



Education

STORAGE CONSOLIDATION WITH IP STORAGE

David Dale, NetApp

- The material contained in this tutorial is copyrighted by the SNIA.
- Member companies and individual members may use this material in presentations and literature under the following conditions:
 - ◆ Any slide or slides used must be reproduced in their entirety without modification
 - ◆ The SNIA must be acknowledged as the 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, or should be construed as legal advice or an opinion of counsel. If you need legal advice or a legal opinion please contact your attorney.
- The information presented herein represents the author's personal opinion and current understanding of the relevant 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.

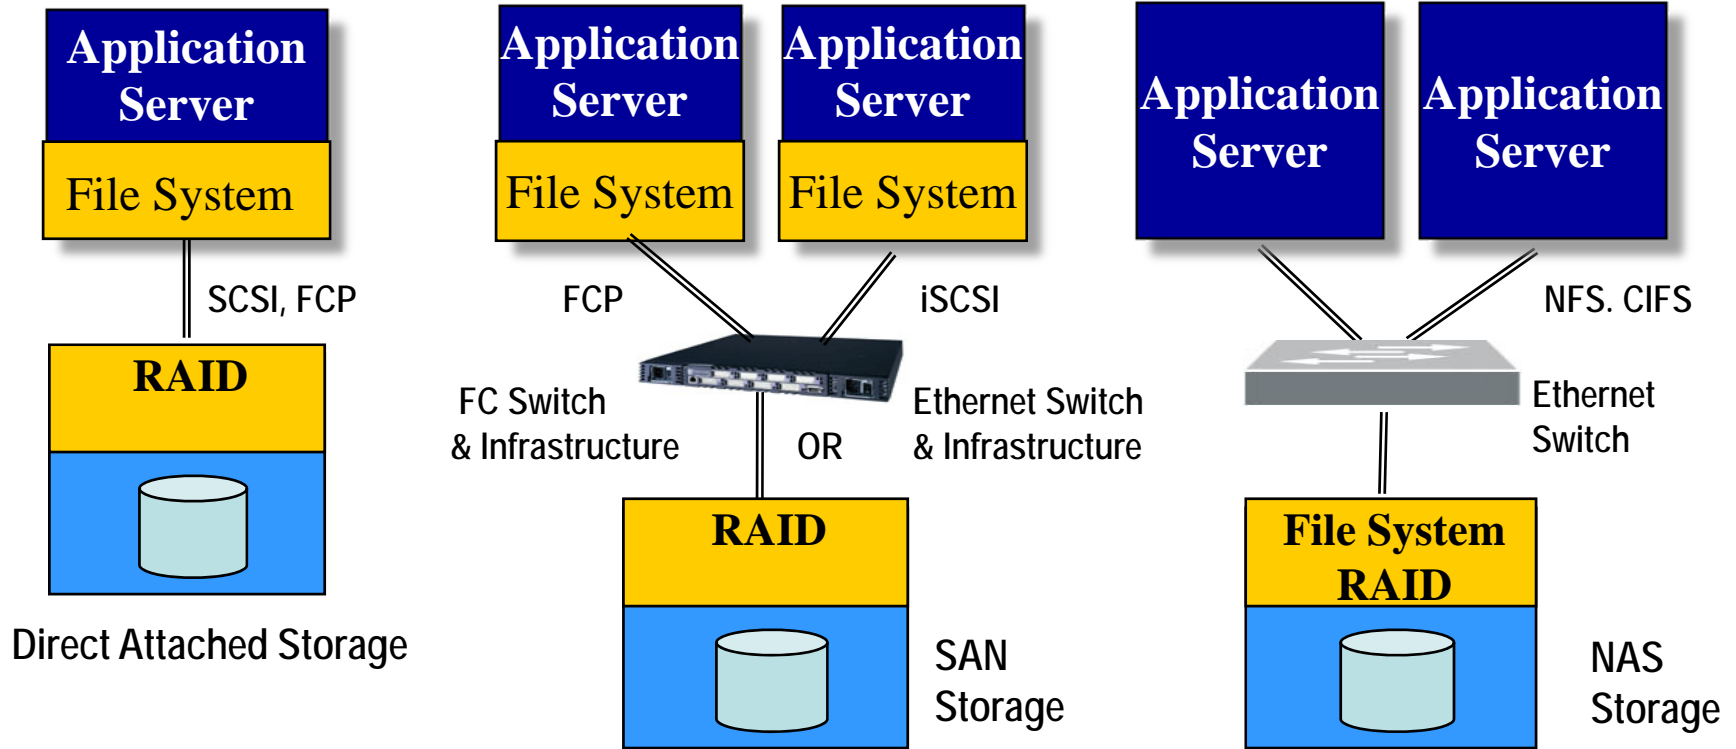
Storage Consolidation with IP Storage

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 higher speed Ethernet 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


- 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

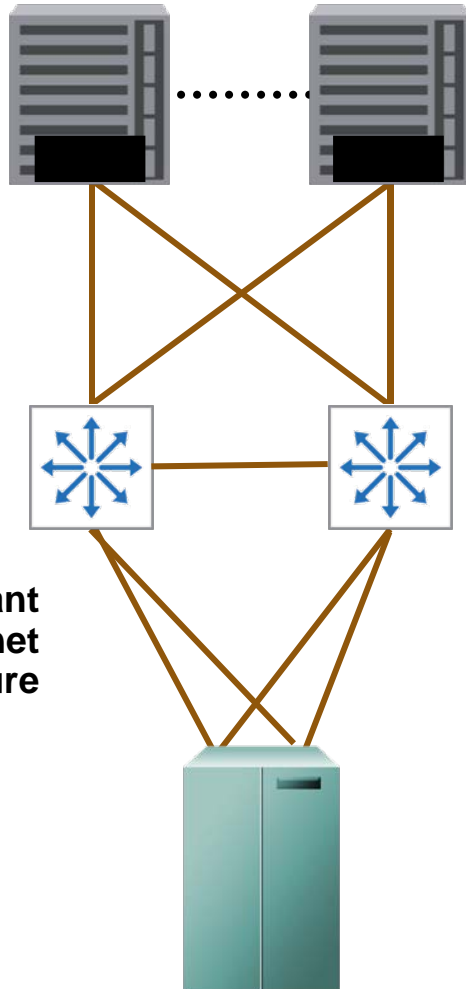


IP SAN Benefits



- **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



Redundant
Ethernet
Infrastructure

Connectivity:

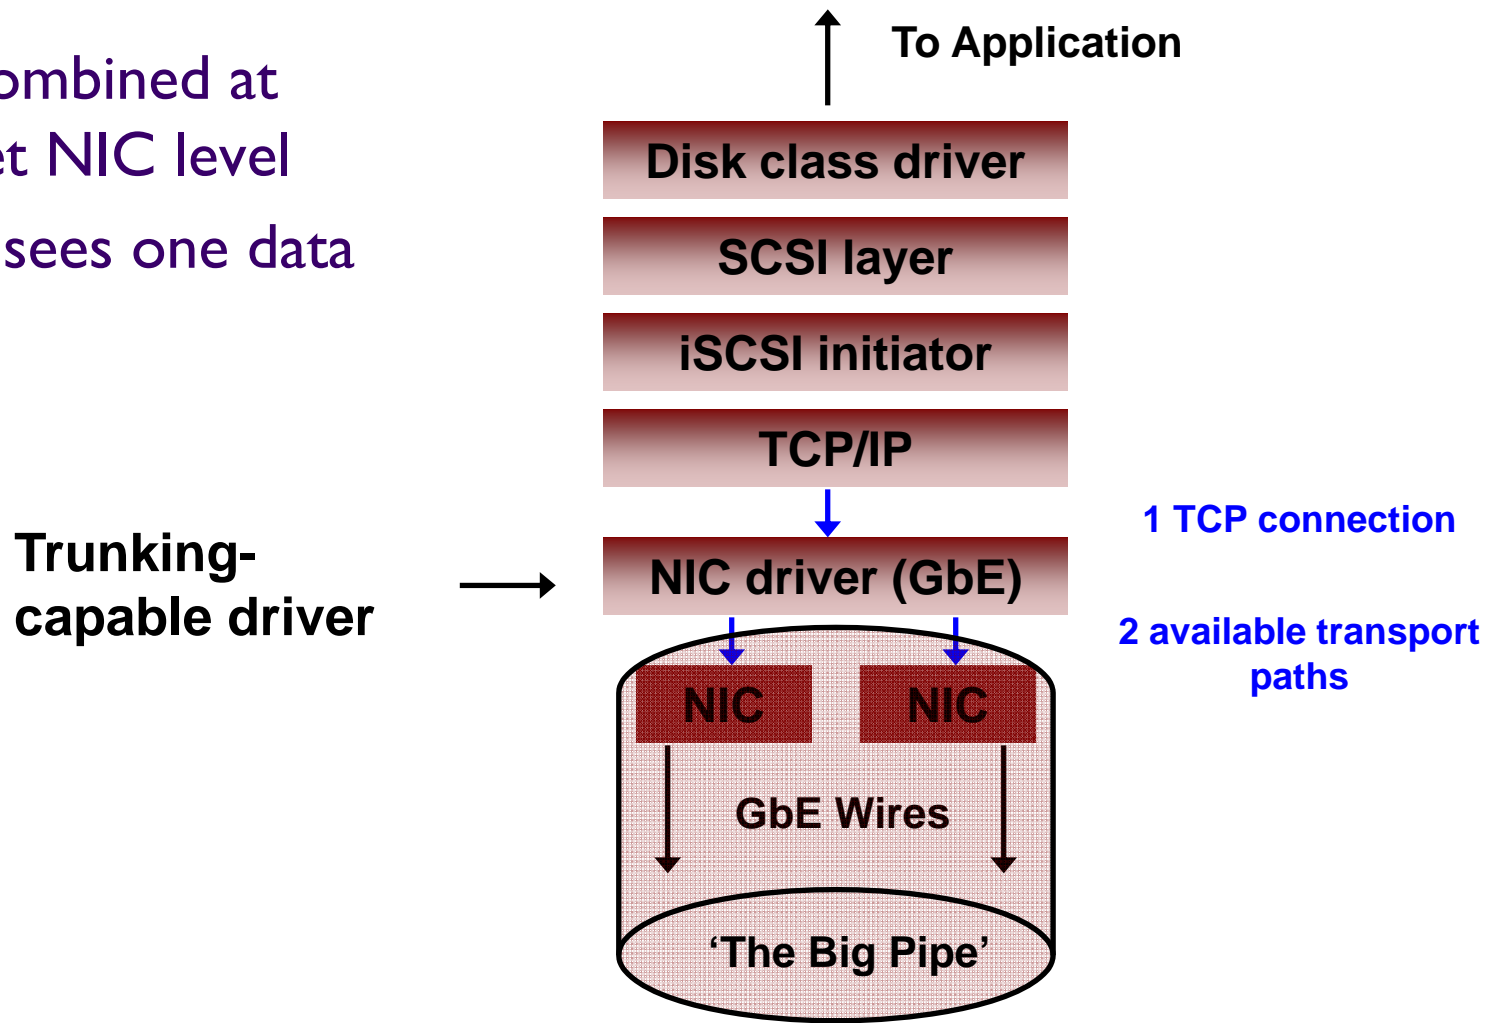
- ▶ 1 Gb or 10Gb Ethernet
- ▶ Jumbo frames (recommended)
- ▶ Link aggregation or MPIO or Multi-Connection Sessions (bandwidth and/or availability)

Security:

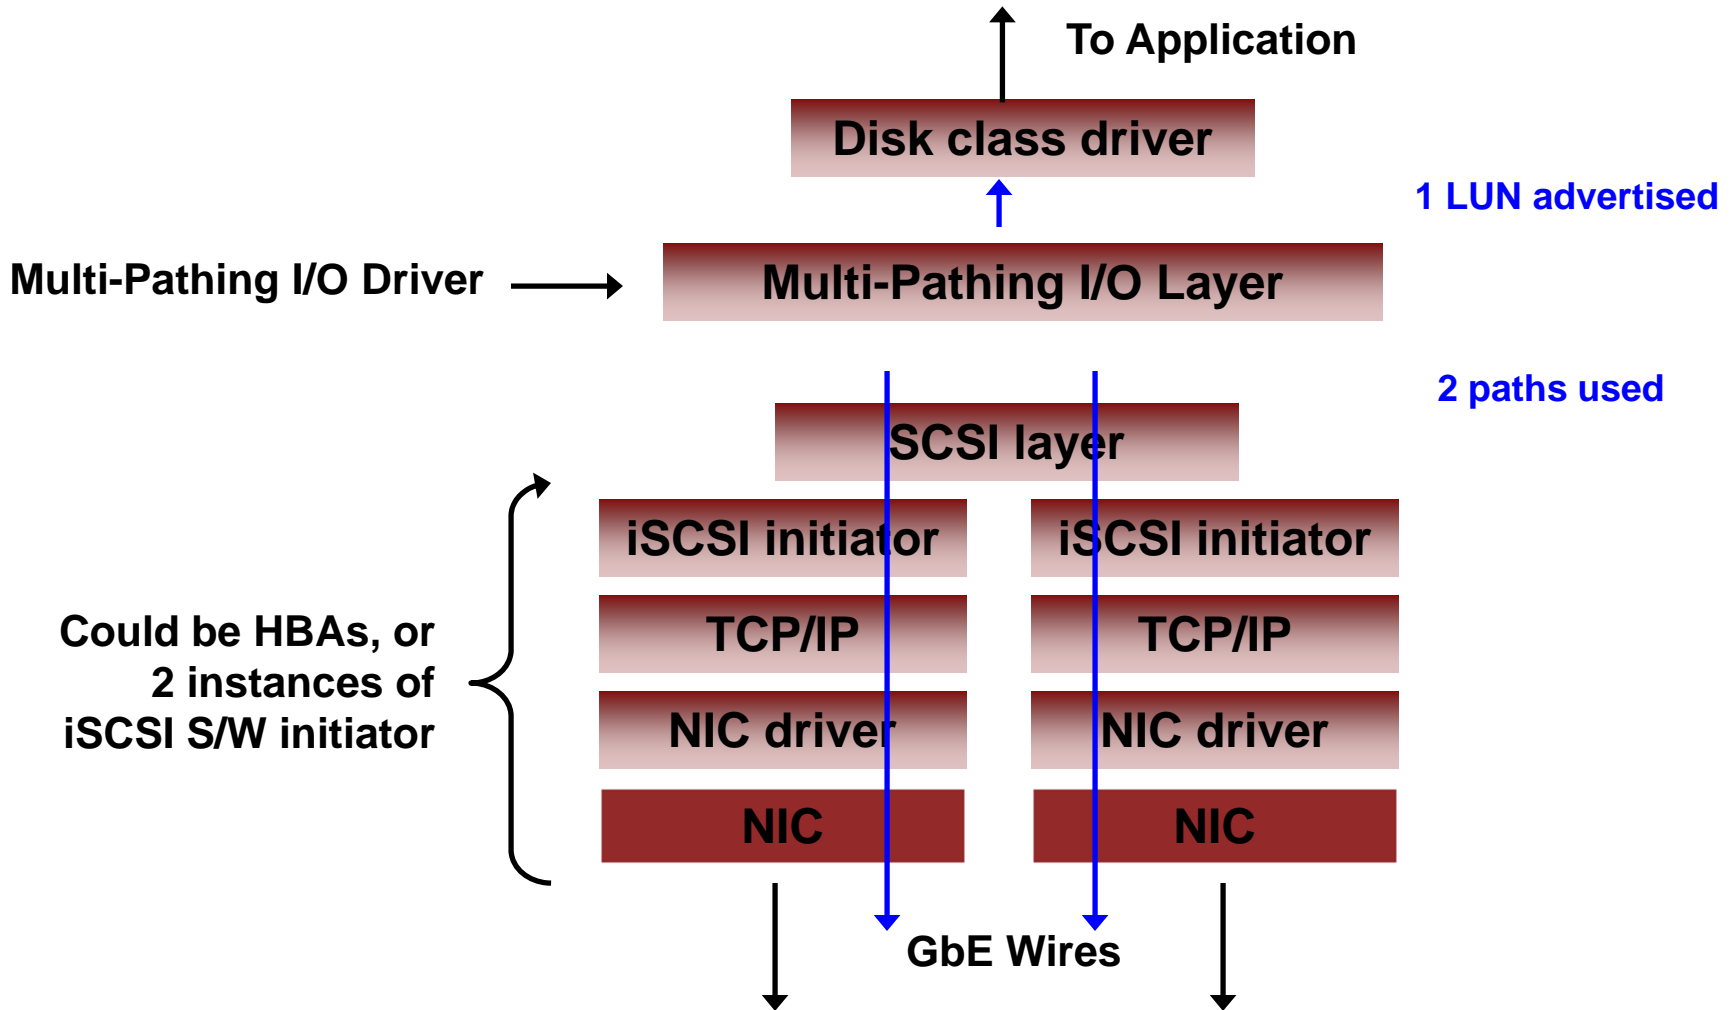
- ▶ Host authentication (CHAP)
- ▶ Private network
 - ▶ Physical
 - ▶ VLAN (zoning)
- ▶ Array LUN masking
- ▶ Optional IPsec
- ▶ Optional key management

HA Options: Link Aggregation

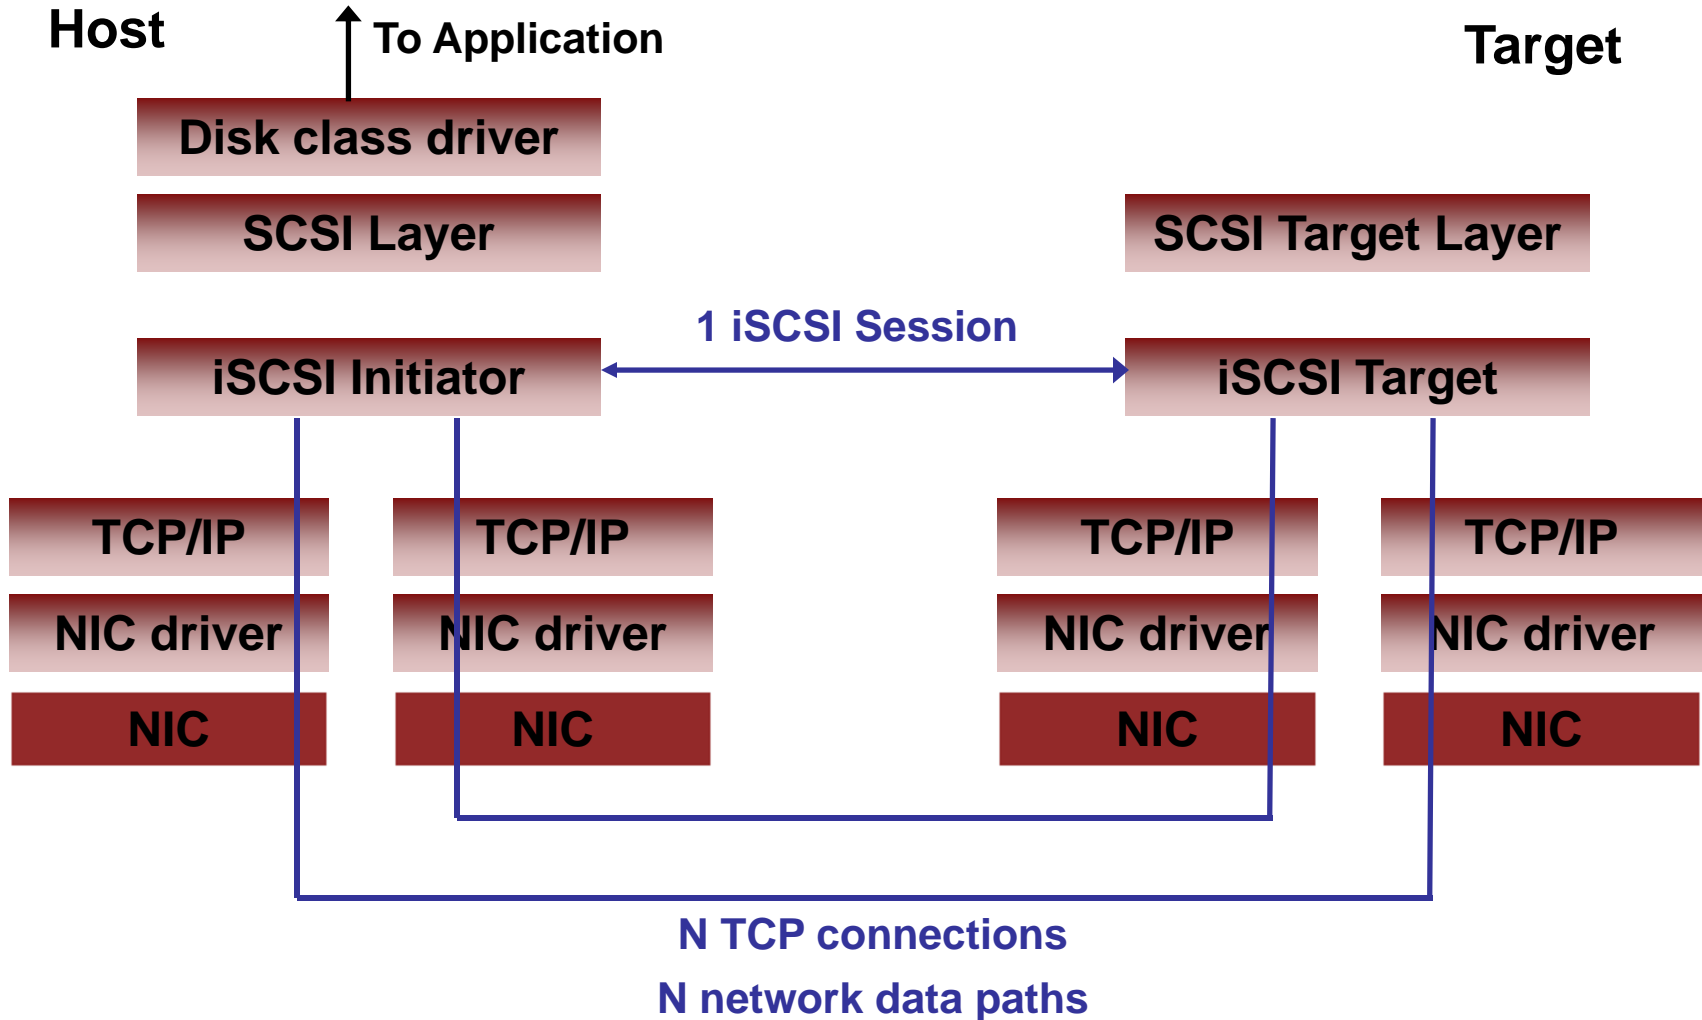
- ▶ Paths combined at Ethernet NIC level
- ▶ TCP/IP sees one data path










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)



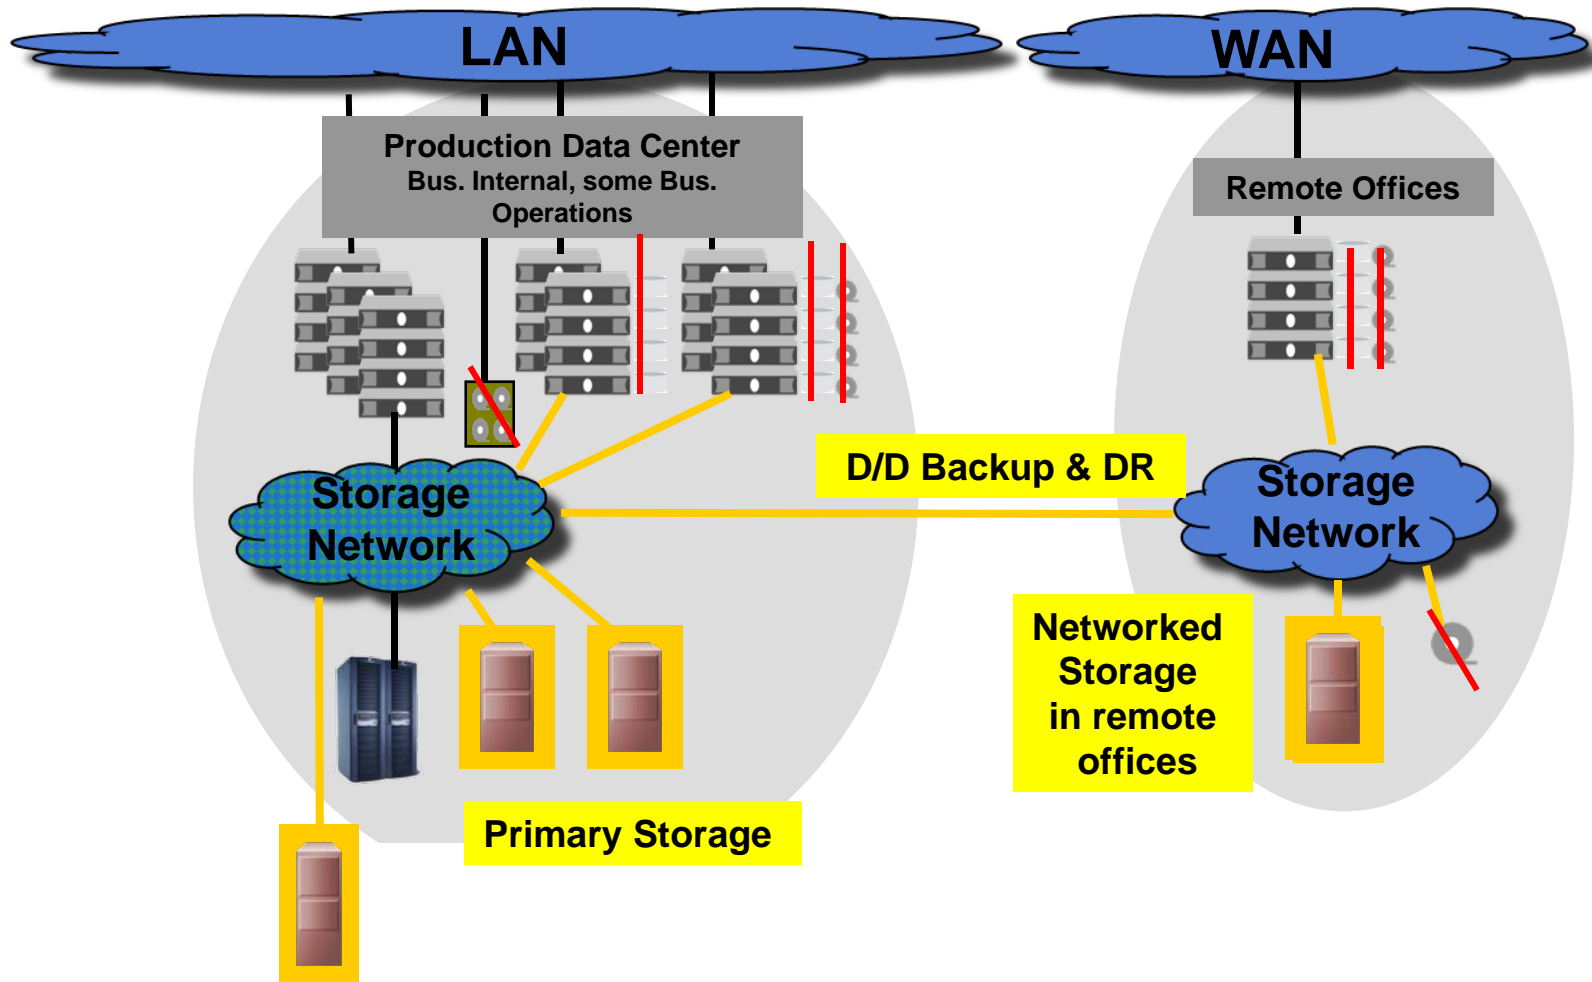
Typical iSCSI Array Capabilities

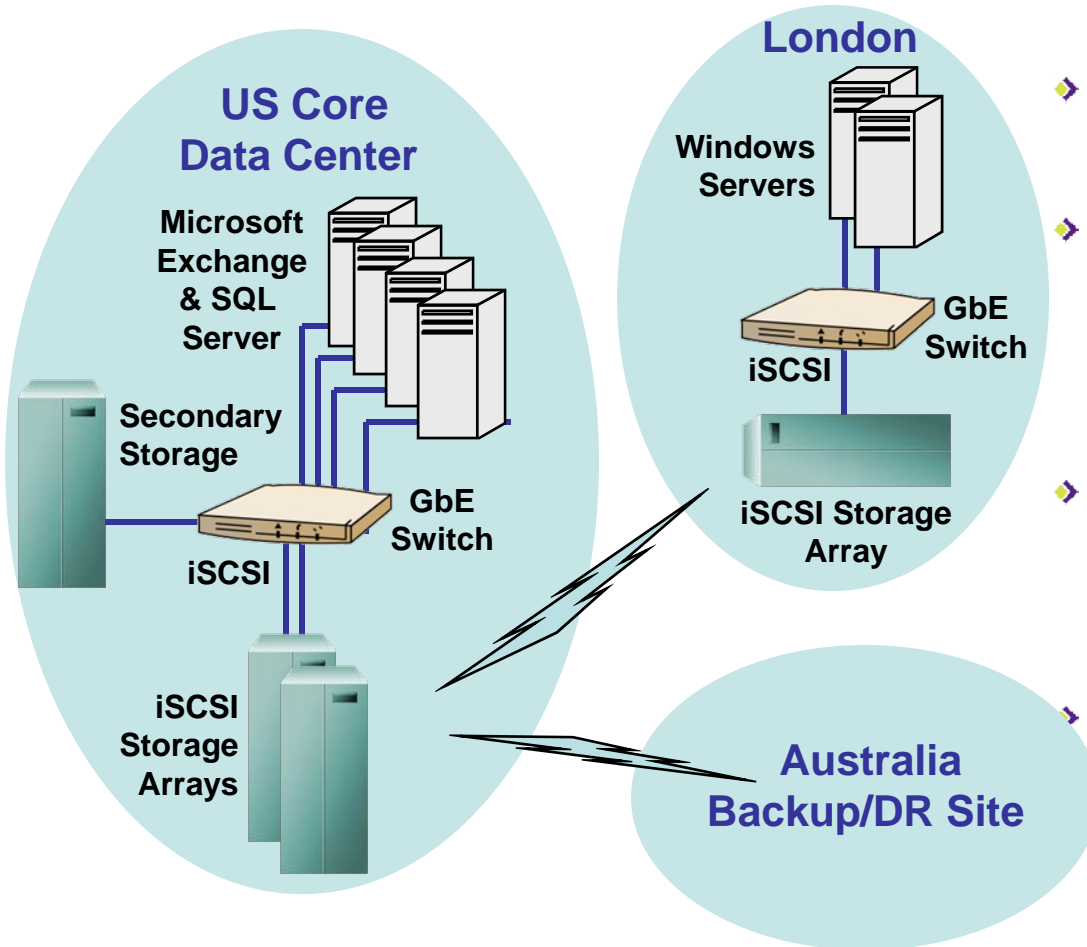
- **Basic storage considerations**
 - ◆ Redundant components
 - ◆ Dual active controllers with failover
 - ◆ RAID
 - ◆ SATA drives; FC drives; SAS drives
- **Storage features**
 - ◆ Point in time copies (Snapshot)
 - ◆ Network Boot
 - ◆ Multi-path I/O for High Availability
 - ◆ Thin provisioning (sparse 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





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



➤ 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.

- 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
 - ◆ 40 Gb and 100 Gb Ethernet
 - ◆ RDMA/10GE (DDR and RDMAP)
 - ◆ iSCSI Extensions for RDMA (iSER)

10 Gigabit Ethernet Background


- ◆ IEEE 802.3ae ratified 2002
 - ◆ Supports all layer 2, 3 and higher network services
- ◆ Broadly deployed in inter-switch links
- ◆ Requirements for host-side proliferation
 - ◆ Affordable Price
 - ◆ Server architecture support
 - ◆ TOE built into standard on-board components
 - ◆ This is now happening
- ◆ Deployment/applications
 - ◆ MAN connectivity plus SONET alternative in WAN
 - ◆ Backbone and port aggregation for 1 Gb LANs
 - › Driven by acceleration in 1 Gb deployment over past 5 years
 - ◆ File and block storage over 10GbE
 - › NFS and iSCSI; FCoE coming

Ethernet beyond 10Gb



Source: John D'Ambrosia, Force 10 Networks
Chair, IEEE P802.3ba Task Force; Ethernet Alliance

Fibre Channel over Ethernet



- An Extension of Fibre Channel onto a 10Gb Ethernet network
- FCoE is a direct mapping of Fibre Channel over Ethernet
- TCP/IP is not required and not present for FCoE
- Preserves ops, control and management environments for the FC layer

IEEE 802.1 DCB Protocol Status

- FCoE requires “lossless” Ethernet
 - ◆ Possible with Ethernet plus some extensions
- The IEEE 802.1 DCB WG is defining these extensions
 - ◆ Priority-based Flow Control (PFC): 802.1Qbb
 - Draft 0.2 available, large consensus
 - ◆ Enhanced Transmission Selection (ETS): 802.1Qaz
 - Including DCBX (DCB eXchange protocol)
 - ◆ Congestion Notification (CN): 802.1Qau
 - Draft 1.2 available, work still in progress
- Standards completion expected 2H2009; products 1H2010
- FCoE I/O Consolidation requirements:
 - ◆ PFC is required,
 - ◆ ETS is highly recommended,
 - ◆ CN is optional (not required for initial FCoE deployments)
- DCB required for multiprotocol support (FCoE and TCP/IP)

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

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

SNIA Education Committee

**David Dale
Suzanne Morgan
SNIA Ethernet Storage Forum iSCSI SIG**