iSCSI Unified Network Storage

Steve Rogers
Adaptec, Inc.
691 S Milpitas Blvd, Milpitas, CA 95035-5473
Phone:+1-510-957-1631
E-mail: steve_rogers@adaptec.com

Presented at the THIC Meeting at the Sony Auditorium, 3300 Zanker Rd, San Jose CA 95134-1940

April 19-20,2005

The Direct Attached Storage Legacy

- Proliferation of Servers on the Network
- Maintenance of Internal Storage
- Inflexible Expansion
- Aging Server & Storage Obsolescence
- Downtime
- Backup and Recovery
- Skill Set of IT Staff
- Dealing with Unplanned Growth
- Poorly Utilized Disk Space

Data Everywhere — Accumulation of Least Used Data on the Most Used Storage Platform

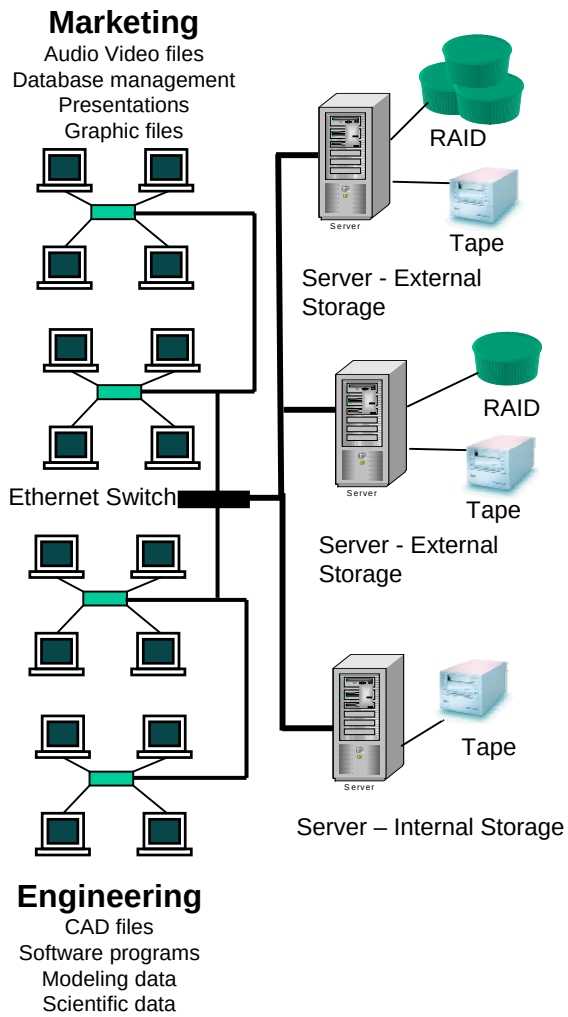

Transition from DAS to Networked Storage

- Storage Independence
- Easier Management
  - Easier Upgrade & Expansion
  - Centralized Management
- Storage Consolidation
- Centralized Backup

iSCSI Has the benefits of Fibre Channel SAN, over Your Existing Ethernet Infrastructure!

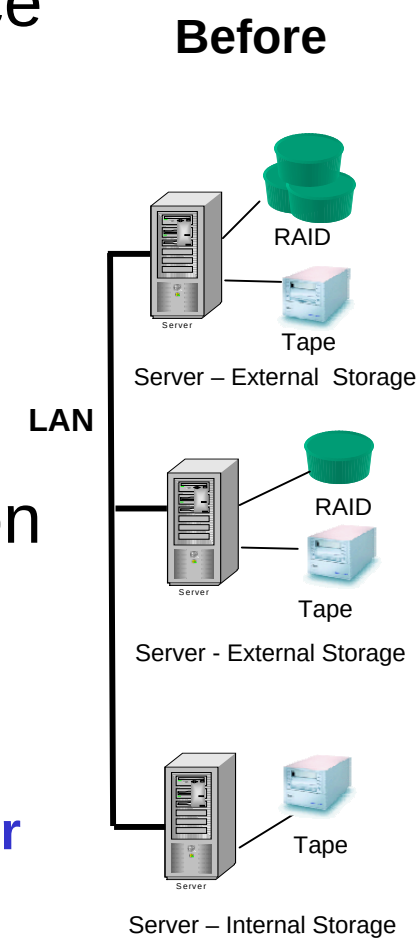

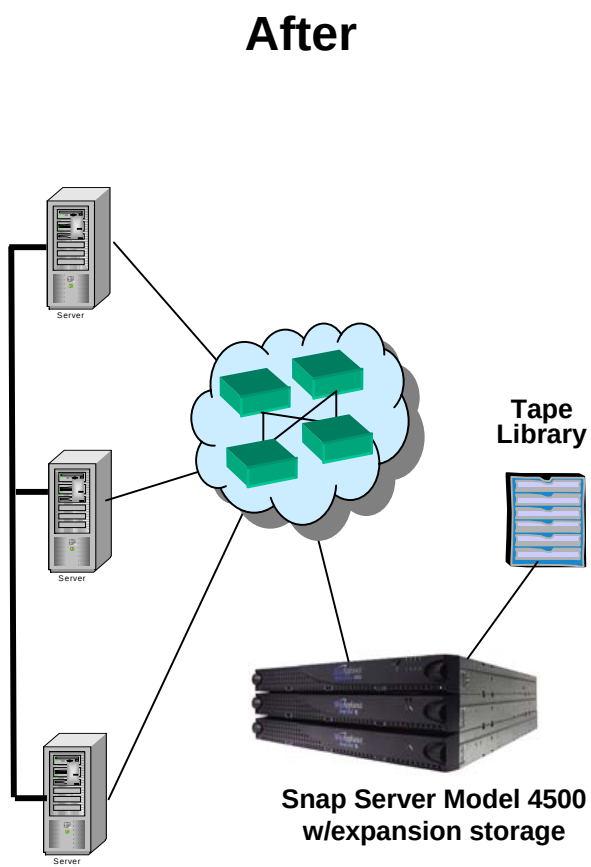

What is iSCSI & How does it work?

- **Leverages TCP/IP Protocol for Transmitting SCSI Data (Block) over an Ethernet Network**
- **TCP/IP is known for Reliable Data Transmission over Potentially Unreliable Networks**
- **iSCSI is a Client/Server Model: Initiator Makes Requests, Target Responds**
- **Each Node has a Unique ID Consisting of the IP Address, the TCP Port Number, and Either the IQN (iSCSI Qualified Name) or EUI (WWN) Name**

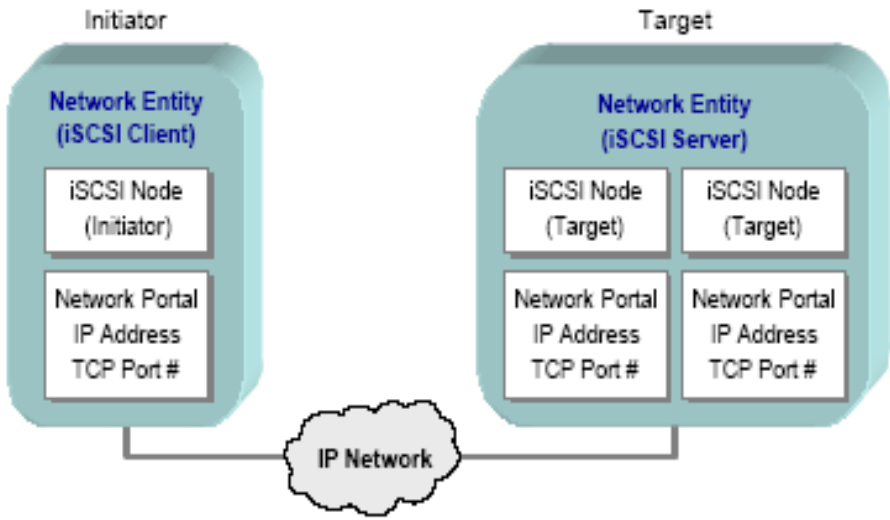

iSCSI Data Transmission

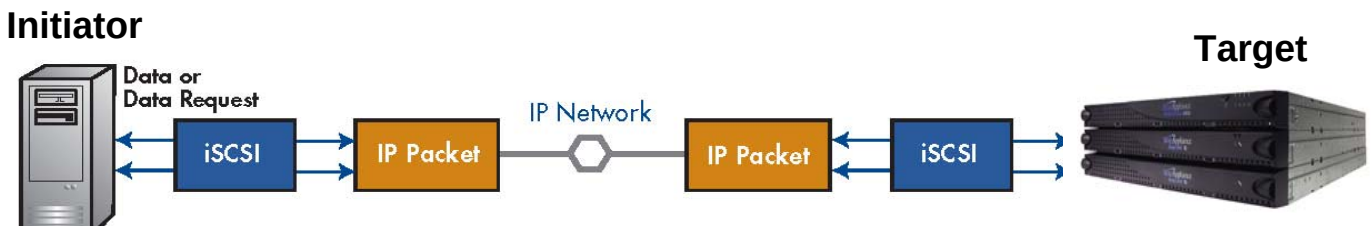

App. Server with ISCSI Initiator driver or HBA

iSCSI Target using Snap Server Model 4500 w/expansion storage

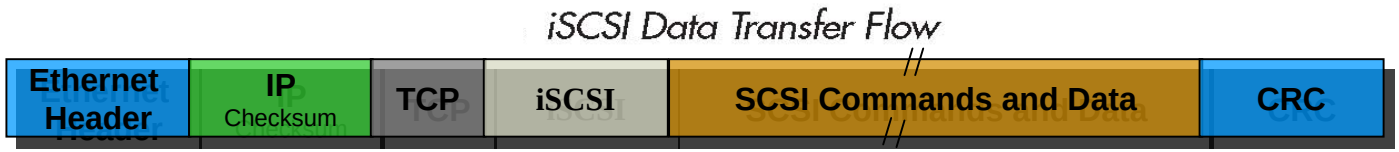

Ethernet Frame with iSCSI Data

iSCSI Disk looks like Local Disk to the Application

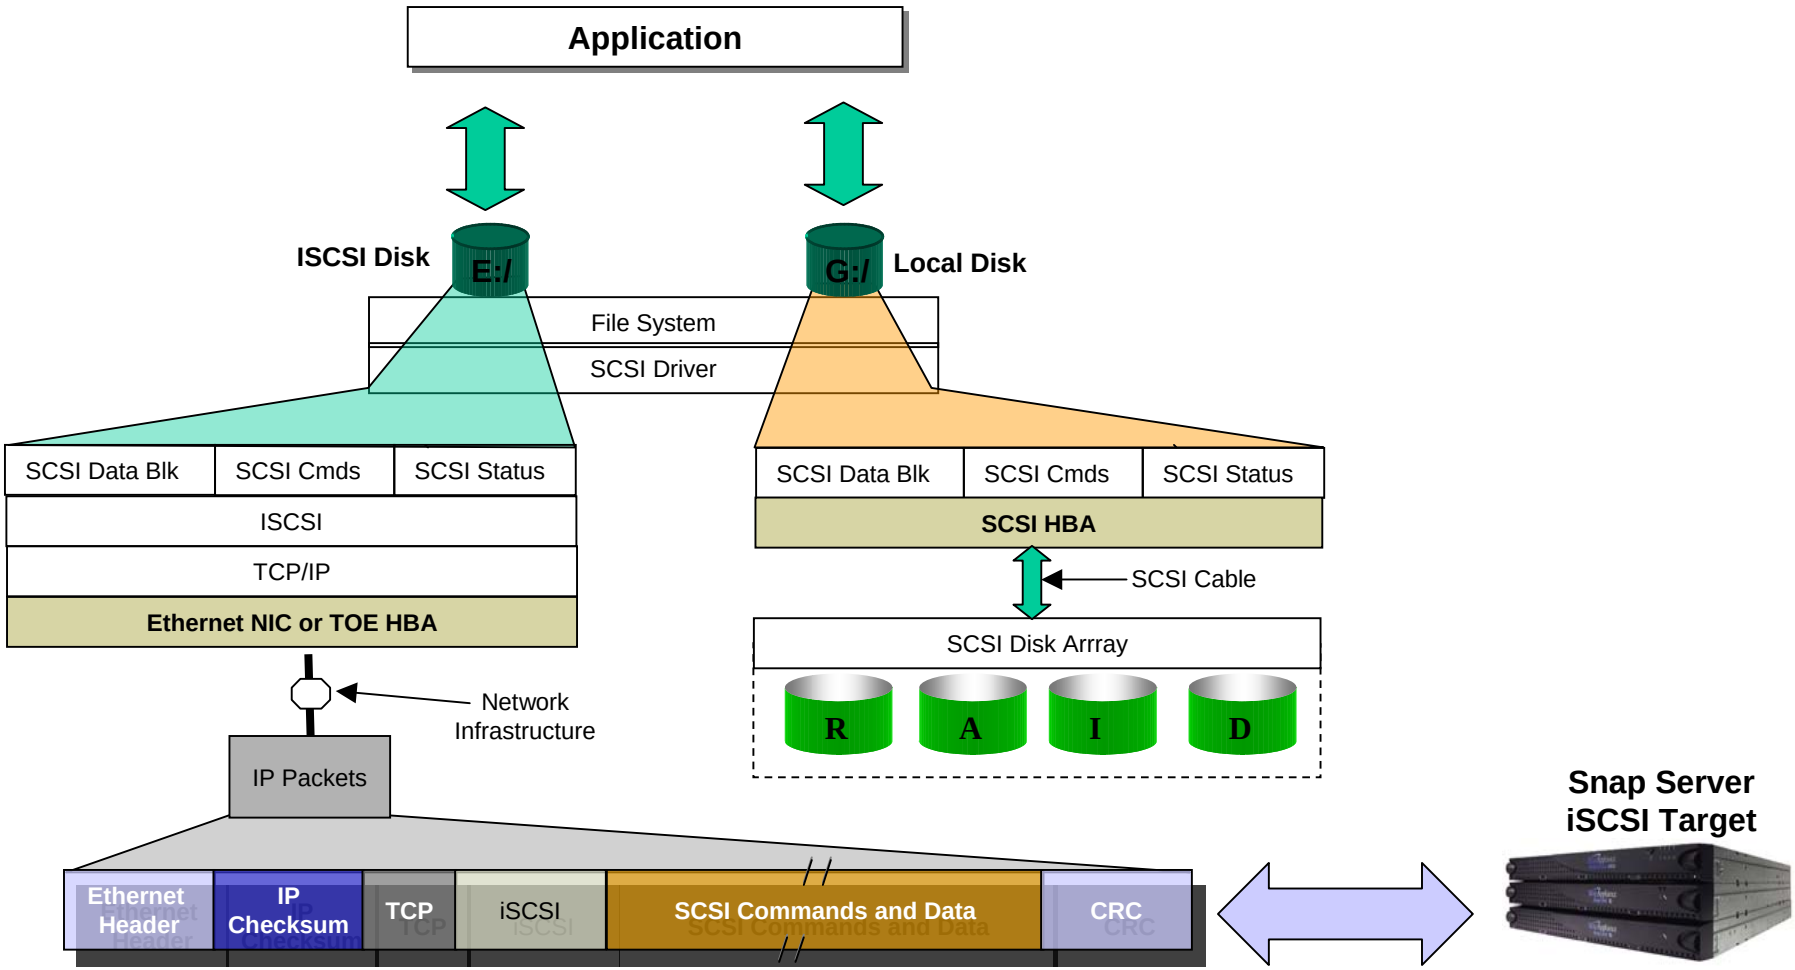

Topology

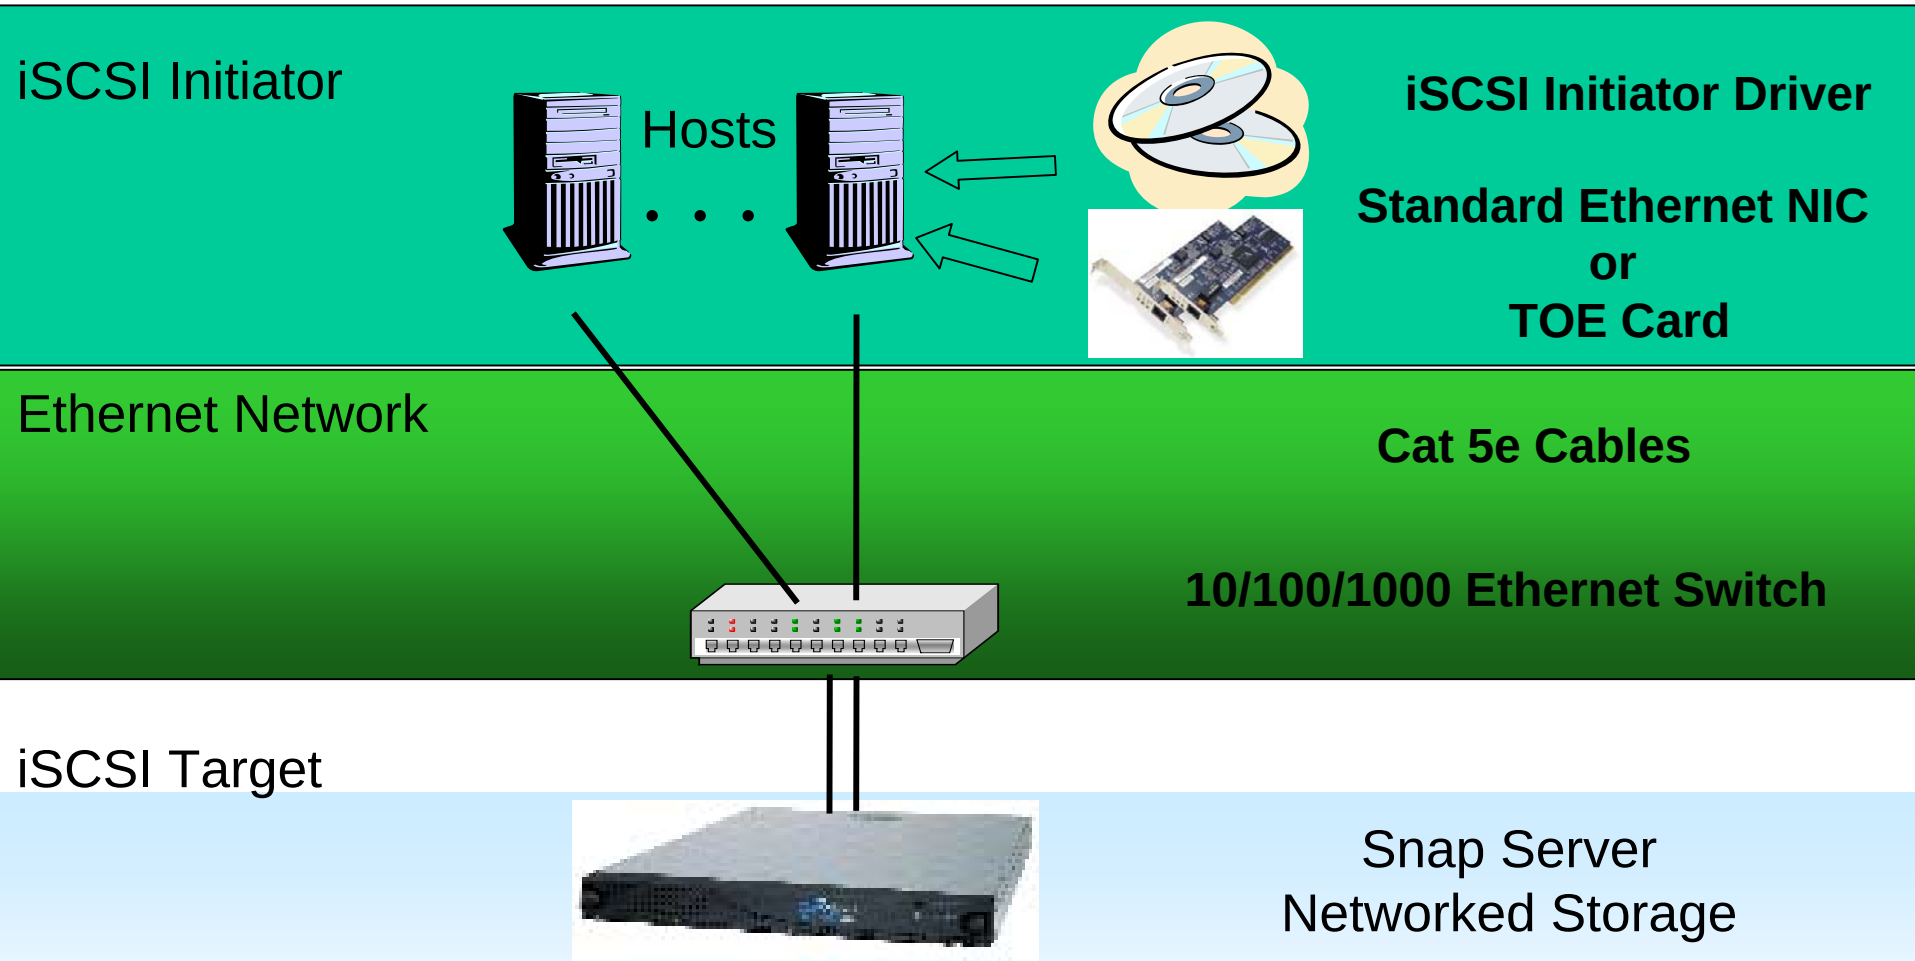
iSCSI Initiator
Hosts
iSCSI Initiator Driver
Standard Ethernet NIC
or
TOE Card
Ethernet Network
Cat 5e Cables
10/100/1000 Ethernet Switch
iSCSI Target
Snap Server
Networked Storage

Initiators – Hardware vs. Software

- **Hardware Initiators:**
  - Prices coming down
  - OS Agnostic
  - Offloads CPU from Processing Network Stack — Some Claim as Much as 5x More CPU Cycles Returned to Server
  - Some Support Dynamic Disks like: Adaptec 7211C
  - Some Have Capability of Booting from an iSCSI Disk
- **Software Initiators:**
  - Free, most downloadable from web.
  - Uses CPU Cycles to Handle Network/iSCSI Stack — Generally Acceptable, Depending on Application Demands

Discovery and Access Security Between Hosts and Targets

- **Simple Access Control**
  - Specify the Target Node Name, IP Address, and Port to the Initiator and Target Devices
- **Advance Access Control - iSNS - Naming Services for Authorized Target Discovery**
  - Authenticates iSCSI Storage Nodes During Logon
  - Store and Distribute X.509 Public-key Certificates
  - Create Discovery Domains

  Similar to FC Zoning

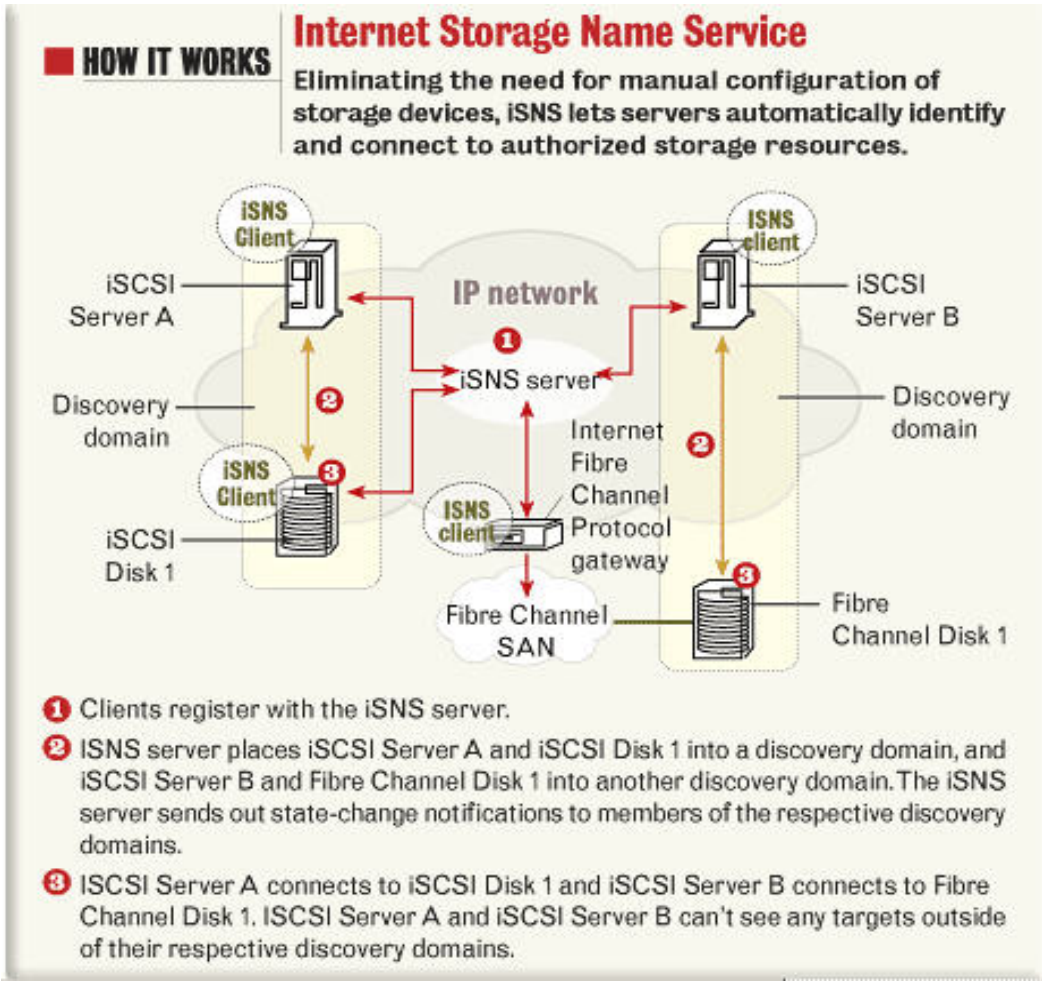

Source: Network World Fusion – iSNS eases management of storage nets 9/01/03

CHAP- An Additional Measure of Security

CHAP (Challenge-Handshake Authentication Protocol) Secure Method for Connecting an iSCSI Initiator to a Target

- At iSCSI Disk Creation, Administrator selects CHAP and Enters a Unique Authentication String
- When the Host/Initiator Requests a Connection to the Target, the iSCSI Target Sends Back a Challenge Message. The Initiator Responds with a Value
- The iSCSI Target Verifies Response by Comparing its own Calculation of the Expected Hash Value
- If There's a Match, Authentication is Acknowledged, and connection proceeds.

iSCSI Addresses SCSI and SAN Constraints

IP SAN	Customer Benefit
Leverage Ethernet Technology	• **Lower Storage Expenditures — Complicated Fibre Channel Fabric not Always Necessary** • **Leverages Existing IT Infrastructure, IT Skills, and Resources** • **IP - Established Standard — Improved Interoperability, Reduces Installation Costs**
Lower TCO	• **Affordable SANs for Small to Medium Business Apps** • **Addresses Today's Budget Constraints and Cost Cutting Programs**
Expand Distance and Capacity	• **Ethernet is Inherently LAN/WAN Capable** • **Capacity can be Added with Minimal Costs** • **Addresses Remote Office and Distributed Storage Needs**

iSCSI Connectivity in Review

- iSCSI SW Driver or HBA is Installed on the Host
- iSCSI Targets Look Like Local Disk to the Host's OS
- Applications Don't Require any Special Awareness
- TCP/IP Facilitates LAN or WAN Connectivity and is Responsible for Packet Data Integrity
- iSCSI Target LUNS are Created on the Target Device
- Each Host Initiator Logs into the iSCSI Target and Attaches to a Particular iSCSI LUN — Host Login can be Done Manually or via an iSNS Server
- Two key implementations for iSCSI:
  - Shared Nothing Implementation
  - or Clustered Implementation

Recap of Evolving Network Storage Architectures

Pros and Cons

Fibre Channel SAN

- ⊕⊕ **Current DAS Replacement**
- ⊕⊕ Established; Proven Performance
- ⊖⊖ Costly: Needs Special Infrastructure and Resources to Manage
- ⊖⊖ Complex, esp. for Remote Storage (FCIP gateway)

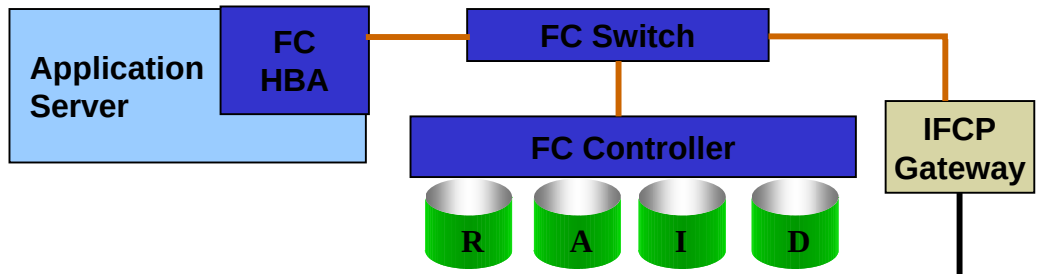

IP SAN

- ⊕⊕ **Looks like DAS to Server**
- ⊕ Lower Infrastructure and Management Costs
- ⊕ Fully Applicable for Remote Storage
- ? Throughput Less than FC – Improves with TOE Acceleration

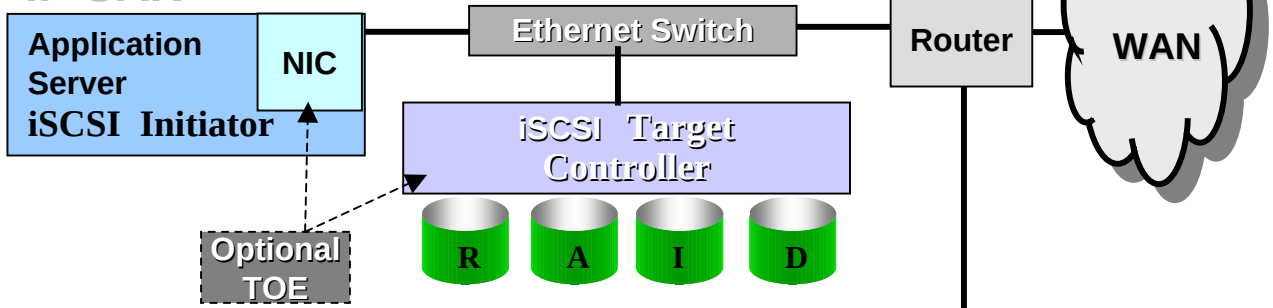

NAS

- ⊕⊕ Simple-to-manage, Inexpensive Off-the-shelf Storage
- ⊕ **"Native" Platform File Sharing**
- ⊕ Fully Applicable for Remote Storage
- ⊖ Single Name Space per device

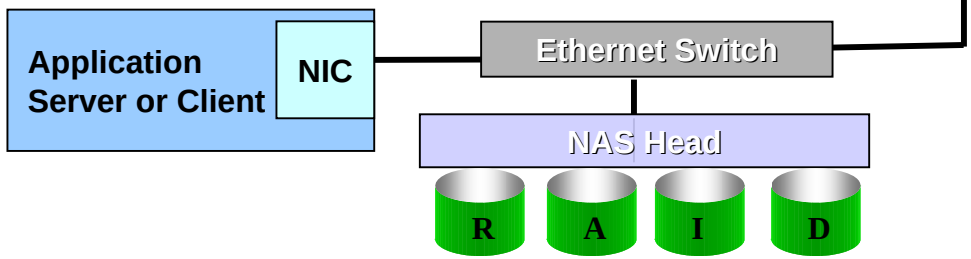

iSCSI / DAS / FC Comparisons

Area	IP-SAN	SCSI (DAS)	FC SAN
Speed	10/100/1000Gigabit FD 10Gigbit (future)	Ultra 160/320 1/2 Duplex	1GB/2GB FD
Discovery Security	Manual logon iSNS CHAP	SCSI Command set None	WW Names Zoning/LUN Masking
Protocol Interconnect	Ethernet Copper / Fibre	SCSI Copper	Proprietary FC Copper / Fibre
infrastructure	Leverages Standard Ethernet Switches/Router/hubs	Internal SCSI Controller and adaptors *SCSI to FC bridges available*	Proprietary FC Switches/Router/Hubs
Plug and Play compatibility	Maturity of Ethernet makes easy compatibility	Mature, but device specific	Requires Stringent Compatibility Testing
Distance	Full WAN Capable with existing infrastructure	Very short - 25m max,	Requires LW Fibre Repeaters and FCIP Gateway
Device limitation	Connections/Target BW Dependent No limit on LUNS/Target	16 Devices Max 32 LUNS/Device	128 Devices/Node loop 16M Devices on Fabric
TCO 1st Year costs Equip Acquisition Costs Cost/MB Management Costs Based on 2TB Storage*	 $15,000 <$.01 $33,000	 $30,000 $.02 $95,000	 $180,000 $.07 $45,000

**Fibre channel costs include, fabric switches, and branded highly available disk arrays*

iSCSI Application Considerations

- **Any Application Using DAS can Utilize an iSCSI Connection to Deliver the Equivalent Raw Disk over Ethernet**
- **End-to-end Connectivity is Dependent on the iSCSI Initiator Driver Target Compatibility**
- **Not all Initiators Work on all Targets — Some Vendors are Making their Initiators Specific to their Targets**
- **iSCSI initiators can only talk to iSCSI Target Devices**
- **Performance Requirements May Make the use of a TOE Card Mandatory**
- **Best Performance Using Gigabit Ethernet**

Snap Server Unified Block and File Architecture

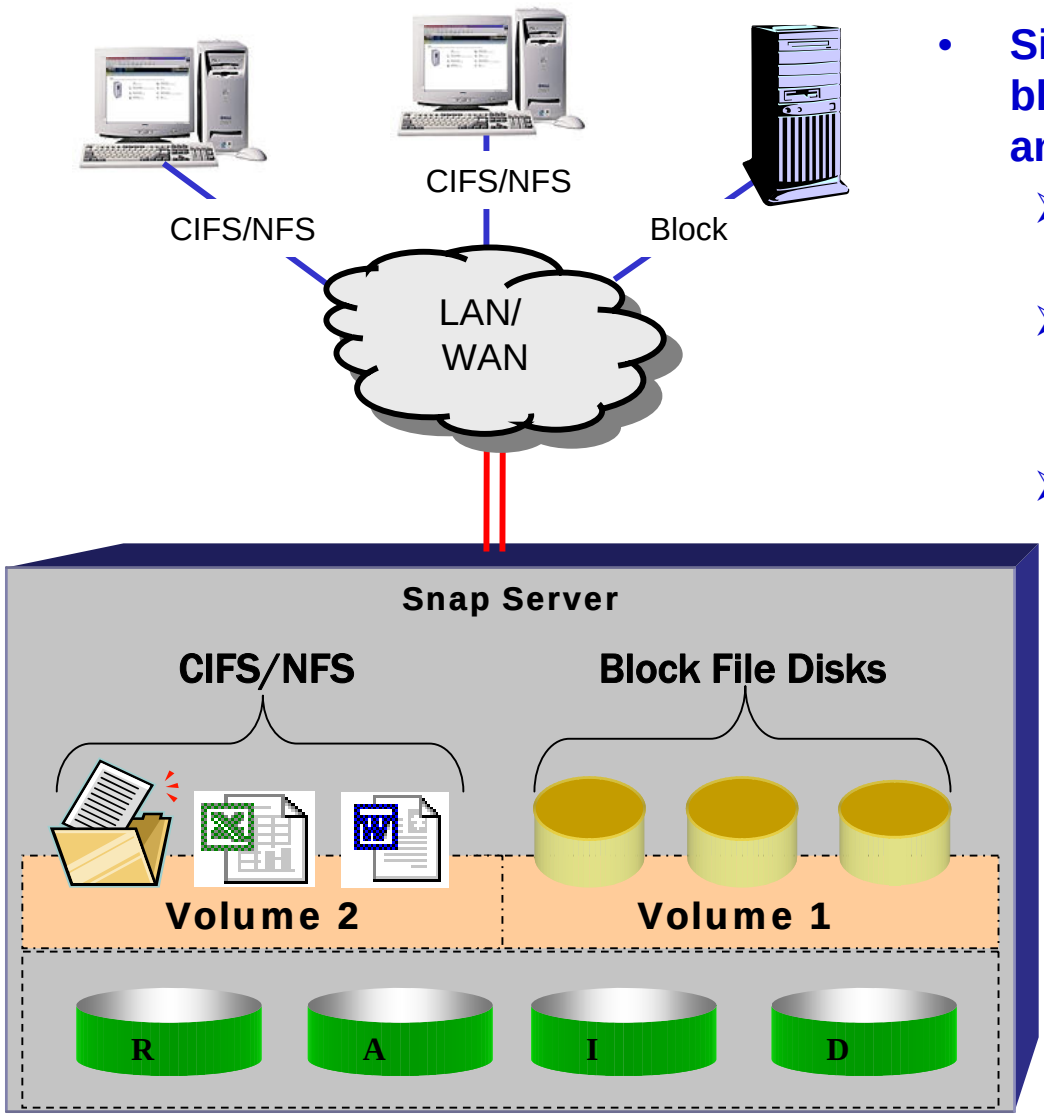

- **Simultaneous file-data (CIFS/NFS) and block-data (iSCSI) to be transported over an IP network. Benefits:**
  - **Leverages existing Ethernet infrastructure and expertise**
  - **Streamlines management by consolidating file and block data on a single device**
  - **Ideal for remote edge-server applications**

Key iSCSI Applications

- Microsoft VSS Shadow Copy Archives
- Microsoft Exchange Archives
- Accelerated Backup/Restore
- Storage Consolidation
- Nearline Archiving Applications
- Primary Storage for small to moderate Database applications
- Mirrored Storage Pairs for Storage Resiliency
- Common Storage Pool for Clustered Servers

VSS Shadow Copy Archive

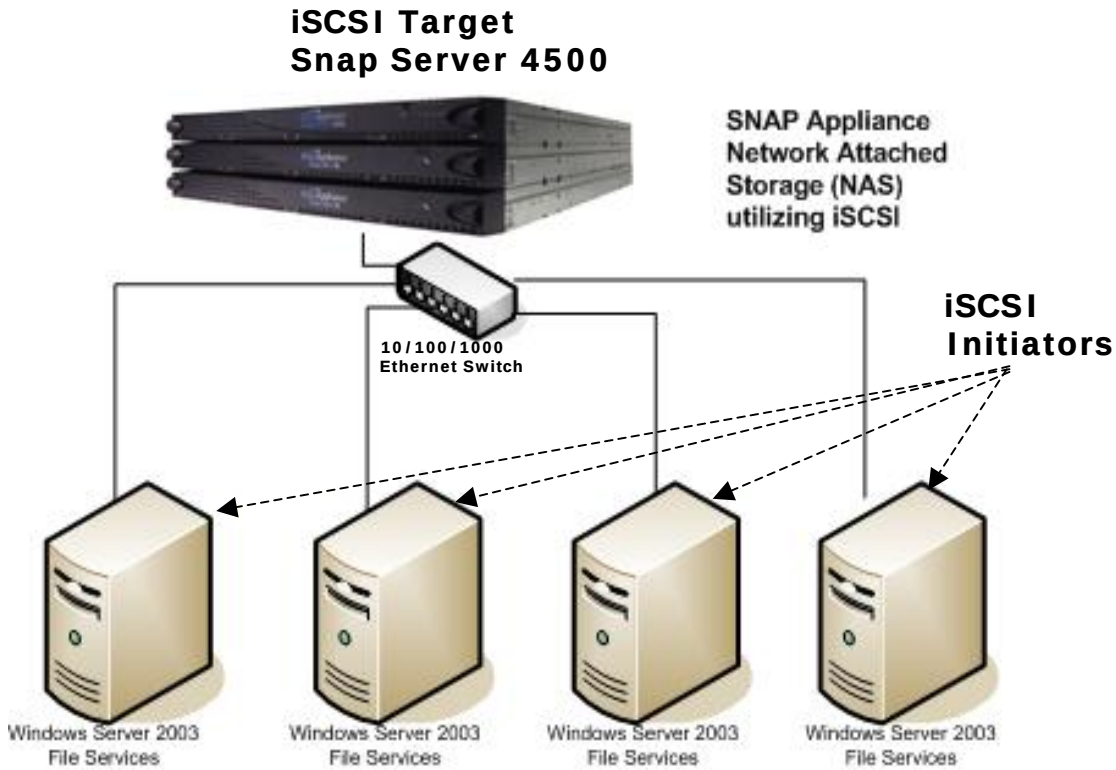

Key Benefits:

- **Saves Recovery Time**
- **User-Accessible Restores**
- **Fast Folder/File Recovery**
- **Easy Access to Previous Versions**
- **Reduces Demand on IT Administrators**

Exchange Message Store Recovery Archive

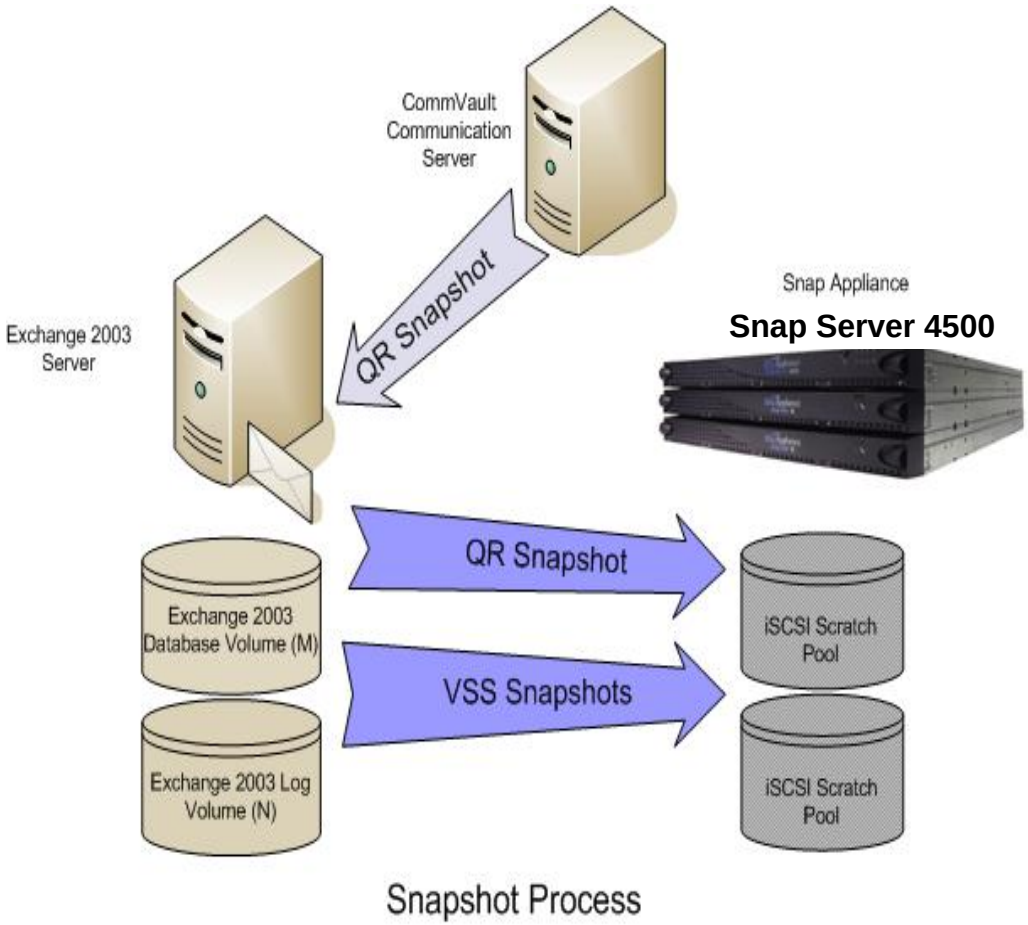

Key Benefits

- Exchange Message Store Recovery in Minutes
- Avoid Necessity of a Clustered Solution.
- Avoid Proprietary Solutions
- Reduced Blackout Time
- Economical Block Storage
- Build Geographic Recovery Points

Using iSCSI with Snap Servers

Application Server Scalability

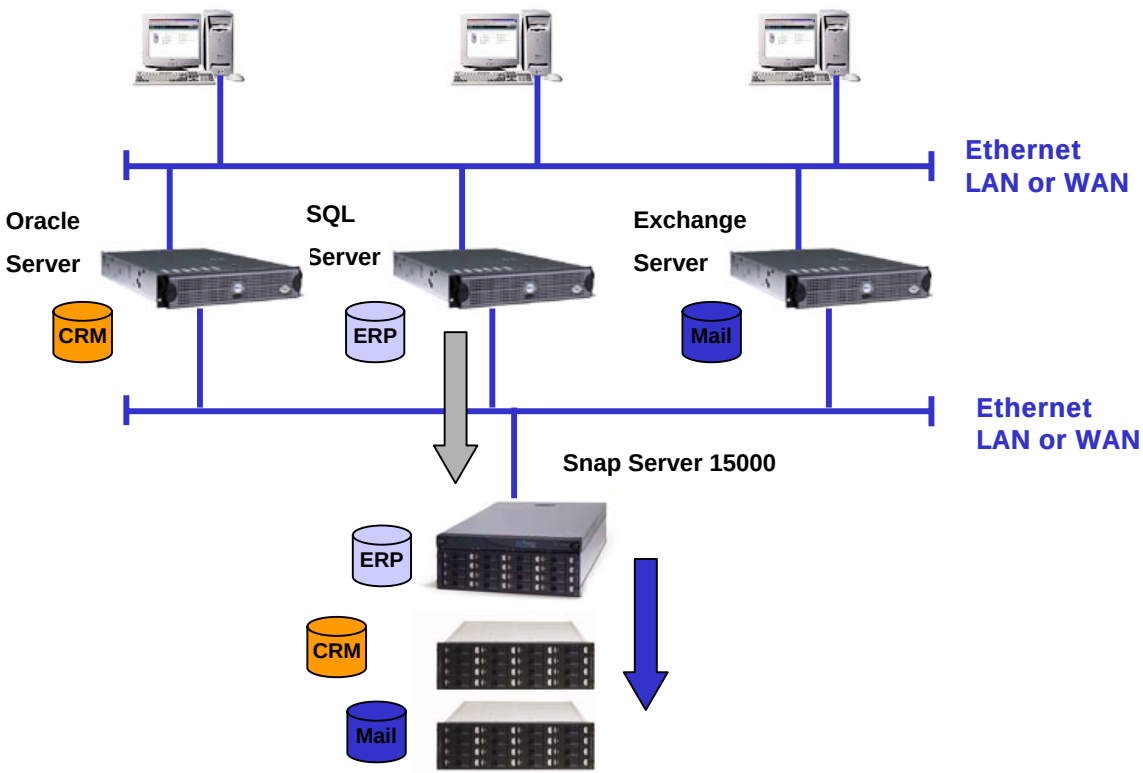

Value Proposition

Migrate Databases to Snap Server and Grow them Without Adding More Local Disk to the Application Server

Easily Scale Storage to Meet Growing Business Needs

Target Application

- Wide Market Applicability Where a Common Storage Pool is Needed
- Exchange Messaging Solutions
- SharePoint Portals
- GreatPlains (Financial)
- CRM Applications
- Databases
- Telco Nearline message Storage Applicability

IP-SAN (Shared Application Store)

Using iSCSI For Block Data and CIFS/NFS For File Data

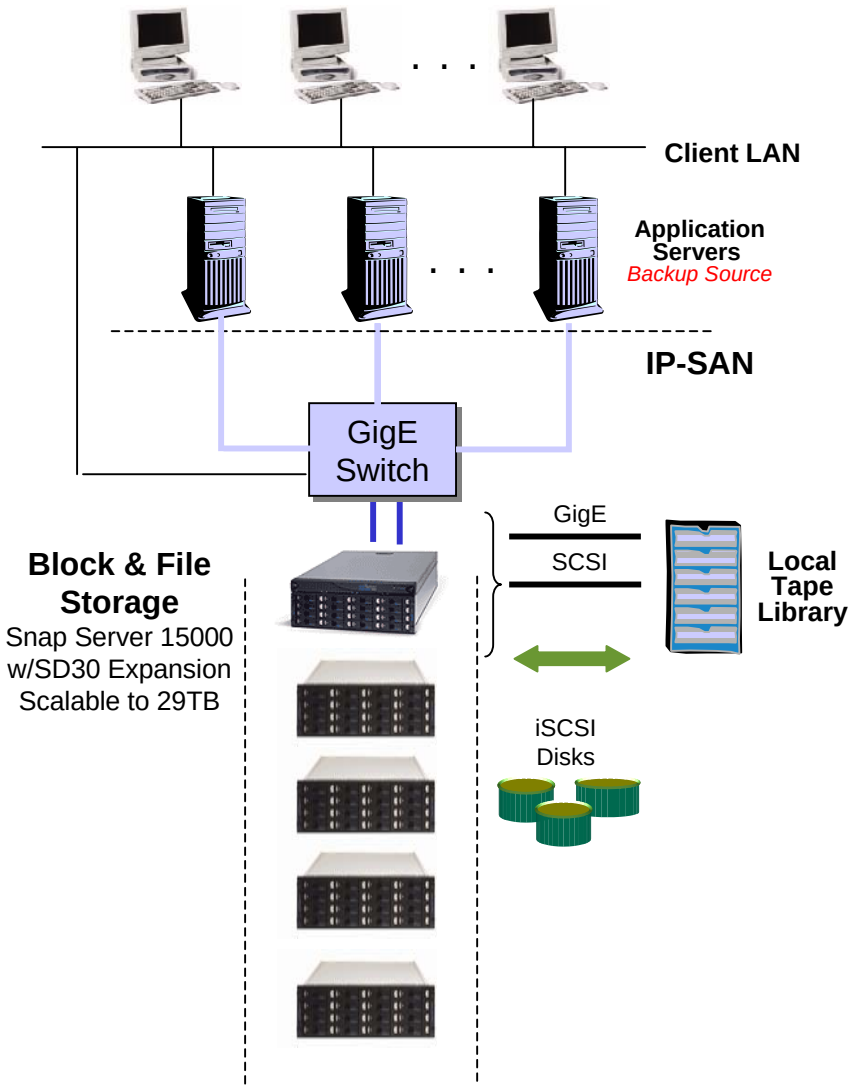

Value Proposition

Scalable, Shared Storage Pool for a Variety of Applications

Easy to Manage, More Affordable than SAN/DAS

Utilizes Inexpensive Ethernet Infrastructure

Target Application

- Wide Market Applicability Where Common Storage is Needed
- Common Block and File Storage Pool
- File Sharing
- Storage Consolidation
- Work Collaboration

Accelerated Backup

Storage Pool for Backup Server

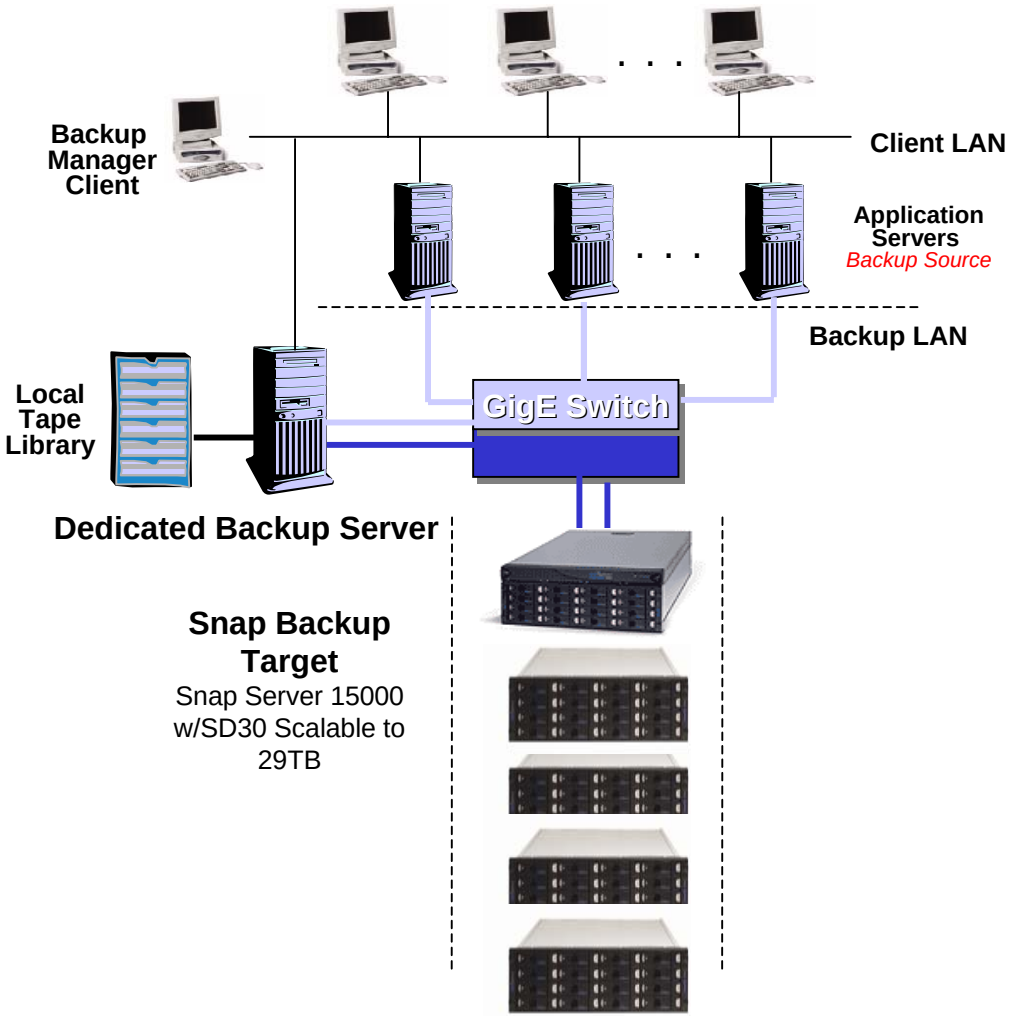

Value Proposition

Accelerated Backup/Restore

Tape Archiving Removed from Critical Backup Window

Comparable to DAS Costs yet Easier to Manage

Integrates into Incumbent Enterprise Backup Software

Target Application

- All Vertical Markets Need Backup and Recovery

Distributed Back Office Storage Consolidation

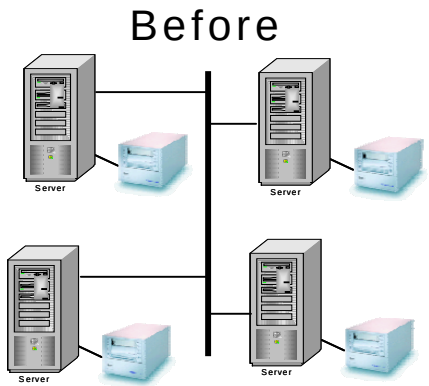

Value Proposition

Each server may have aging disks that can't expand.

Cost-effective Solution for Environments with light to moderate performance Requirements

Native BakBone Self-contained Backup

Extends the Life of Servers with Limited, Storage

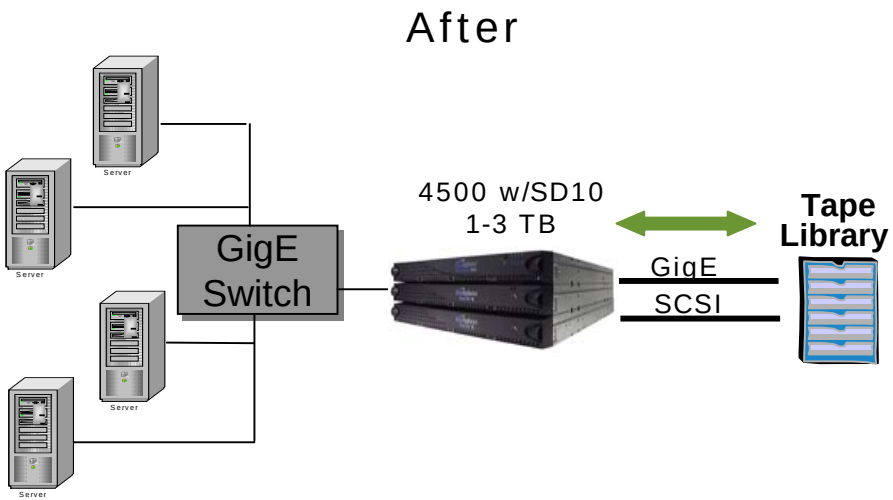

Target Application:

- Back Office
- SMB
- Retail
- Departmental

Host Mirrored Storage Pool for Application Server

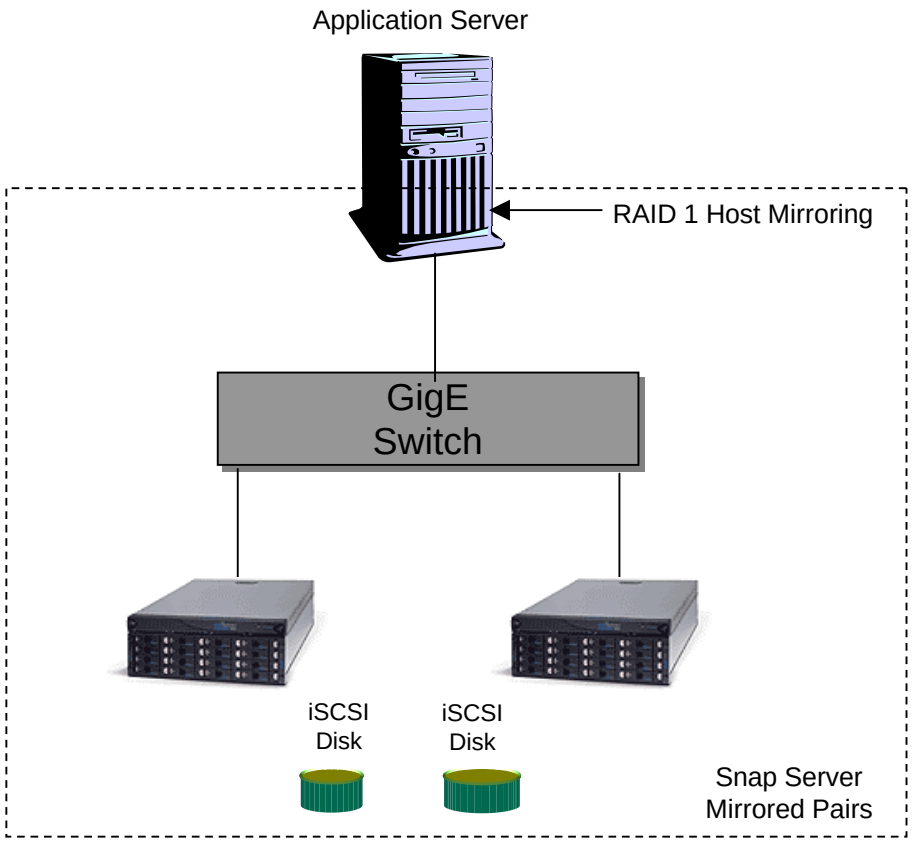

Value Proposition

More Simple and Cost-effective than Clustering

RAID 1 (Mirroring) at Host Protects Against Failure of Either of the Mirrored Storage Resources

Mirroring is an Easy-to-implement Redundancy Solution with a Minimal Performance Penalty

Target Application

- Any Direct Attached Storage Application with Cost Sensitivities and Data Redundancy Requirements

Common Storage Pool for Clustered Servers

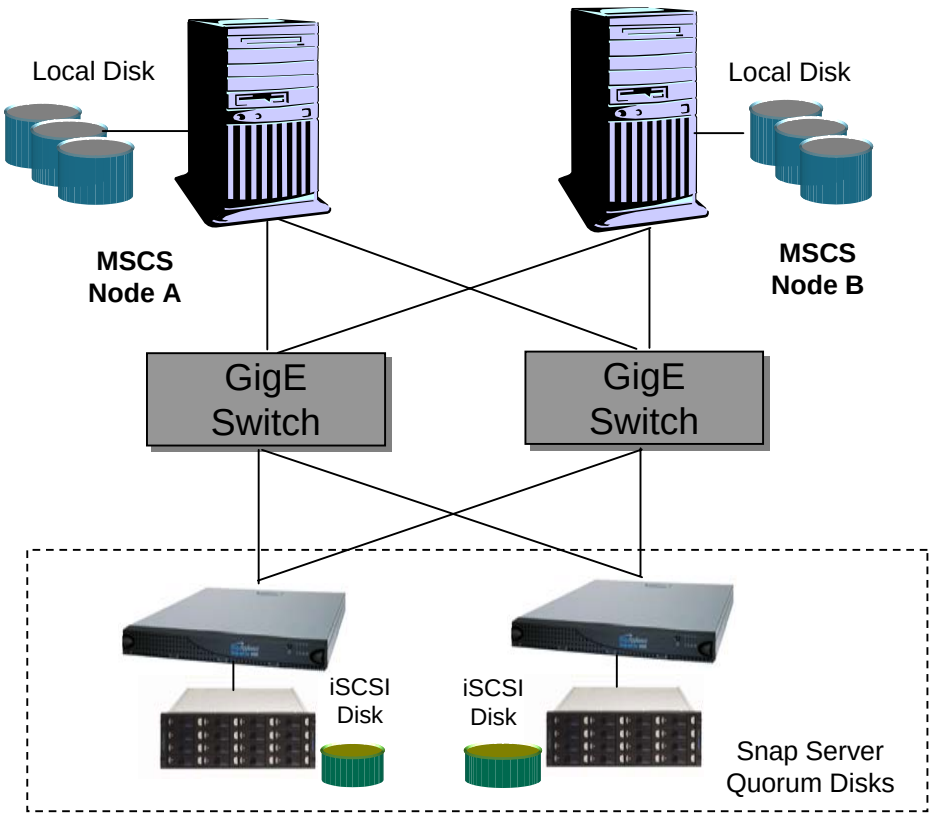

Value Proposition:

Scalable, High Availability Shared storage pool for a variety of applications. Uses Multi-Path I/O drivers for Resiliency of common storage pool. Easy to manage, cheaper than SAN/DAS by utilizing inexpensive Ethernet infrastructure and disks.

Target Application

- **Clustered Applications using local DAS storage.**
- **HA applications where Fibre Channel is cost prohibitive.**
- **Wide market applicability where a common storage pool is needed**
- **Clustered Exchange messaging solutions**

iSCSI Value Proposition

- **Ideal for Storage Consolidation**
  - LAN-free Backup
  - An Effective Solution for Extending the Life of Older Servers
- **Improves Storage Utilization, Scalability and Availability**
- **Lowest Initial Investment for any SAN Solution**
  - IP SANs use Ethernet, Substantially Lowering Acquisition Cost
- **Lowest TCO Among Storage Topologies to Date**
  - Lower Initial Costs
  - Reduces Training and Staff Costs Due to Abundance of Ethernet Expertise
- **Easy-to-use, Deploy, and Understand**
  - iSCSI Removes Complexity – Uses Gigabit Ethernet wiring and switches, a widely deployed and understood technology

Snap Servers - Ready for IP SANs

All Support BOTH Block (iSCSI) and File (CIFS/NFS) Services
Online Expandability of iSCSI Disks and Data Volumes

				
Model	4200	4500	4500 w/SD30	18000 w/SD30
Native Capacity	640	1TB – 3.6TB	1.6TB - 12TB	6TB to 30TB
Drives	**4 ATA-100 Hot swappable**	**4 ATA-100 Hot swappable**	**4 ATA-100 Per Chassis Hot swappable**	**Head: 4 ATA-100 Exp: 14 SATA Hot swappable**
RAID	**5,1,0 JBOD Global Hot Spare**	**5,1,0 JBOD Global Hot Spare**	**5,1,0 JBOD Global Hot Spare**	**5,1,0 JBOD Global Hot Spare**
Network Connections	**Dual 10/100/1000**	**Dual 10/100/1000**	**Dual 10/100/1000**	**Dual 10/100/1000**
Local Backup	**Ultra 160 SCSI BakBone NetVault**	**Ultra 160 SCSI BakBone NetVault**	**Ultra 160 SCSI BakBone NetVault**	**Ultra 160 SCSI BakBone NetVault**

Snap Appliance Software Solutions

- Replication of File Data
- Snapshots
- Unified, Multi-Snap Server Management
- Embedded Backup Server - using BakBone NetVault Workgroup Edition
- Enterprise Backup Agent Support for Veritas, Legato, CA, and BakBone
- NDMP Support

Benefits of Deploying Snap Servers

- **Unified Storage — Block and File in One Appliance!**
  - Easy Installation, Migration, and Management
  - Appliance Model
  - Integrates into Existing Environments
- **Scalability**
  - Seamlessly Upgrade Without Data Migration
  - Expand with Mixed RAID Sizes
  - Retain Enterprise Storage Investment
- **Simplified Management**
  - Easy, Browser Management
  - Unified, Multi-server Management
- **Full-featured Backup and Restore**
  - Virtual Tape Library Capability
  - Reduced Management with Centralized Backup
  - Increase Speed and Accuracy of Data Recovery

Reduced Overhead
Lower TCO
Integrates With Existing Infrastructure

Windows Logo Certification for iSCSI