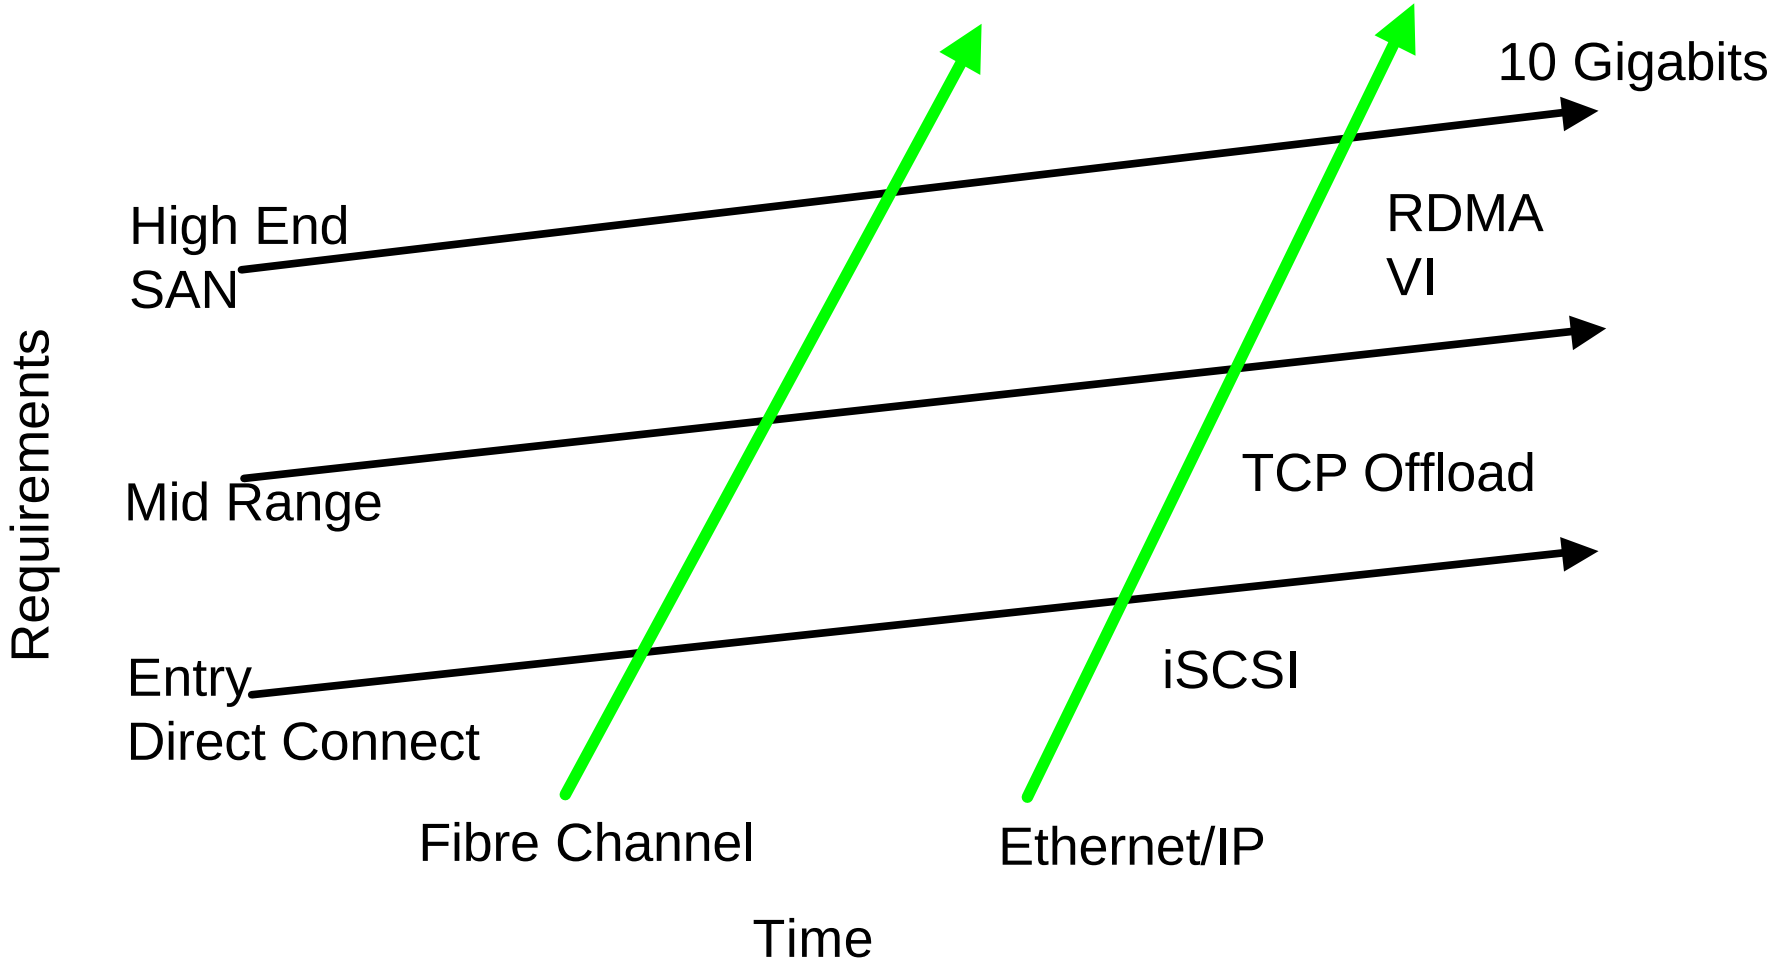

Adapted from Clayton Christenson, The Innovator's Dilemma

Today's Choices

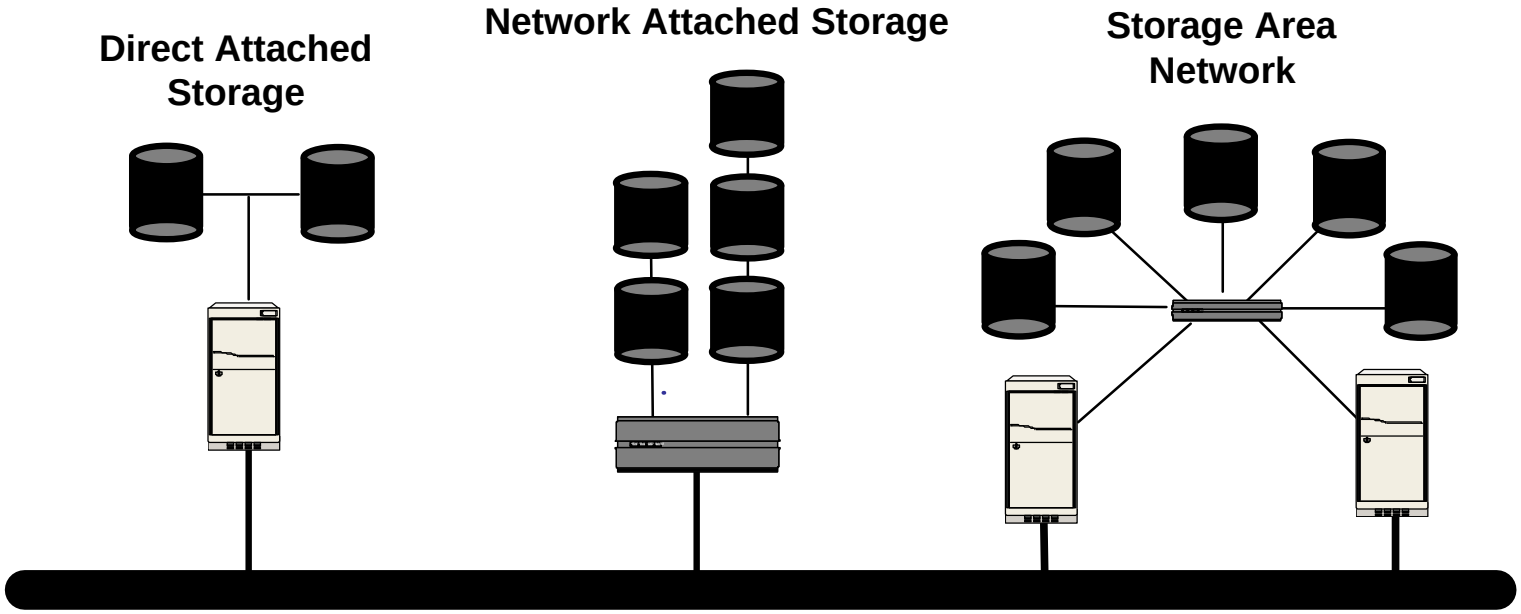

- Not available if the server is busy or down
- Cannot scale beyond the server performance limits
- Difficult to manage across multiple servers

- Head becomes bottleneck
- Multiple NAS heads adds complexity
- File (static) only

- Requires dedicated FC network
- Requires new tools and training
- Management intensive
- Distance limitations
- Block (dynamic) only

Today's Storage Landscape

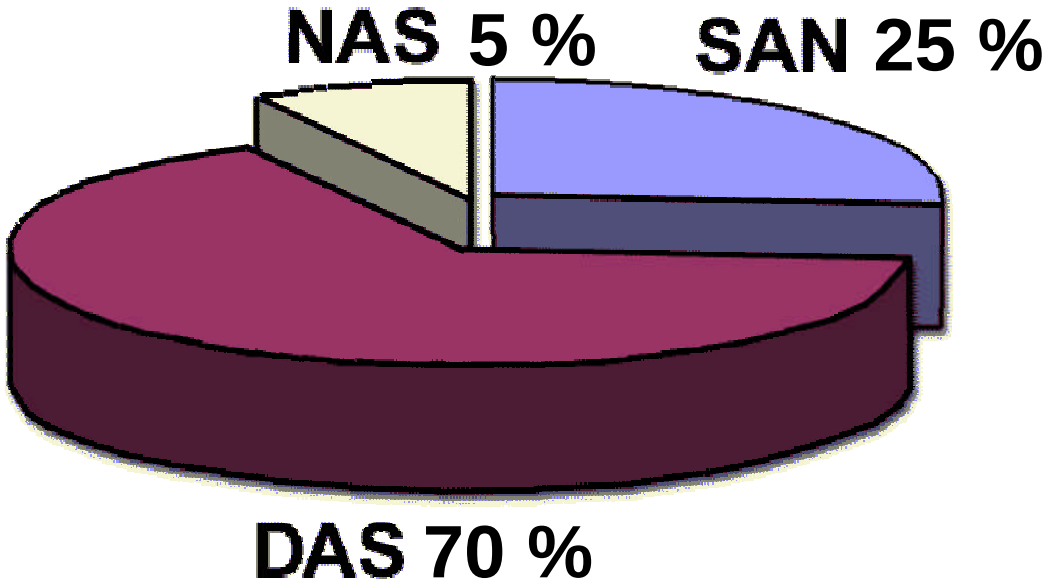

What is iSCSI

- Technology: Standard format for block-oriented storage on IP
- Market: Propelled by expected steady, sharp price declines of IP networking equipment
- Driver: Potential buyers expect rising competition
- Driver: Flexibility in using same equipment for LANs and SANs
- Driver: Particularly attractive with inexpensive drives
- Trend: The convergence of storage and data networking
- Implication: A disruptive technology – opportunity for some, threats to others

IP-Storage vs. Fibre Channel

Why IP Storage?

- Performance will surpass Fibre Channel:
  - Next generation TCP/IP offload engines with RDMA (memory to memory) data movement
  - 10 Gig-E
- No distance limitation
- Ethernet is a widely deployed and well understood technology
- Not just the Fortune 200 (as is Fibre Channel)
- Companies do not have to retrain for TCP/IP networks
- Allows the creation of a single network using familiar standards
- Brings Interoperability & Ethernet economics to storage
  - Fibre Channel will never reach Ethernet economies of scale
- Enables, any server to any storage access

Flexible IP SAN Storage

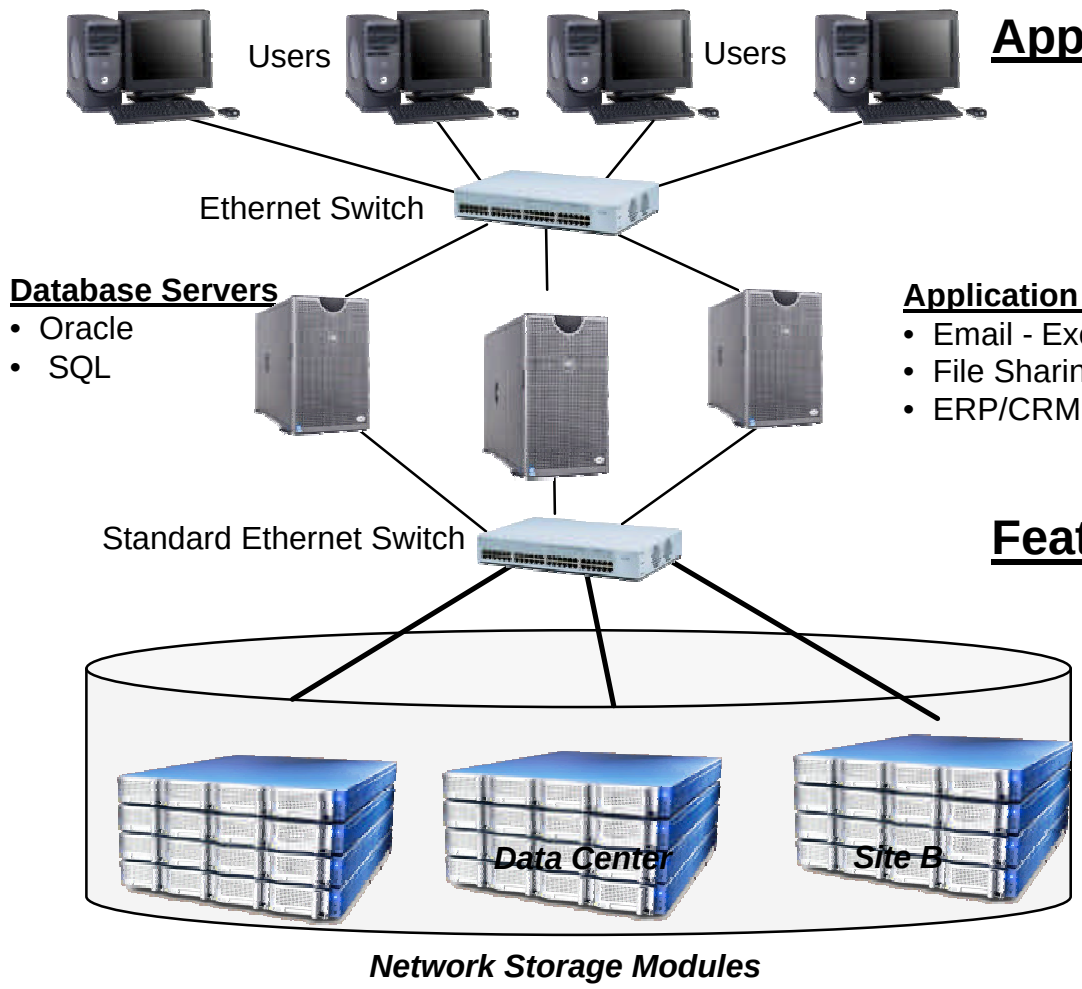

Applications:

- **Server consolidation**
- **Disk-based backup**
- **Business continuity**

Features:

- Distributed Modular Architecture
- Std. Ethernet Network
- Granular Scalability
- Pooling & Virtualization
- Adv. Management Features
  - Snapshot
  - Local/geographic replication
  - Hot sparing – load balancing

LeftHand's Leading Solution

Our technological competitive advantage is delivering the following features all simultaneously within a single solution.

1. Block-level Networked Storage
2. Network-wide Storage Virtualization, locally & geographically
   - Unlimited scalability of storage pools
3. Configurable Replication, 2-way or more
4. Full Support for Shared Volumes
5. High Availability
   - Self Healing, Hot Spare and Automatic Failover & Recovery
   - Coherency maintained with automatic data re-synchronization of down or failed storage modules
   - On-the-Fly Snapshot & Layout Modification

Software Technology Uniqueness

- **First advanced virtualization software technology that is based on clustered/distributed computing technology.**
  - Cluster node auto-discovery protocol
  - Distributed storage clusters using quorum management
  - Advanced cluster group communication protocols based on Ensemble
  - 3-Phase commit protocol for transactional integrity

- **Unique characteristics of LeftHand's distributed system**
  - Self load balancing
  - No cluster management bottleneck - performance scales nearly linearly as storage modules are added to the network
  - Write-all/read-any asynchronous protocols: no STOMITH
  - Support for external lock manager
  - Compatible with Distributed File Systems

Reliable & Scalable Block-Level IP SANs

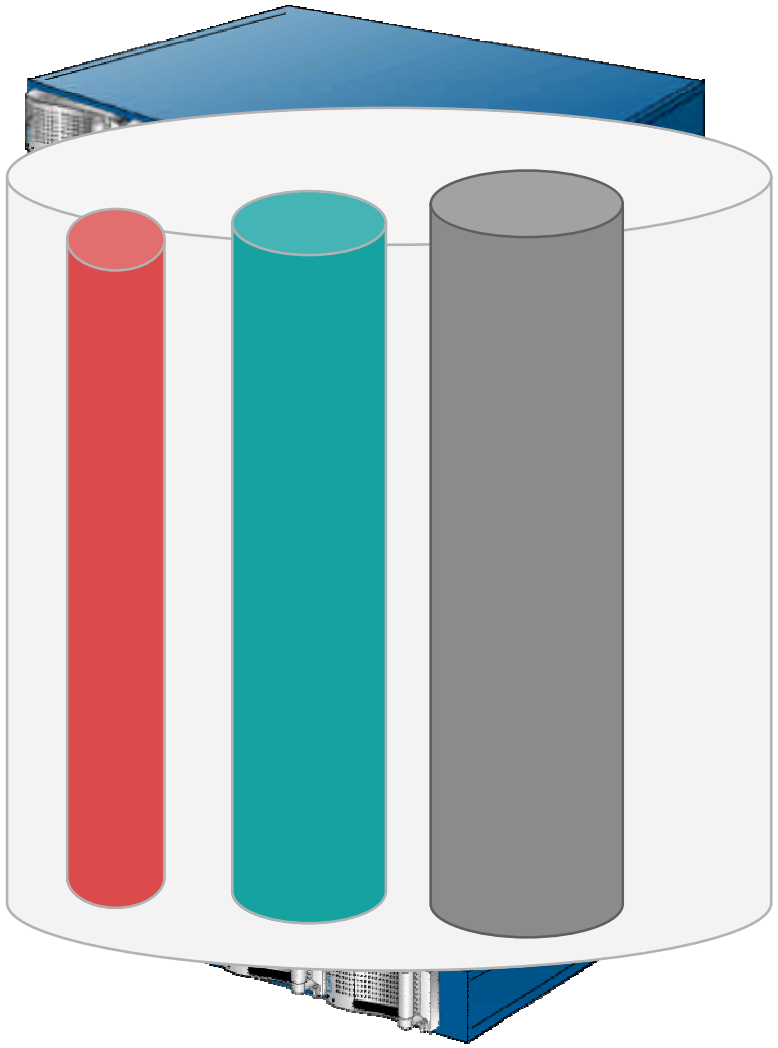

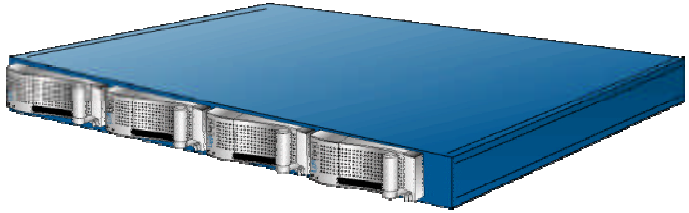

- Patented distributed storage virtualization & clustering SW
- Add storage capacity safely and on-the-fly
- No server downtime
- No single points of failure, not in data path
- Virtual volumes striped/ re-striped automatically across distributed storage pool
- Intel-based storage arrays

LeftHand Network Storage Modules

- Intel-based, distributed network architecture

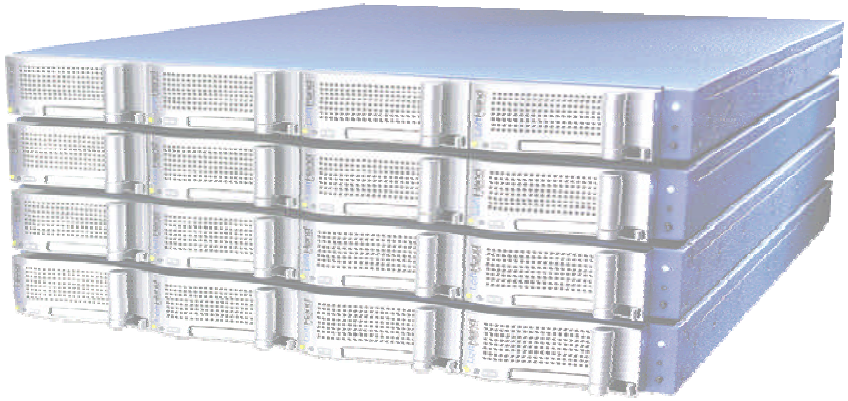

- Flexible 1U/2U form factors:

NSM 100 (500 GBs) / NSM 200 (1.2 TBs)

LeftHand SAN/iQ Console

Drag-and-drop to create virtual storage pool within and across SAN enclosures

Create/Expand LUNS on-the-fly

Drag & drop additional disk capacity

Built-in Security & SNMP alerting

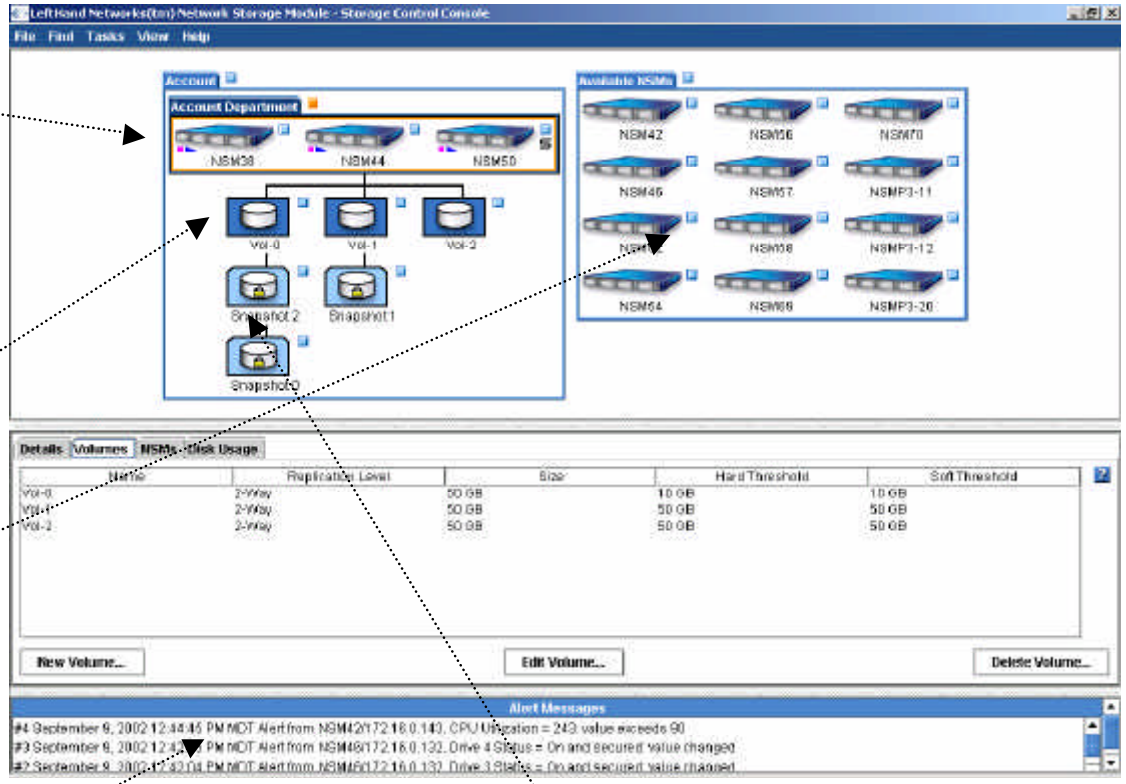

Create snapshots & replicate data across multiple SANs across LANs and WANs

LeftHand SAN/iQ Remote IP Copy

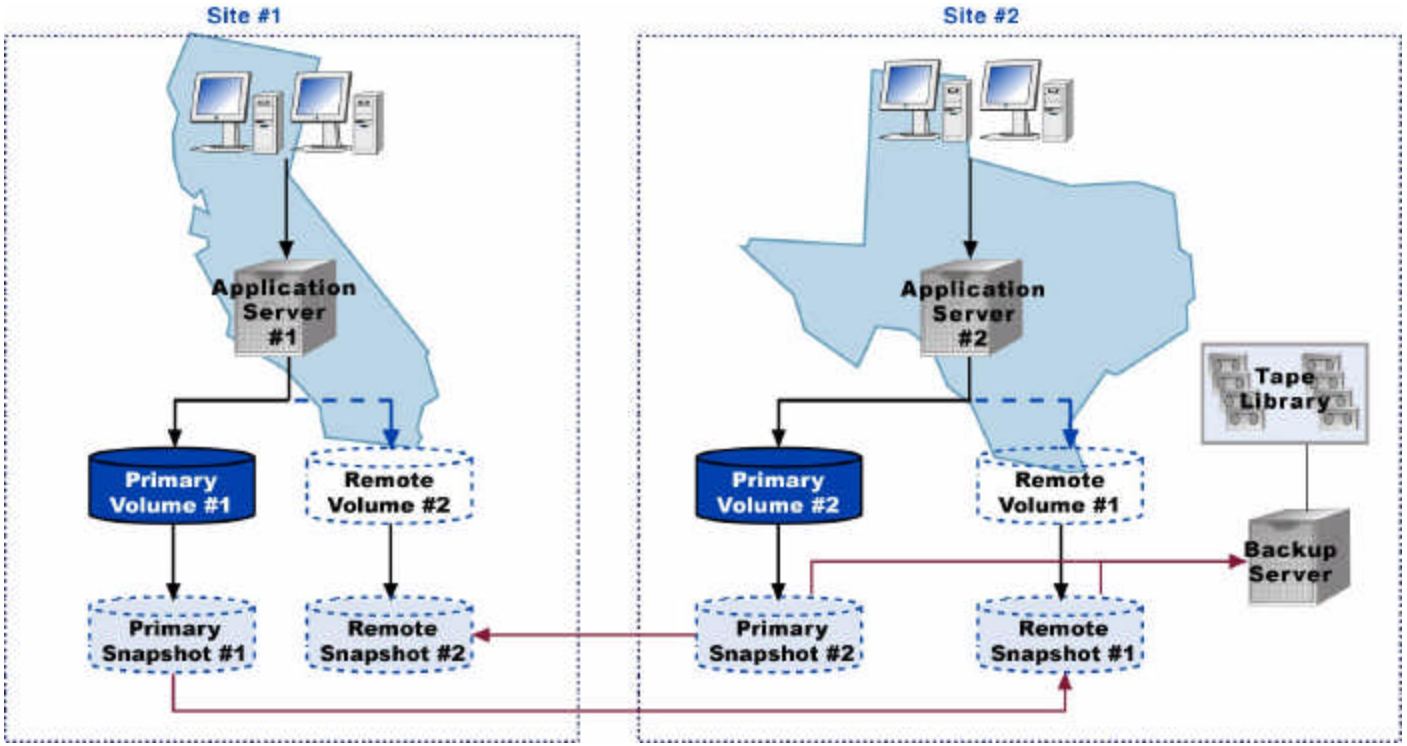

- Cross-replication between sites for local or WAN Data Protection
- Ad hoc or scheduled replication
- Incremental updates minimize bandwidth and telecom costs
- Interrupted updates restart, rewrite previous 10%

Remote IP Copy

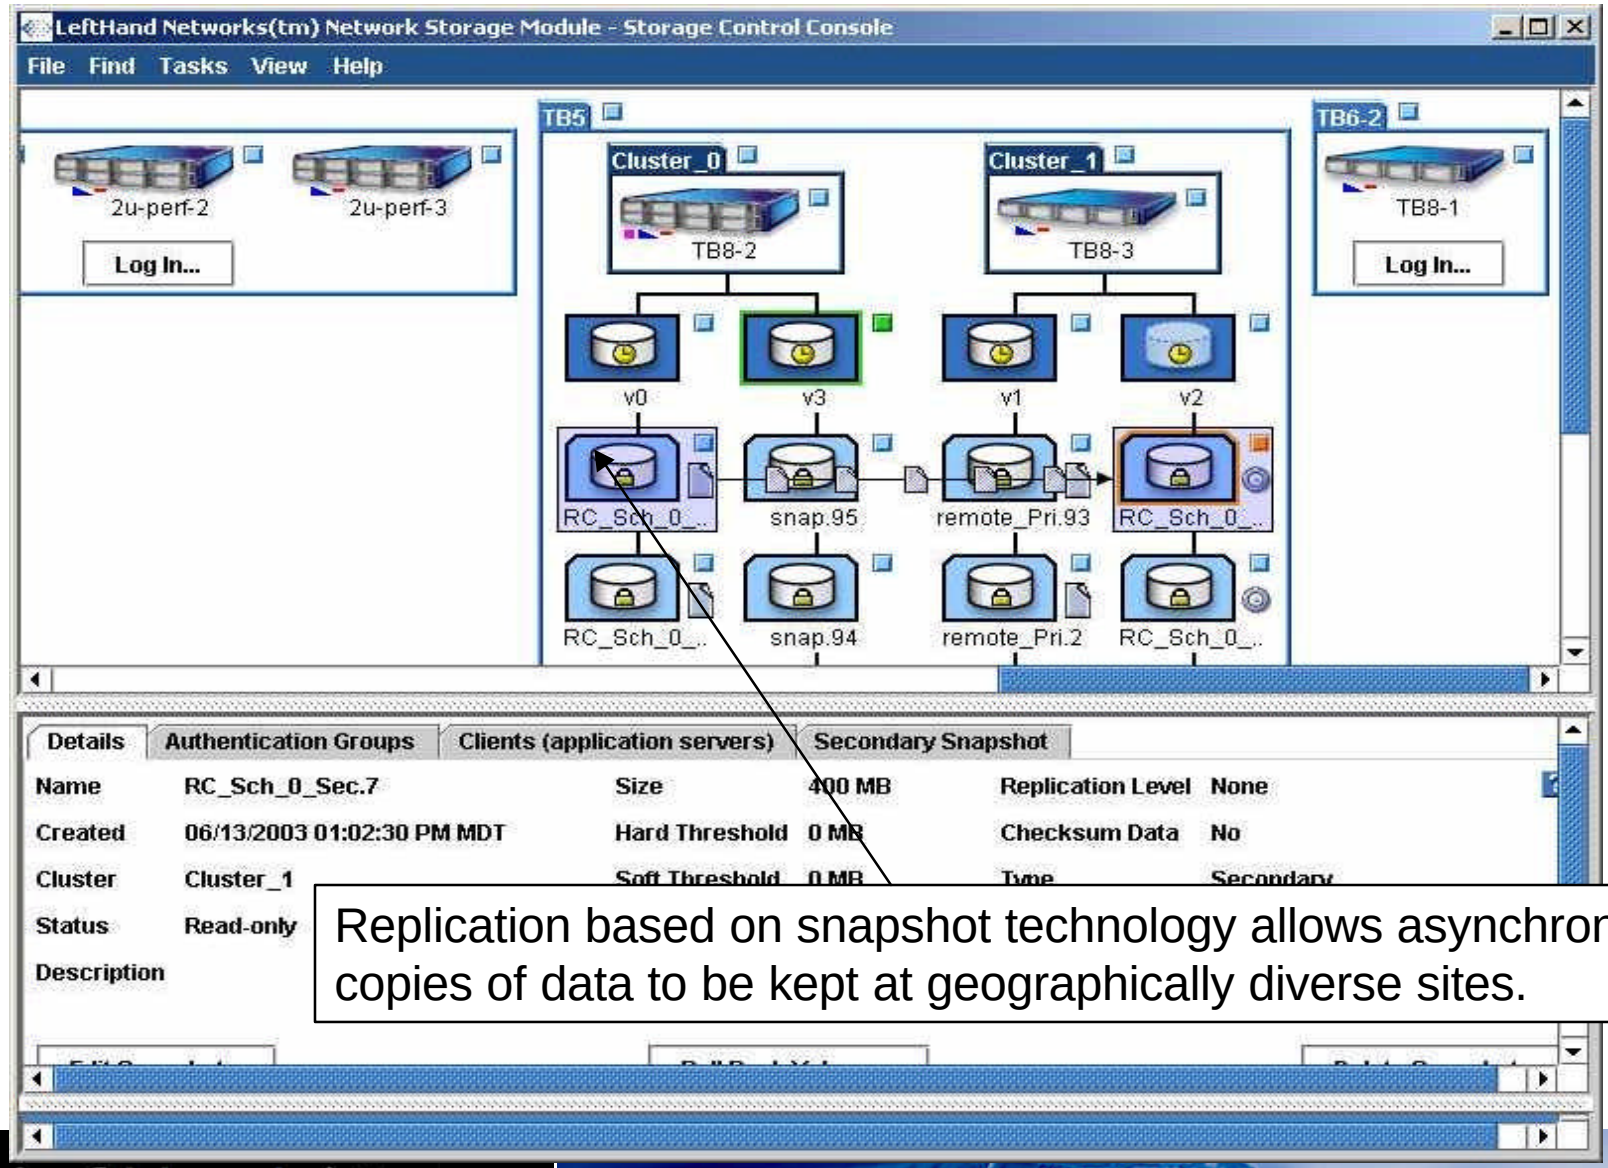

Replication based on snapshot technology allows asynchronous copies of data to be kept at geographically diverse sites.

Provisioning Services

Simplified administration:

Size volume & replication levels

Ability to over provision volume size

Managed by thresholds

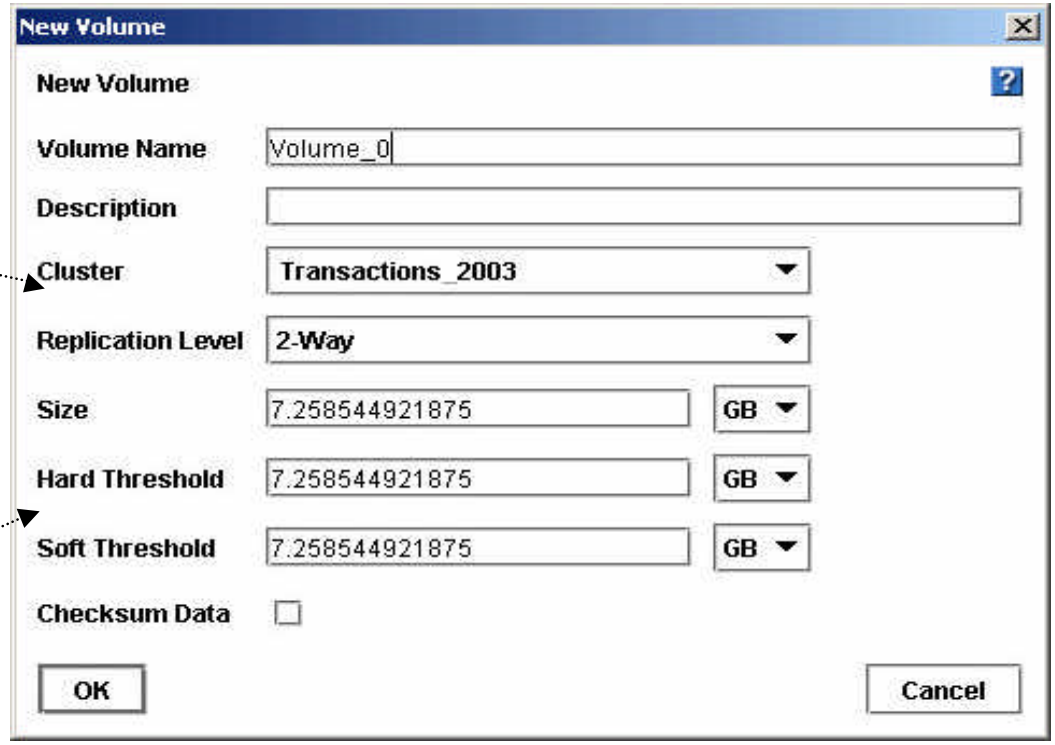

Chain De-clustering

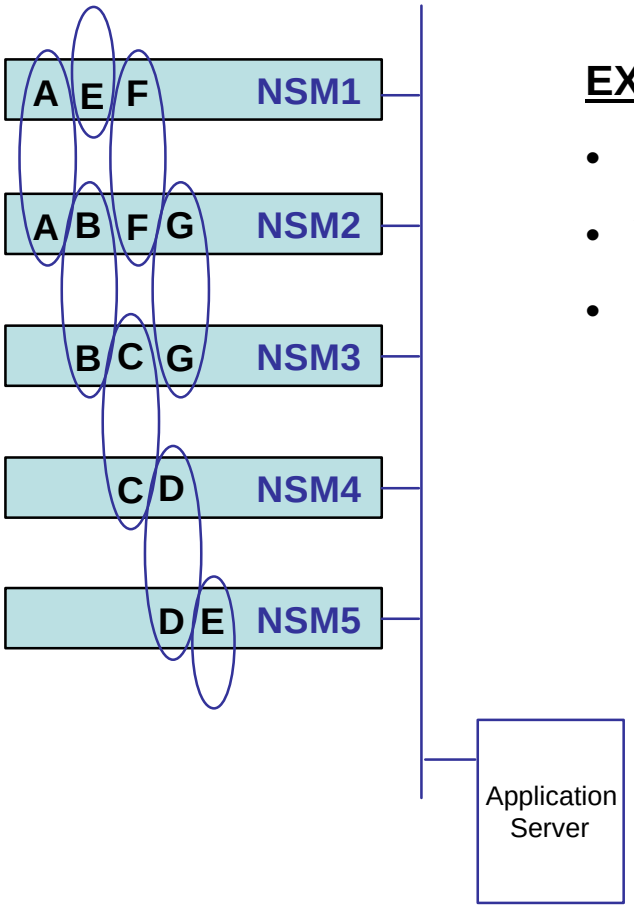

EXAMPLE

- 5 NSM Cluster
- Volume 1 is striped across 5 individual NSMs
- Replication level (chain level) = 2

Chain De-clustering

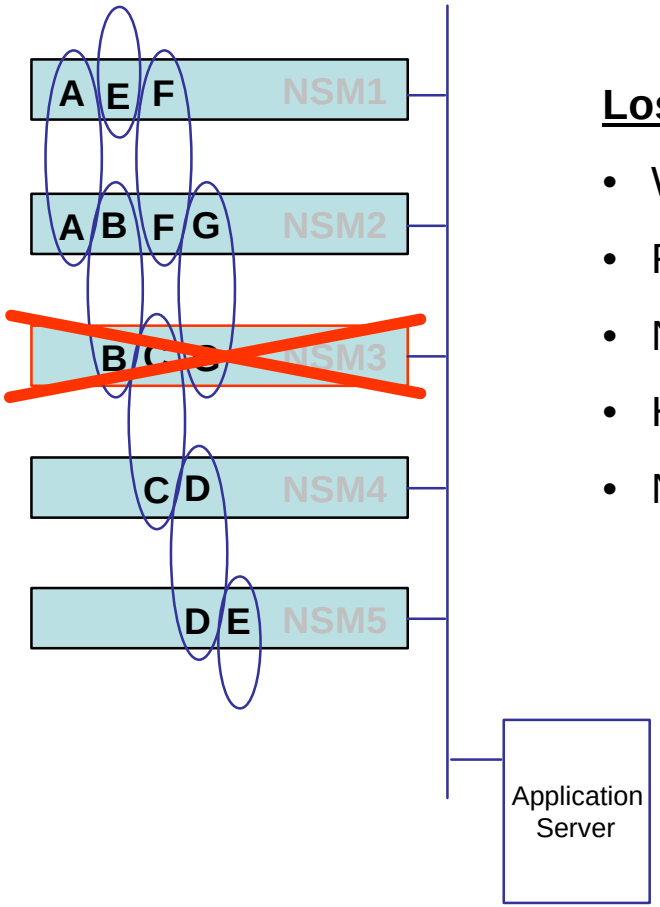

Loss of an NSM within the cluster

- Write access speed remains constant
- Read access speed decreased by only 1/N
- No parity calculation required
- Hot spare NSM is brought on line (if available)
- No interruption of application access to data

Data Protection Services

Programmable Snapshots

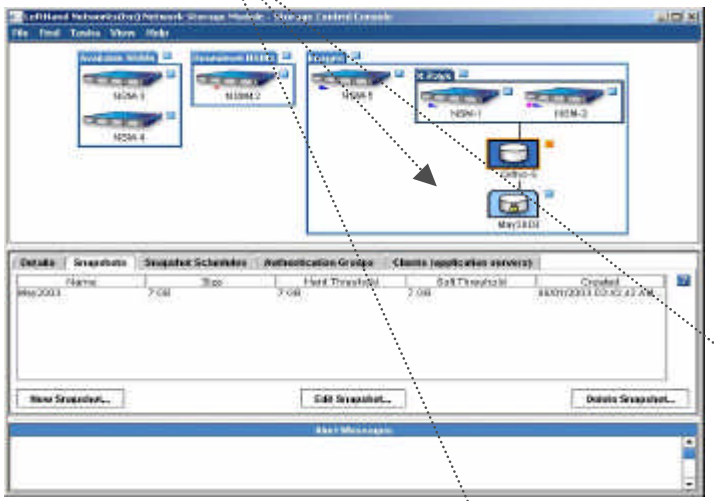

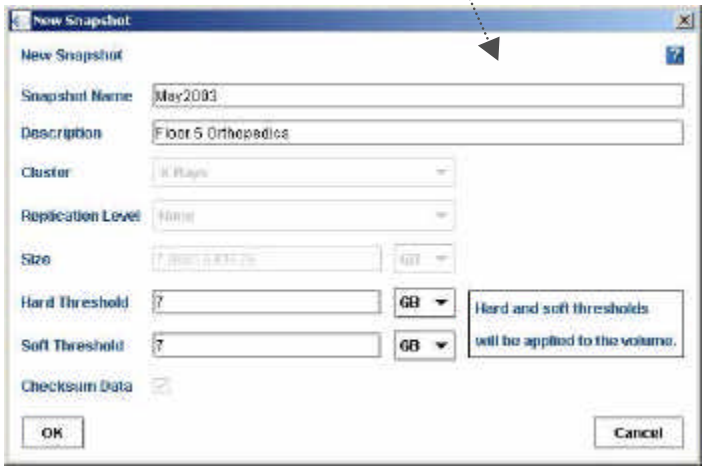

Create a copy of your volume, scheduled or ad hoc, without impacting performance

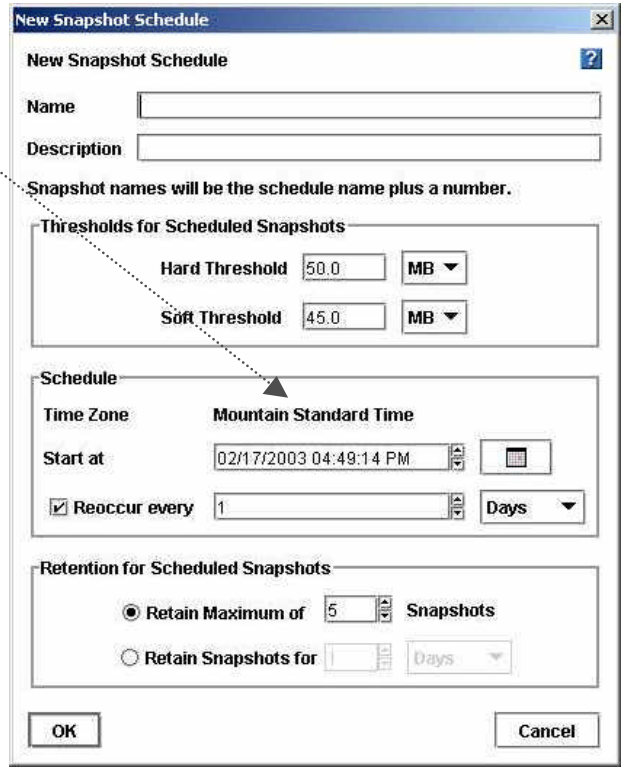

Security Services

Limits access to each volume to specific server IDs

Set to:
No access
Read-only access
Read-write access

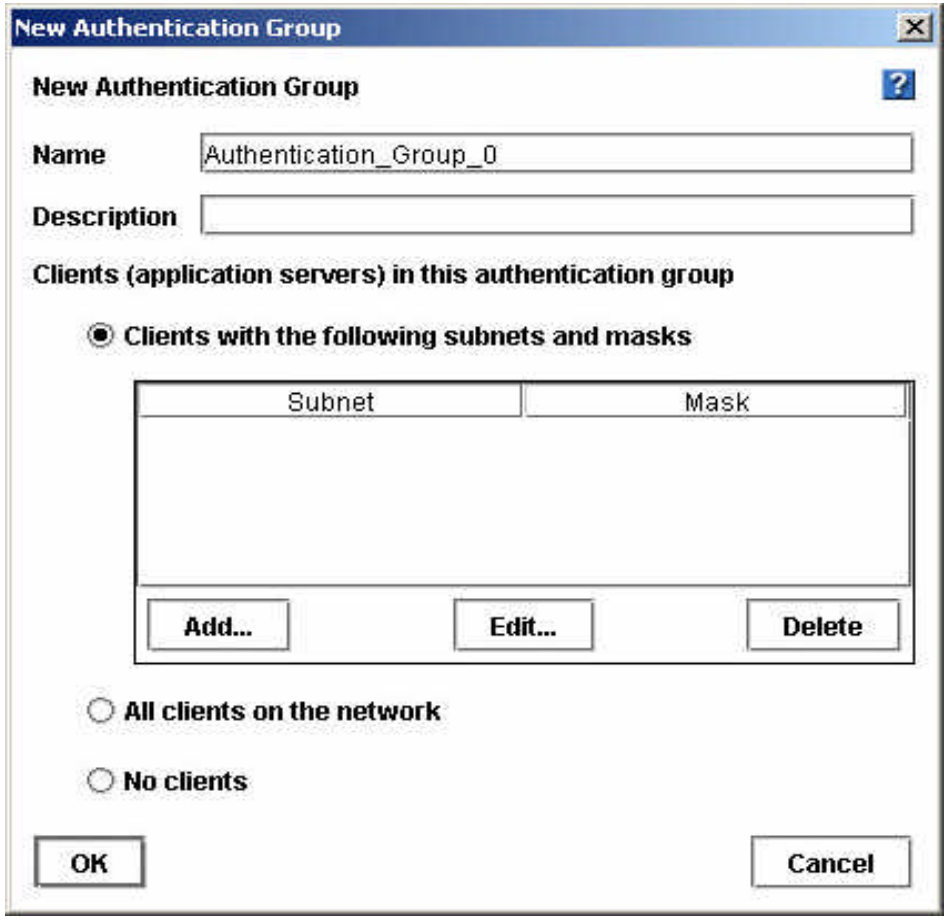

Managing an IP-SAN Disk

```
-- Add Device 0 to system
# aebsvm --add 0
Loading Device 0:
Loading Configuration for Device 0:
Starting Device 0 ( Volume_2:Linux_Quorum ):SUCCESS

-- Block Device Created
# ls -l /dev/aebs/disk0
brwxr-xr-x    1 root     root     177,   0 Nov  7 05:36 /dev/aebs/disk0

-- Make an ext3 filesystem
# mke2fs -j  /dev/aebs/disk0

-- Make mount point and mount volume
# mkdir /mnt/Volume_2
# mount /dev/aebs/disk0 /mnt/Volume_2

-- Display mount point information
# df -h  /mnt/Volume_2
Filesystem           Size  Used Avail Use% Mounted on
/dev/aebs/disk0      441G   33M  418G   1% /mnt/Volume_2
```

Entry Level Cluster Performance

Clients running IOMETER application

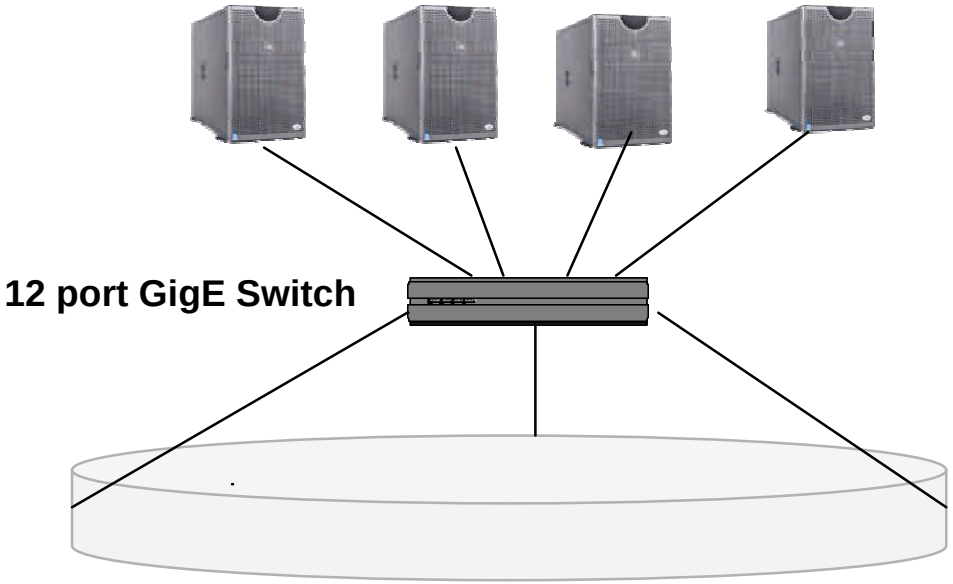

- **Peak write throughput of 155 MB/s**
- **Peak write IOPS of 5000**
- **100% cache hit write IOPS of 16,000 4k blocks**

Company Overview

- **Mission: To be the leader in IP-based networked storage**
  - *Simplifying Full-Featured IP SANs*
- **Founded in 1999, HQ in Boulder, CO**
- **Well capitalized venture-backed $39M to date**
  - Sprout Group, Sequel Partners, Boulder Ventures, Vista Ventures, Portage Ventures, Ironside Ventures, Wasatch Ventures, New World Ventures, Garage Technology Ventures
- **Experienced management team**
  - Dell, HP, Compaq, Maxtor, GE, Lucent, Quantum, StorageTek
- **Sells turnkey Intel-based IP SANs - LeftHand's strategic focus is software**
- **Sell to mid-tier and emerging enterprise companies - Indirect channels distribution model**
  - Storage and Systems VARs and Resellers
- **First full revenue year in 2002, 300% growth in 2003 YTD**
  - 250 installations
- **Partnerships: Microsoft OEM & Certified Partner, Red Hat Ready Partner, Oracle Solution Partner, Veritas, Sun iForce partners**

Storage As It Should Be

LeftHand Networks brings simplicity to complex, full-featured storage networks