



Education

# **STORAGE CONSOLIDATION WITH IP STORAGE**

David Dale, NetApp

- The material contained in this tutorial is copyrighted by the SNIA.
  - Member companies and individuals may use this material in presentations and literature under the following conditions:
    - ◆ Any slide or slides used must be reproduced without modification
    - ◆ The SNIA must be acknowledged as source of any material used in the body of any document containing material from these presentations.
  - This presentation is a project of the SNIA Education Committee.
  - Neither the Author nor the Presenter is an attorney and nothing in this presentation is intended to be nor should be construed as legal advice or opinion. If you need legal advice or legal opinion please contact an attorney.
  - The information presented herein represents the Author's personal opinion and current understanding of the issues involved. The Author, the Presenter, and the SNIA do not assume any responsibility or liability for damages arising out of any reliance on or use of this information.
- NO WARRANTIES, EXPRESS OR IMPLIED. USE AT YOUR OWN RISK.**

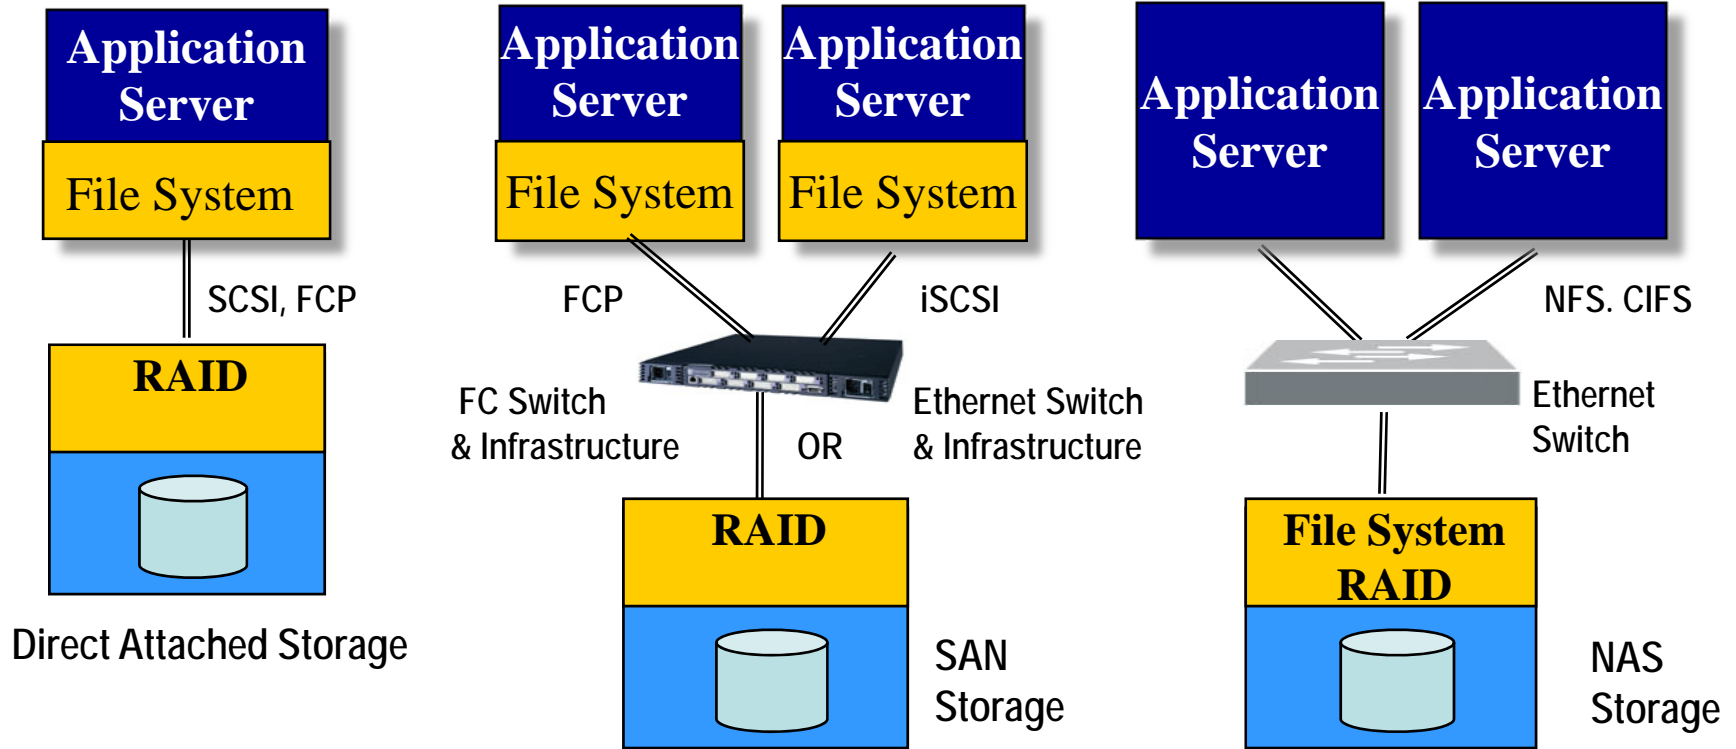
## Storage Consolidation with IP Storage

This session will appeal to IT managers, administrators and architects interested in best practices and deployment considerations of storage consolidation solutions available with IP Storage technologies today.

This presentation, an update to a very popular SNIA Tutorial, outlines the benefits of networked storage, contrasting the different options. It then goes into detail on iSCSI-based SAN configurations, capabilities, options and best practices, including contemporary iSCSI storage features. It then explains via customer case studies how these capabilities deliver benefits to IT organizations. Finally, emerging and future technologies and capabilities are considered, including 10GbE and FCoE.

- Storage Consolidation: Advantages of SANs
- Storage Area Networking with IP Storage
  - ◆ Topologies
  - ◆ Benefits
- iSCSI-based SANs
  - ◆ Connectivity and security
  - ◆ Host support
  - ◆ Performance
  - ◆ Typical array capabilities
- Typical iSCSI Deployment Deployments
  - ◆ Where IP Storage Fits
  - ◆ Typical Reasons for Deployment
  - ◆ Deployment Case Studies
- What's Next
- Summary

# Storage Technologies



- Server-based data management
- No resource sharing
- No data sharing
- Works with all apps


- Server-based data management
- Resource sharing
- No data sharing
- Works with all apps

- Storage-based data management
- Resource sharing
- Data sharing
- Works with qualified apps


# What SANs Deliver

- ◆ Value of Storage Networking
  - ◆ Improved reliability and reduced cost of backup
  - ◆ Improved scalability of storage capacity and performance
  - ◆ Simplified storage provisioning
  - ◆ Improved data availability
- ◆ Top reasons for deploying a SAN
  - ◆ Back-up
  - ◆ Storage consolidation
  - ◆ Satisfy on-going demands for additional capacity
  - ◆ Performance
  - ◆ Disaster recovery
  - ◆ New project or application deployment

# iSCSI SAN Options

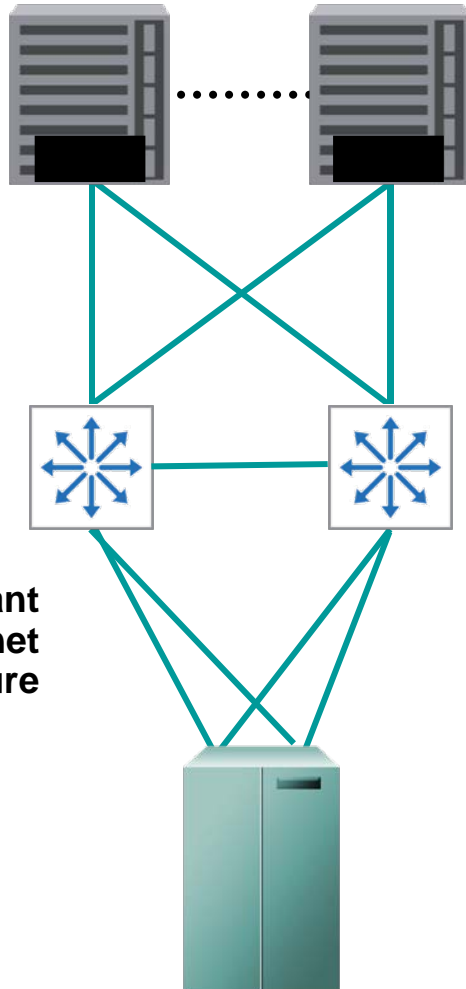


# IP SAN Benefits



- **Standard SAN storage**
  - ◆ Block storage access
  - ◆ Supports all apps
  - ◆ Transparent migration from direct attached storage
- **Lower TCO than FC**
  - ◆ Zero host connection cost
  - ◆ Less costly infrastructure
  - ◆ Easier to manage
- **Leverages IP Expertise**
  - ◆ Expertise in existing staff
  - ◆ Robust well-understood management software
  - ◆ Easily enables remote integration of data assets

## Host Systems



## Connectivity:

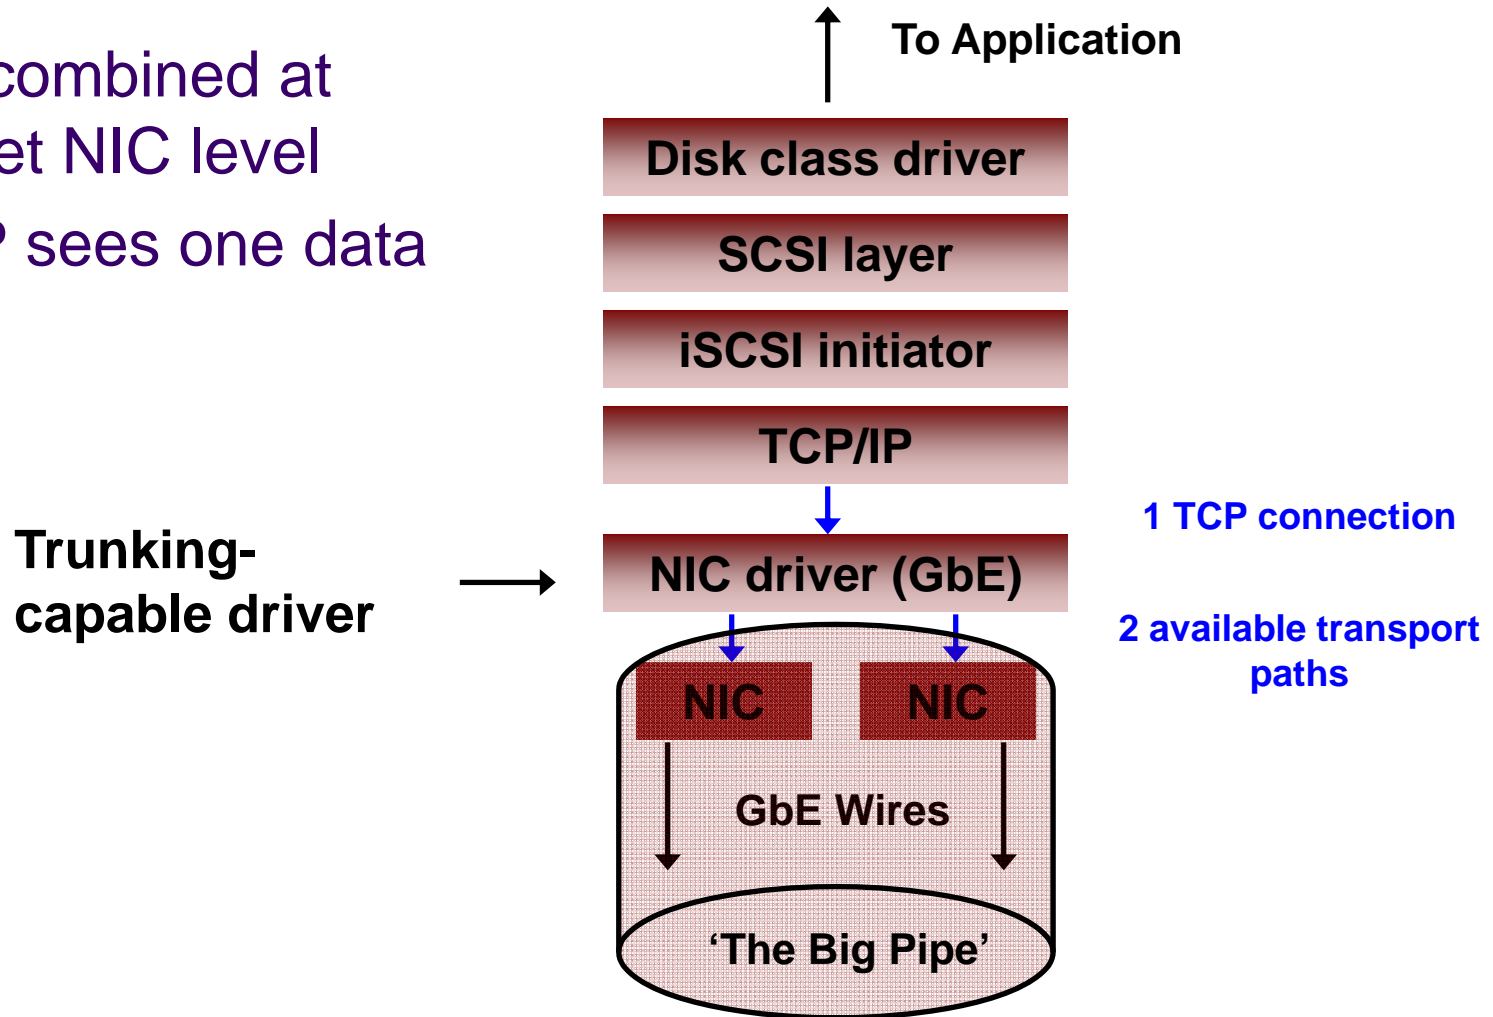
- ▶ 1 Gb or 10Gb Ethernet
- ▶ Jumbo frames (recommended)
- ▶ Link aggregation or MPIO or Multi-Connection Sessions (bandwidth and/or availability)

## Security:

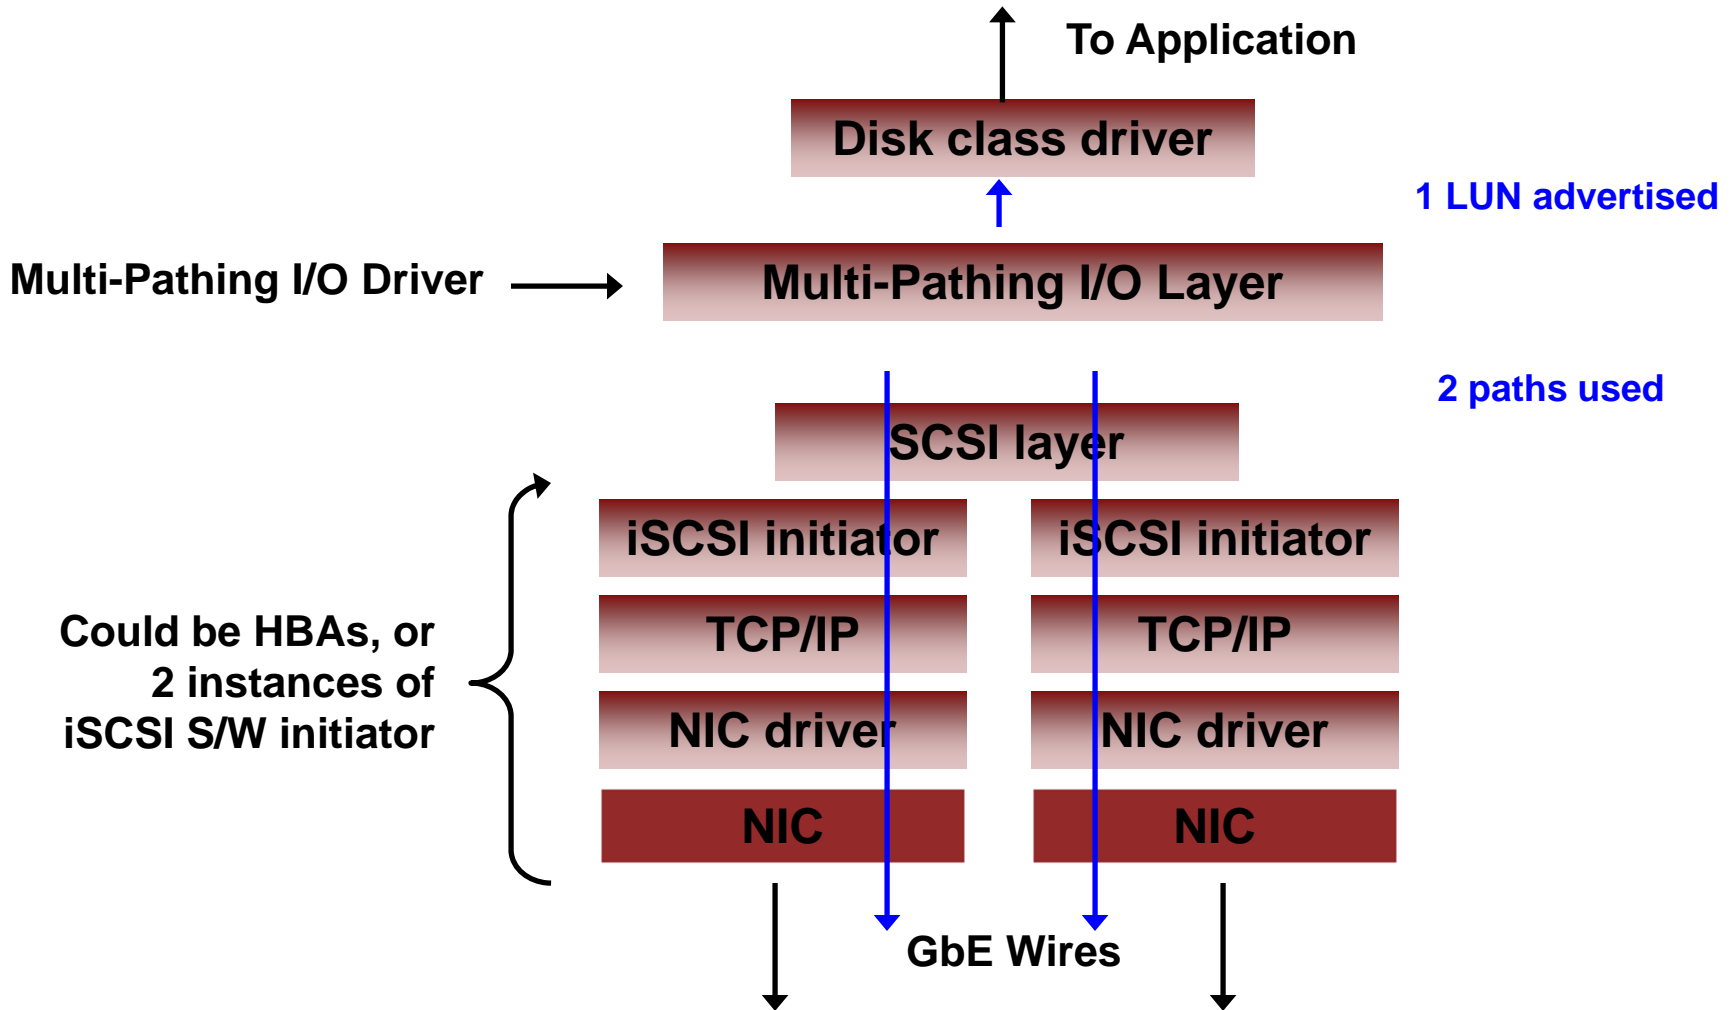
- ▶ Host authentication (CHAP)
- ▶ Private network
  - ▶ Physical
  - ▶ VLAN (zoning)
- ▶ Array LUN masking
- ▶ Optional IPsec
- ▶ Optional key management

# HA Options: Link Aggregation

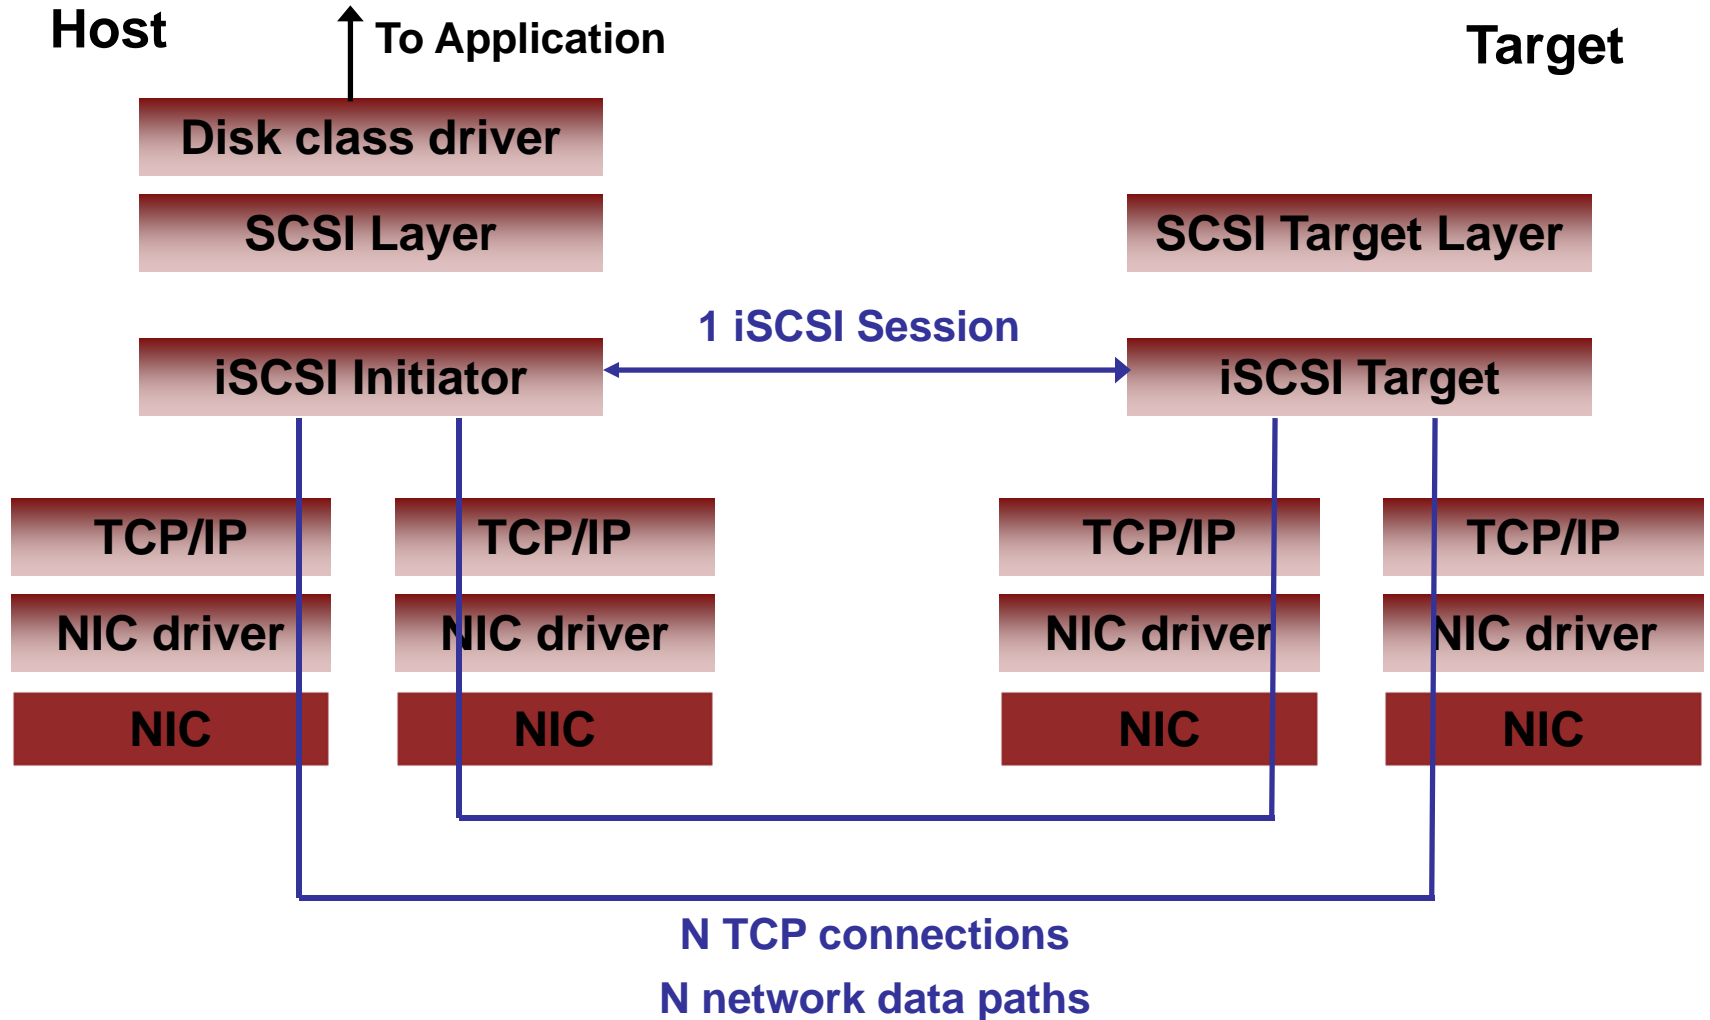
- ▶ Paths combined at Ethernet NIC level
- ▶ TCP/IP sees one data path










# HA Options: Multi-Pathing I/O



# HA Options: Multi-Connection iSCSI Sessions




# IP SAN Host Support

OS	Initiator	Certified	Multi-pathing	Cluster
	Hardware, Software	<input checked="" type="checkbox"/>	MPIO, MCS	Yes
	Hardware, Software	<input checked="" type="checkbox"/>	Trunking, MPxIO	Yes
	Software	<input checked="" type="checkbox"/>	PV Links	TBD
	Software	<input checked="" type="checkbox"/>	Trunking	TBD
	Hardware, Software	<input checked="" type="checkbox"/>	Trunking; MPIO	Yes
	Hardware, Software	<input checked="" type="checkbox"/>	Trunking, MPIO	Yes
Novell. NetWare.	Software	<input checked="" type="checkbox"/>	Trunking	Yes
	Software	<input checked="" type="checkbox"/>	Trunking	Yes

# I/O Performance Considerations

- **Software iSCSI initiator + standard NIC (port)**
  - ◆ Host CPU overhead (up to 500MHz)
  - ◆ Low cost (free download)
  - ◆ Adequate for vast majority of applications
  - ◆ Most popular solution today (85% deployments)
- **TCP Offload (TOE) NIC**
  - ◆ Lower host CPU overhead
  - ◆ Uses OS-native iSCSI software initiator
- **iSCSI HBA**
  - ◆ Overhead similar to Fibre Channel HBA
  - ◆ Rarely used

## ➤ Enterprise Strategy Group Validation study (4/04)

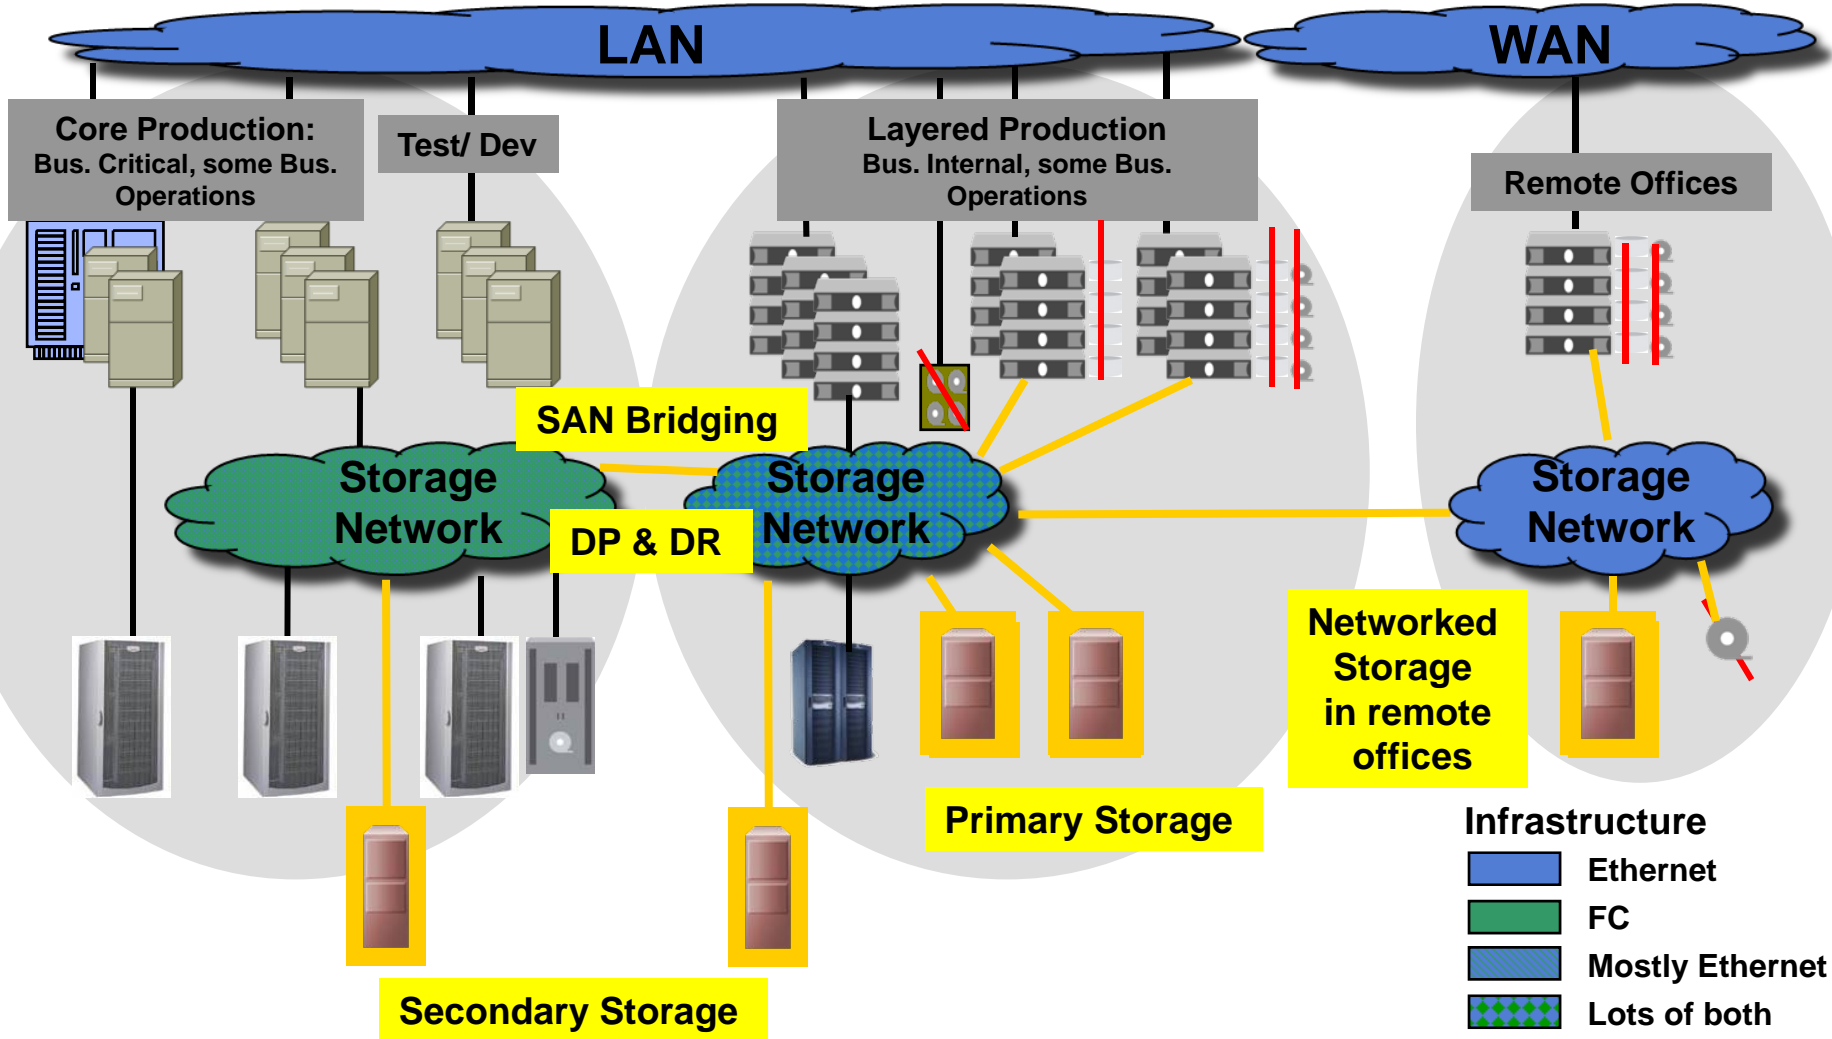


# Typical iSCSI Array Capabilities

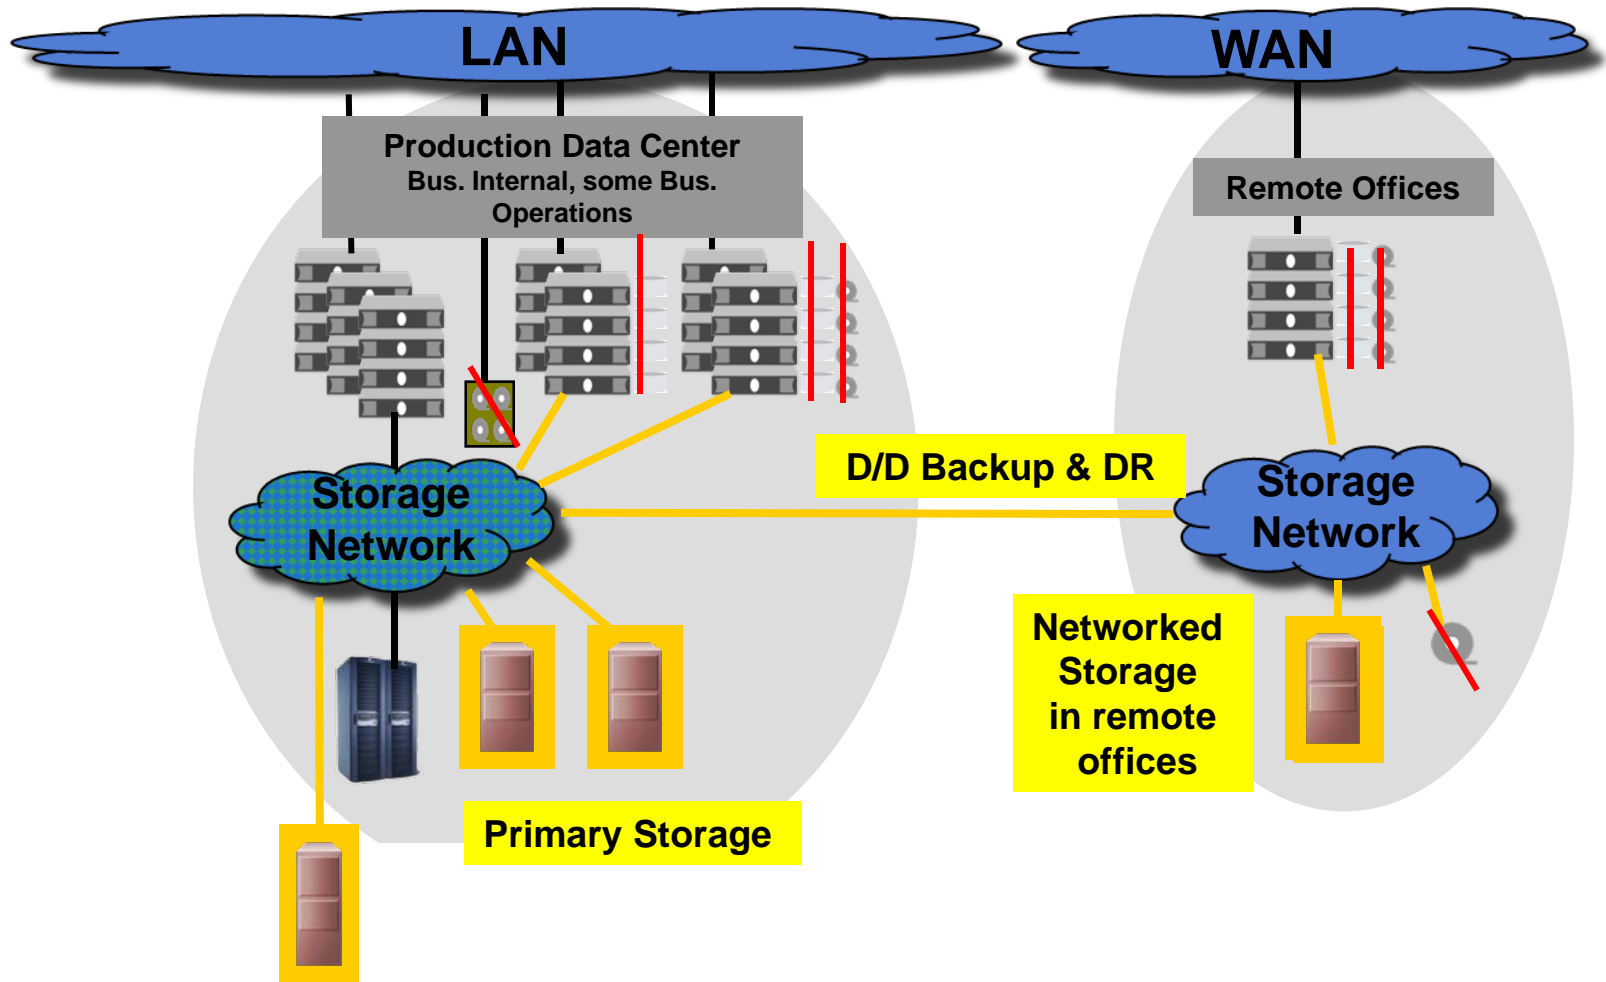
- Basic storage considerations
  - ◆ Redundant components
  - ◆ Dual active controllers with failover
  - ◆ RAID
  - ◆ SATA drives; FC drives; SAS drives
- Storage features
  - ◆ Point in time copies (Snapshot)
  - ◆ Network Boot
  - ◆ Multi-path I/O for High Availability
  - ◆ Thin provisioning (sparse allocation)
  - ◆ Remote data copy
  - ◆ Asynchronous mirroring for disaster recovery
- Growth/scalability/configurability
  - ◆ Capacity
  - ◆ Performance
  - ◆ Host integration



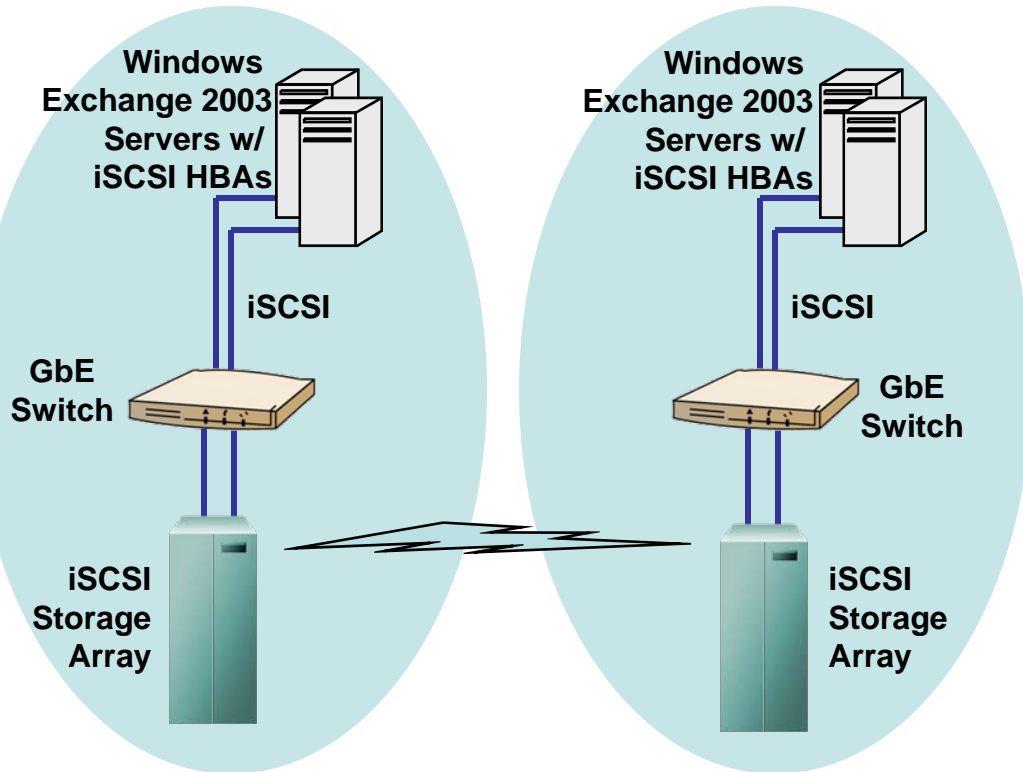
# Where IP Storage Fits - Large Enterprise



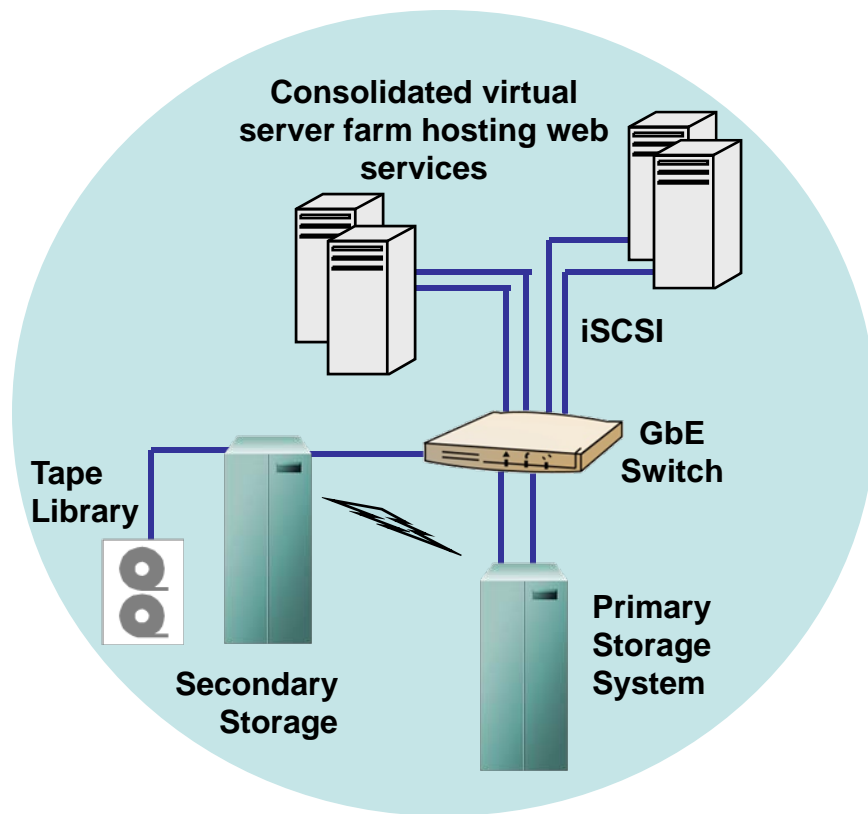
# Where IP Storage Fits - Medium/Small Enterprise



## Consolidated Regional Exchange 2003 Environment




- Application
  - *Microsoft Exchange 2003*
- Pain Points
  - *Exchange environment on 32 servers*
  - *Reliability issues with DAS*
  - *Unacceptable recovery time (18-24 hours)*
  - *No Disaster Recovery solution*
- Solution
  - *Consolidation using IP-SAN*
  - *Async mirroring between sites*
- Benefits
  - *No downtime or outages*
  - *Exchange recovery now <25 minutes*
  - *Support costs cut by one third*
  - *Affordable*
  - *Disk-to-disk backup cuts admin costs by 90%*
  - *Diskless boot improves server support environment*



- Application
  - ◆ *Web hosting services*
- Pain Points
  - ◆ *Rapid growth*
  - ◆ *Outgrowing “green” data center*
  - ◆ *Very poor server utilization*
  - ◆ *Disruptive backup process*
- Solution
  - ◆ *Virtual servers w/ IP SAN*
  - ◆ *Disk-to-disk backup for HA and DR*
- Benefits
  - ◆ *Virtual server environment provides the flexibility to host additional clients and increase revenue potential.*
  - ◆ *Server and storage consolidation reduced data center power consumption by 60 percent.*
  - ◆ *Replacement of 120 white box servers with four SMP servers reduced cooling costs and data center footprint.*
  - ◆ *Cost savings and cost avoidance enabled pursuit of additional environmental conservation solutions.*

# Case Study



## Application

- ◆ SQL Server databases
- ◆ Microsoft Exchange

## Pain Points

- ◆ Challenge to meet SLAs with direct attached storage environment
- ◆ Affordable SAN storage for SQL Server databases
- ◆ Affordable DR solution

## Solution

- ◆ IP SAN in each location
- ◆ Multiple snapshot backups per day to US core data center
- ◆ Async mirror to DR site

## Benefits

- ◆ High performance solution
- ◆ Simple, cost-effective storage network
- ◆ DR between existing data centers
- ◆ Enhanced ability to meet demanding airline customer SLAs

- Mainstream deployment in multi-OS host environments (Windows, Linux, Unix)
- Mainstream deployment in virtual server environments (VMWare, Xen, Hyper-V)
- Broad deployment in small blade server environments
- Deployment in large grid (blade) scale-out environments with 10GbE backbone
- Mainstream deployment with 10GbE infrastructure
- Standards in process
  - ◆ FCoE
  - ◆ RDMA/10GE (DDR and RDMAP)
  - ◆ iSCSI Extensions for RDMA (iSER)

# 10 Gb Ethernet Background

- High-speed networking solutions for LAN, MAN, WAN and SAN environments
- IEEE 802.3ae ratified 2002
  - ◆ Supports all layer 2, 3 and higher network services
- Optical and copper physical media supported
  - ◆ 10GBASE-E, 10GBASE-L, 10GBASE-S
  - ◆ 10GBASE-CX4; 10GBASE-T
- Requirements for proliferation – now happening
  - ◆ Affordable Price
  - ◆ Server architecture support
  - ◆ Standard TOE-supporting software architecture
  - ◆ TOE built into standard on-board components

# 10 Gb Ethernet Today

## ➤ Deployment/applications

- ◆ MAN connectivity plus SONET alternative in MAN
- ◆ Backbone and port aggregation for 1Gb LANs
  - Driven by acceleration in 1Gb deployment over past 4 years
- ◆ Storage over 10GbE
  - NFS and iSCSI; FCoE coming

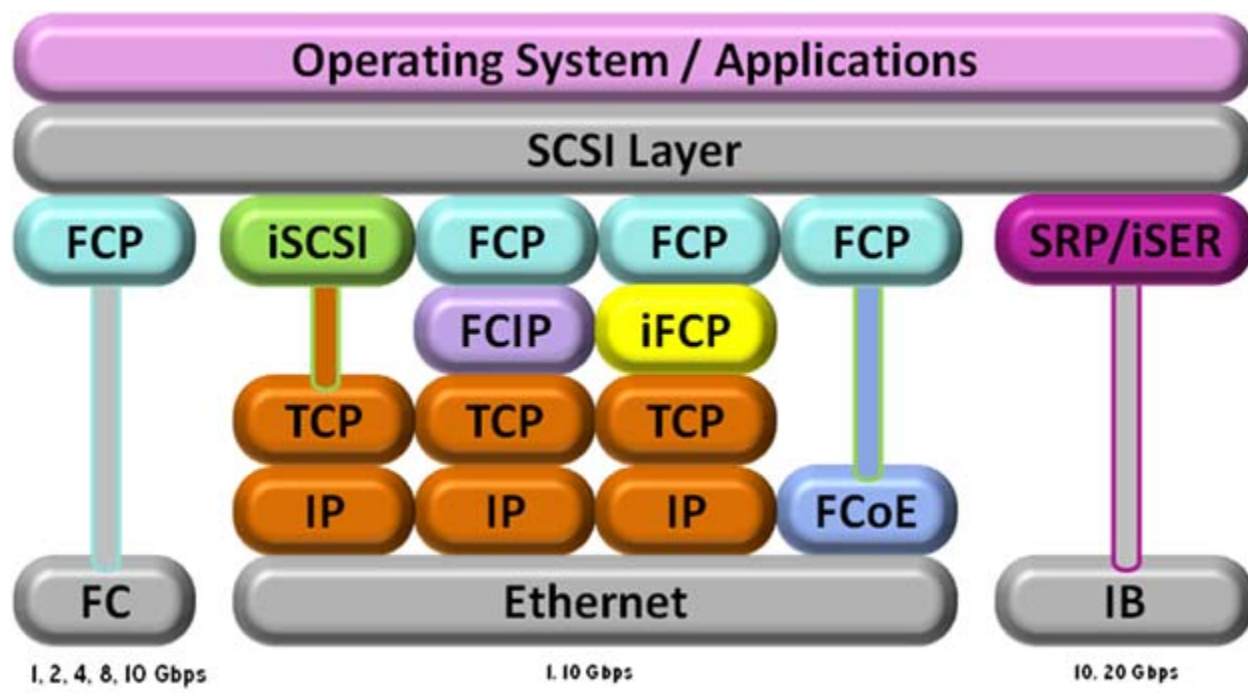
## ➤ Pricing

- ◆ Greatly reduced pricing over past 5 years
- ◆ Cost effective aggregation alternative to 10x 1Gb

## ➤ Architecture and System Support

- ◆ High-speed I/O (PCI-X and PCI-express) now mainstream
- ◆ Defacto standard software TOE support with “TCP Chimney”
- ◆ 10Gb Ethernet connections available on many storage products

# Fibre Channel over Ethernet (FCoE)



- FCoE (Fibre Channel over Ethernet)

- An Extension of Fibre Channel onto a 10Gb Ethernet network
- FCoE is a direct mapping of Fibre Channel over Ethernet
- Unlike other protocols such as iSCSI, FCIP, and iFCP, TCP/IP is not required and not present for FCoE
- Preserves the operational, control, and management environments for the FC layer

# FCoE Primer

## ➤ Status

- ◆ Protocol in development in INCITS T11 committee
- ◆ “Data Center Ethernet” enhancements in development in IEEE 802.3 committee

## ➤ Rollout dependencies

- ◆ FCoE-only deployment can use existing Ethernet features (Pause)
- ◆ Unified (FCoE and TCP/IP traffic) deployment requires DCE enhancements
- ◆ Operating system support and qualification

## ➤ Benefits

- ◆ Enables unified data center network
  - › Fibre Channel and TCP/IP traffic over 10Gb Ethernet
  - › Maximum flexibility;
  - › Reduced power and cooling
- ◆ Investment protection for existing FC environments
  - › Reduced training and administrative costs
  - › Same Management Model as Native FC

# Summary - IP Storage

- Sophisticated storage consolidation solutions for low-end and mid-range server environments
- Takes advantage of existing IT knowledge base
- Provides simpler, more affordable SAN infrastructure
- Improves data availability and performance
- Integrates distributed data and resources
- Solutions are deployed in many thousands of companies around the world
- Ultimately provides one technology for connecting clients, servers & storage devices

- Please send any questions or comments on this presentation to SNIA: [trackstorage@snia.org](mailto:trackstorage@snia.org)

**Many thanks to the following individuals  
for their contributions to this tutorial.**

*SNIA Education Committee*

**David Dale, NetApp  
Suzanne Morgan, Microsoft  
SNIA IP Storage Forum Members**