



Education

# Server and Storage Consolidation with iSCSI Arrays

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.**

## Server & Storage Consolidation with iSCSI

This session will appeal to IT managers, administrators and architects interested in best practices and deployment considerations of storage consolidation solutions available with iSCSI-based systems today.

This presentation, an update to a very popular SNIA Tutorial, outlines the benefits of networked storage, contrasting deployment models. . After a summary of iSCSI-based SAN benefits, the presentation provides a detailed description of iSCSI SAN configurations, capabilities, options and best practices. As iSCSI has a natural affinity to virtualization, it also covers virtual server environments. Finally, the presentation describes typical deployment scenarios, and emerging developments, including higher speed Ethernet and FCoE.

- **Storage Consolidation with iSCSI SANs**
  - ◆ Advantages of networked storage
  - ◆ iSCSI SAN benefits
- **iSCSI –based Storage Area Networks**
  - ◆ Host connectivity and security
  - ◆ Boot from SAN
  - ◆ High availability
  - ◆ Quality of service
  - ◆ iSCSI performance
- **Deployments and Futures**
  - ◆ Typical array capabilities
  - ◆ Server virtualization with iSCSI
  - ◆ Typical deployment scenarios
  - ◆ Emerging new capabilities
- **Summary**

# Why Move to Networked Storage

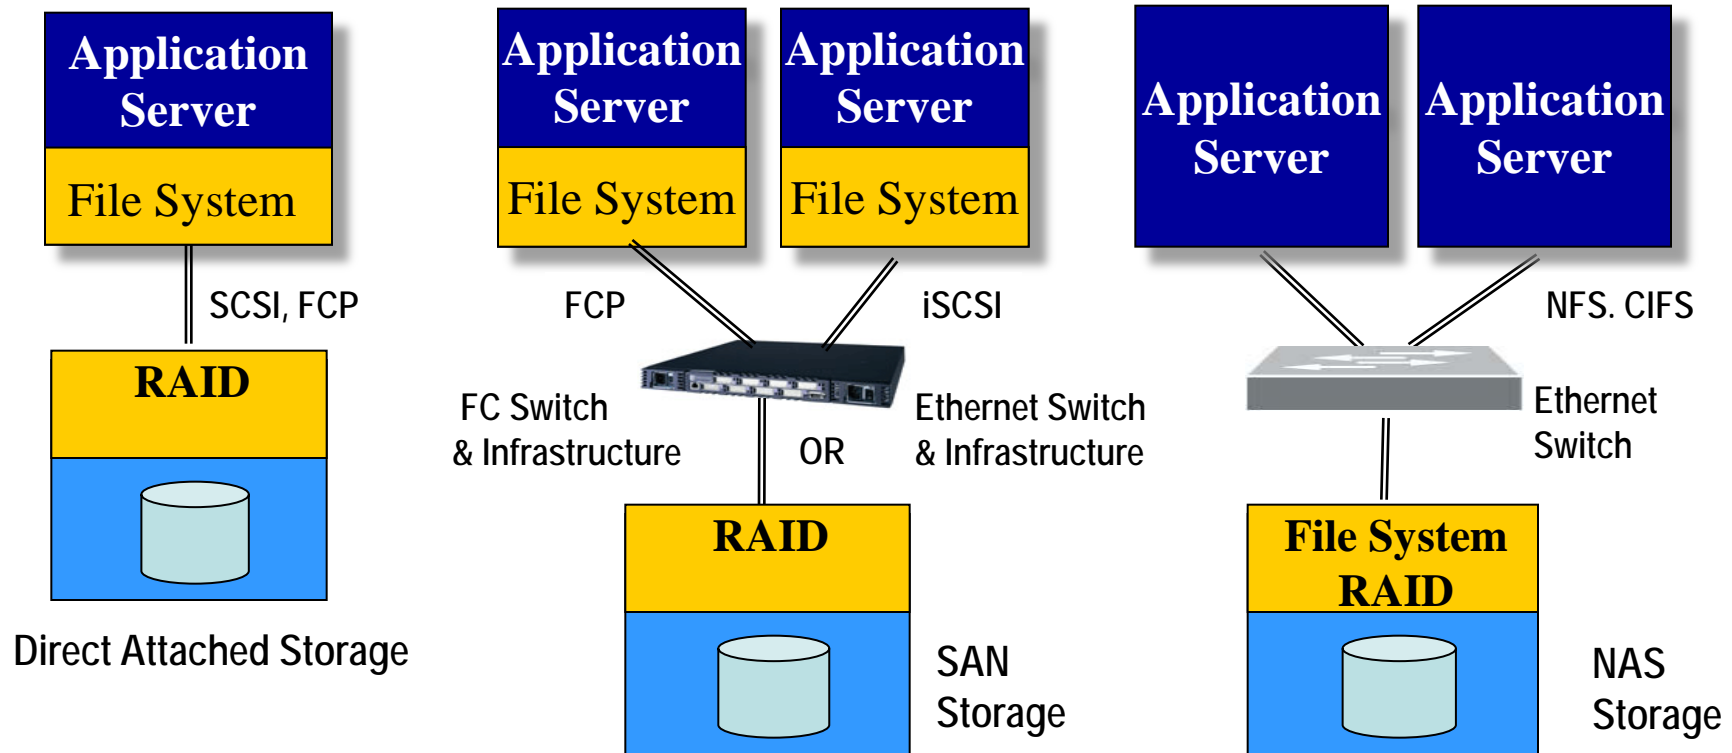
## ➤ Value of Storage Networking

- ◆ 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


# Storage Technologies Compared



- **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**



- **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

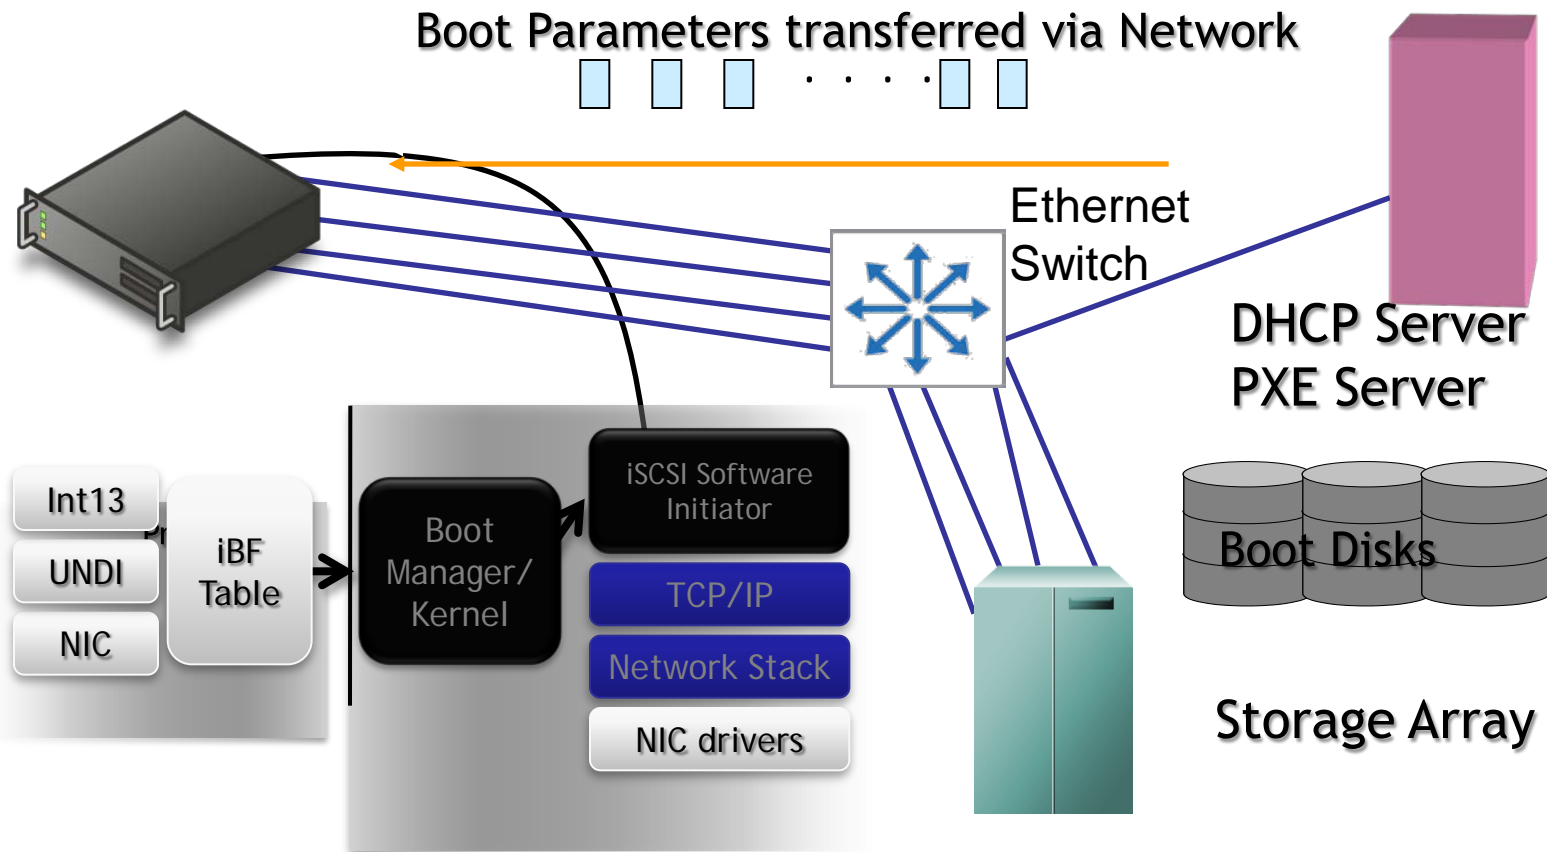


# iSCSI Host Support

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

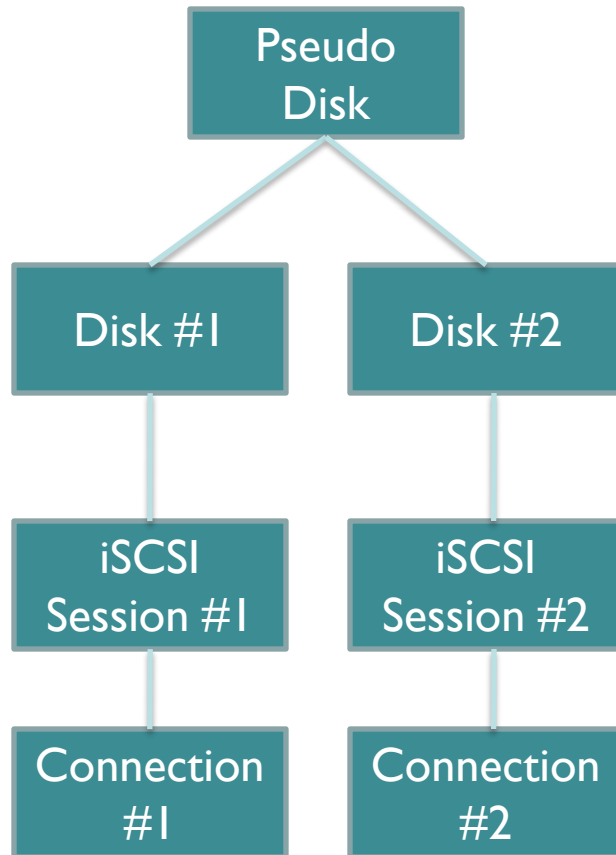
- ✓ **Simplify Server Hardware Upgrades/Repurposing**
- ✓ **Easy to configure temporary or test servers**
- ✓ **Simplify Server OS Imaging**
- ✓ **Simplify disk drive replacement**
- ✓ **Centralize storage monitoring and management**
- ✓ **Boot from iSCSI HBA or from native OS initiator**
- ✓ **Industry standard implementation of iBFT**
- ✓ **Removes spinning media from server**
  - **Important for power saving strategies and reduced maintenance**

# iSCSI Boot with S/W Initiators




# iSCSI : High Availability Options

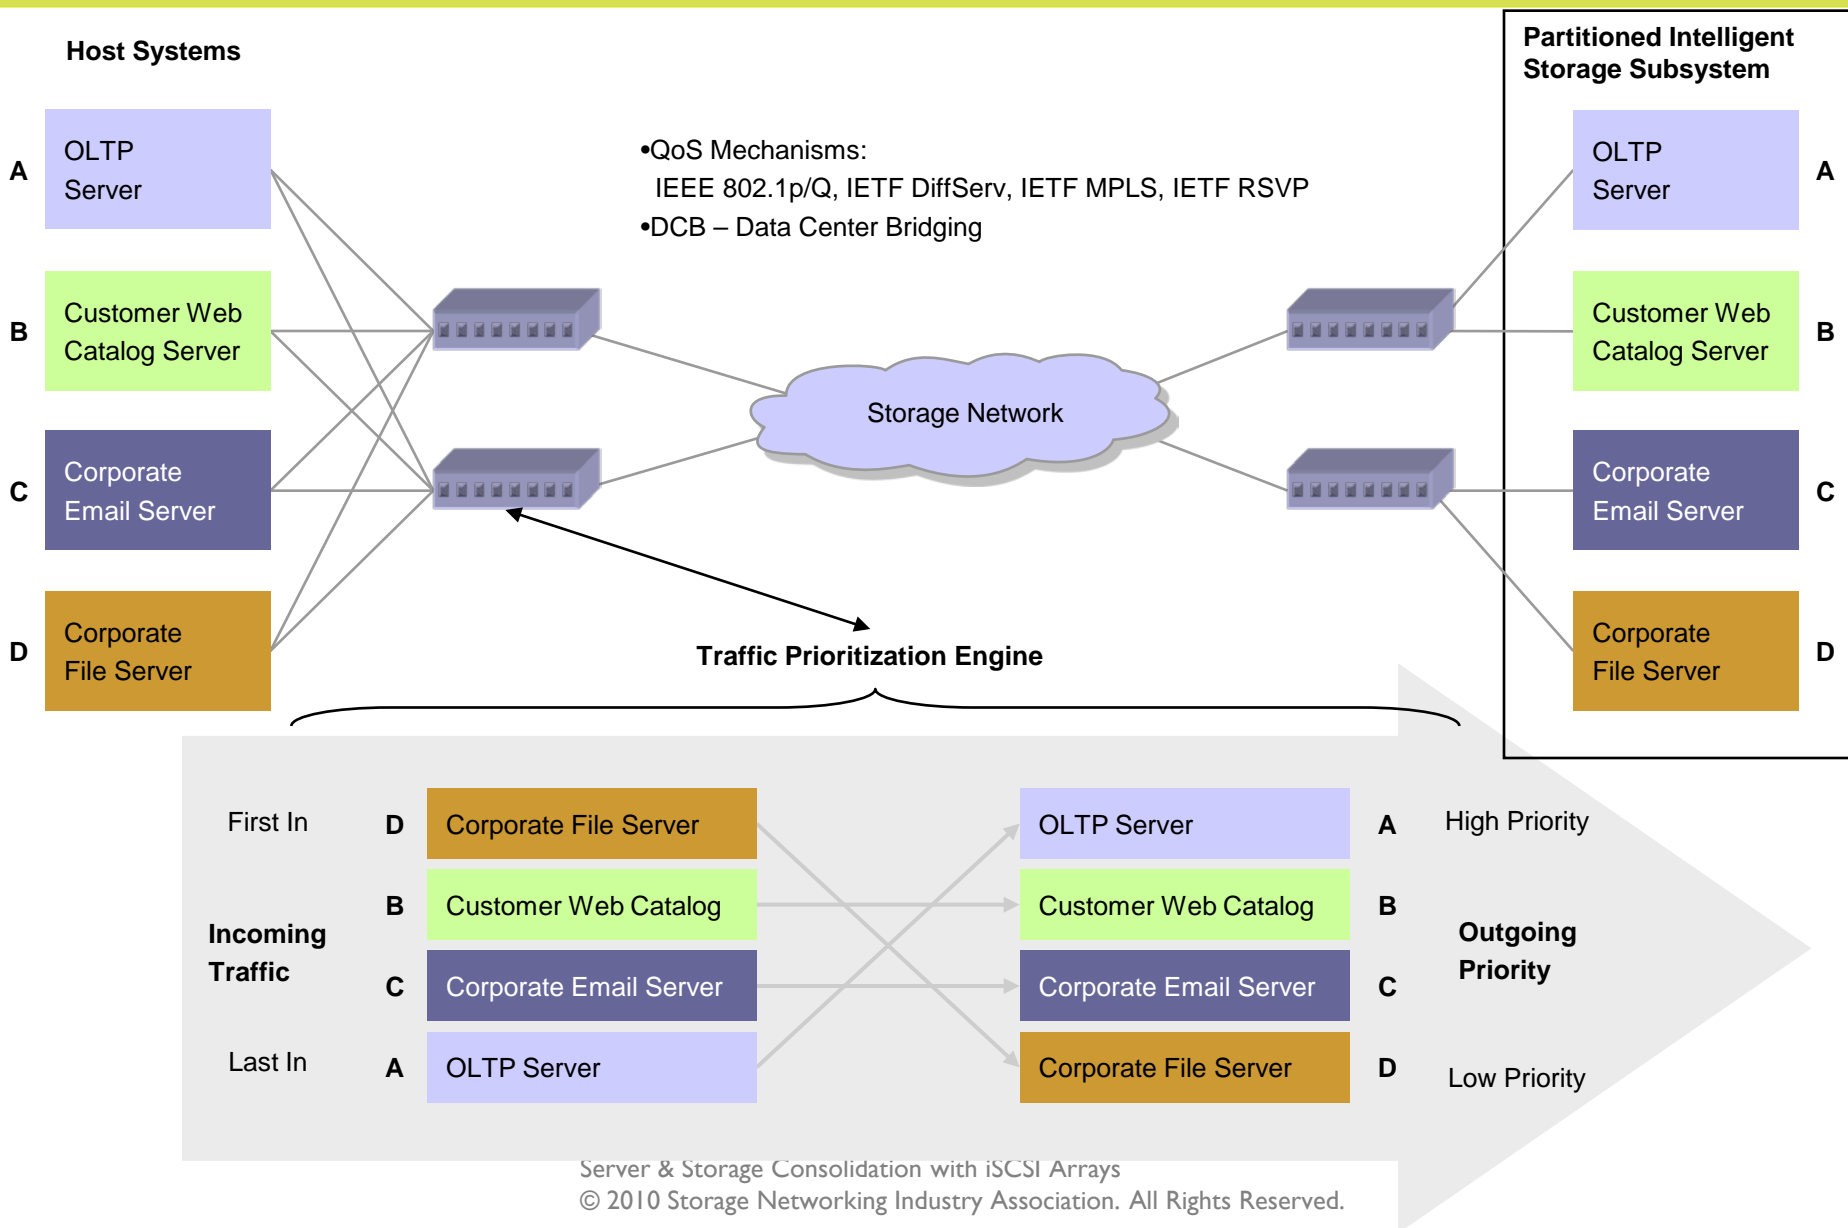
## MPIO




## MCS



# Quality of Service Policies

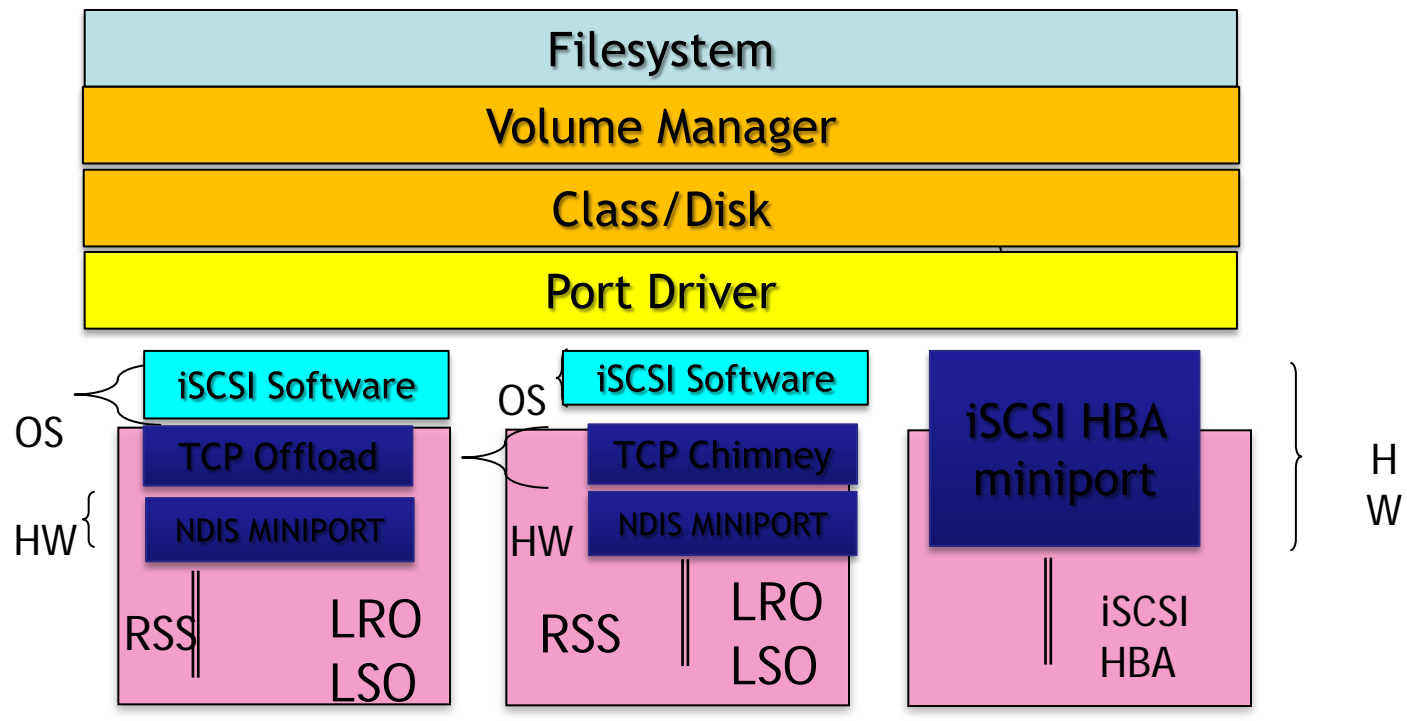


# iSCSI : High Availability



# Demystifying iSCSI Performance

- iSCSI Protocol is not limited in performance, only by underlying bus speed
  - iSCSI operates at 10Gig Wire Speed today with software initiators
  - HBAs may reduce CPU utilization for some workloads
  - Performance scales with Ethernet speeds 10 Gb, 40 Gb, and up



# 10 Gigabit Ethernet

- IEEE 802.3ae ratified 2002
- Broadly deployed in inter-switch links
- Host-side proliferation now happening
  - ◆ Affordable Price
  - ◆ Server I/O architecture support
  - ◆ Standard component on commodity servers
  - ◆ Offload built into on-board components, supported by operating systems
- Deployment/applications
  - ◆ Backbone and port aggregation for 1 Gb LANs
  - ◆ File and block storage over 10GbE



- Multi-Core changes the game
  - Multi-core Processors scale iSCSI software initiator performance
  - Performance is no longer limited by processing power of HBA engine
  - iSCSI Digest Offload directly to CPU hardware instruction set

# 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


# Server and Storage Virtualization offer VM mobility

## Scenario Description:

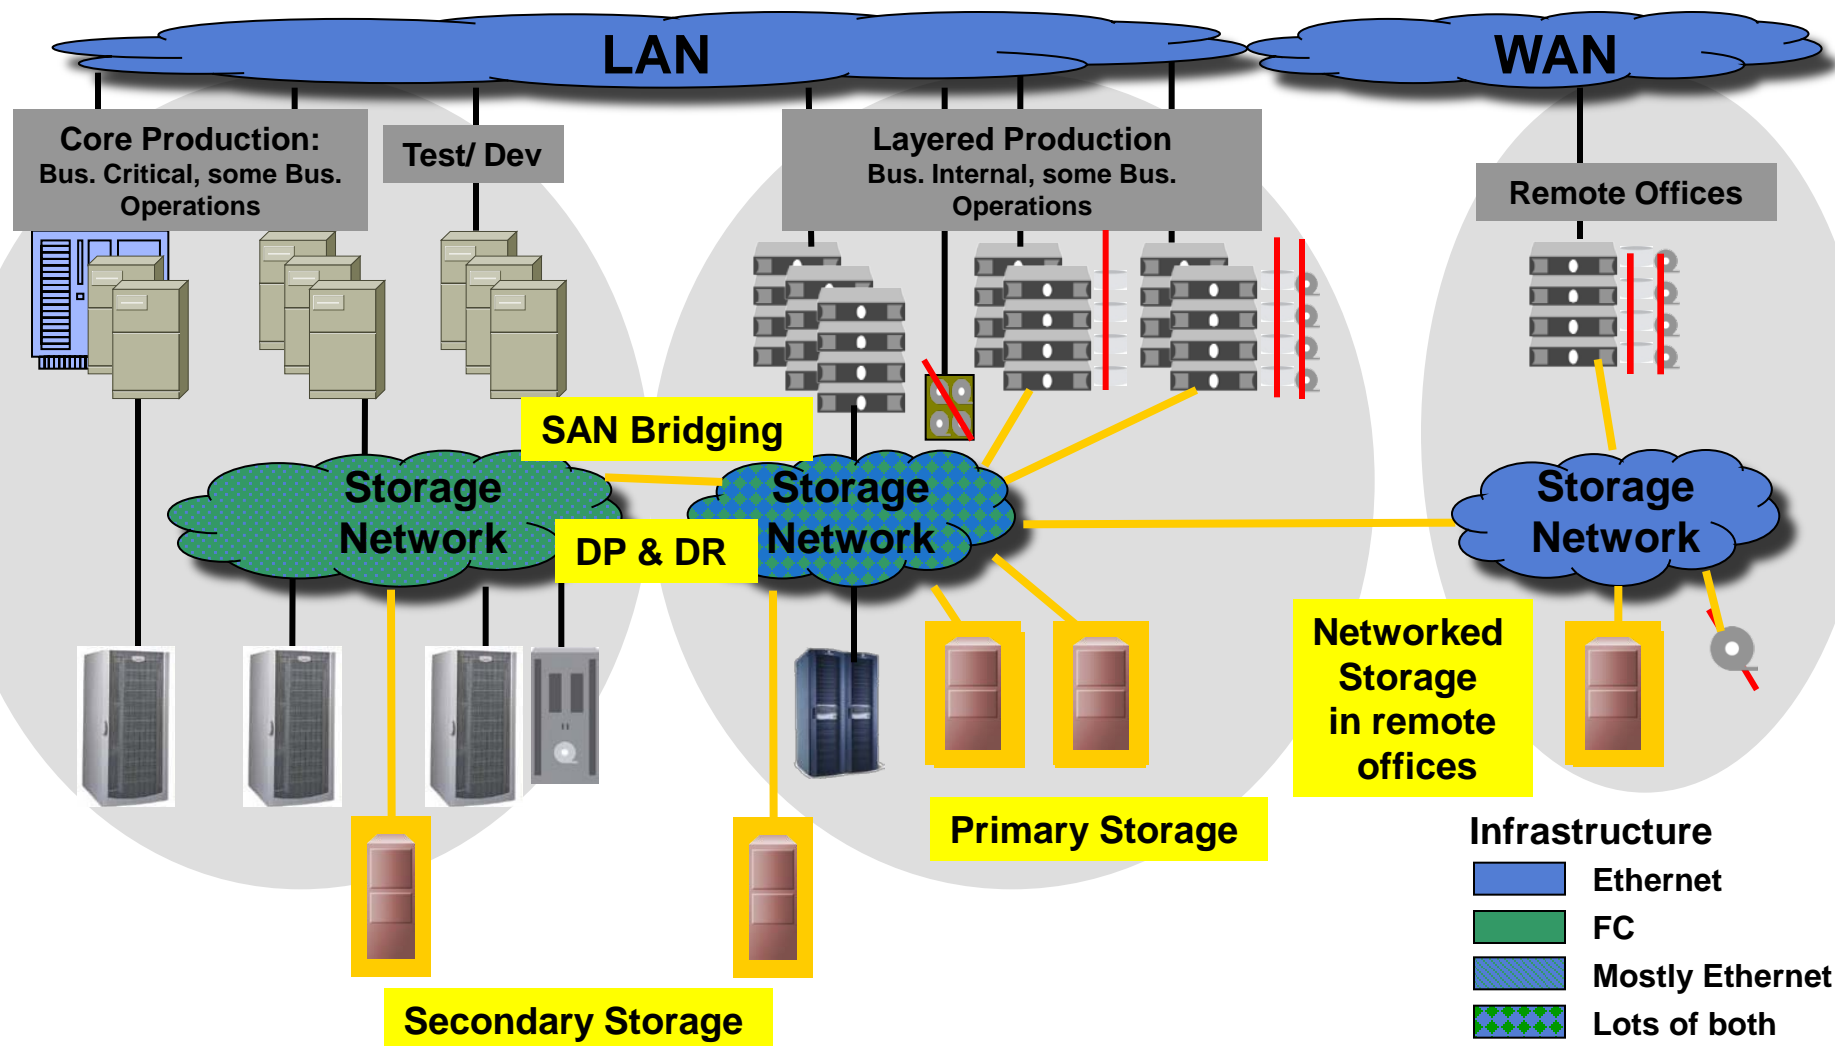
- Manage high availability with multipathing and Live migration for planned downtime situations, such as hardware and software maintenance
- Failover individual virtual machines (VMs) to other hosts within a
- Use iSCSI software initiator within the guest connected to iSCSI storage to provide guest clustering
- Nodes in cluster can be active-active

## Virtualization Benefits:

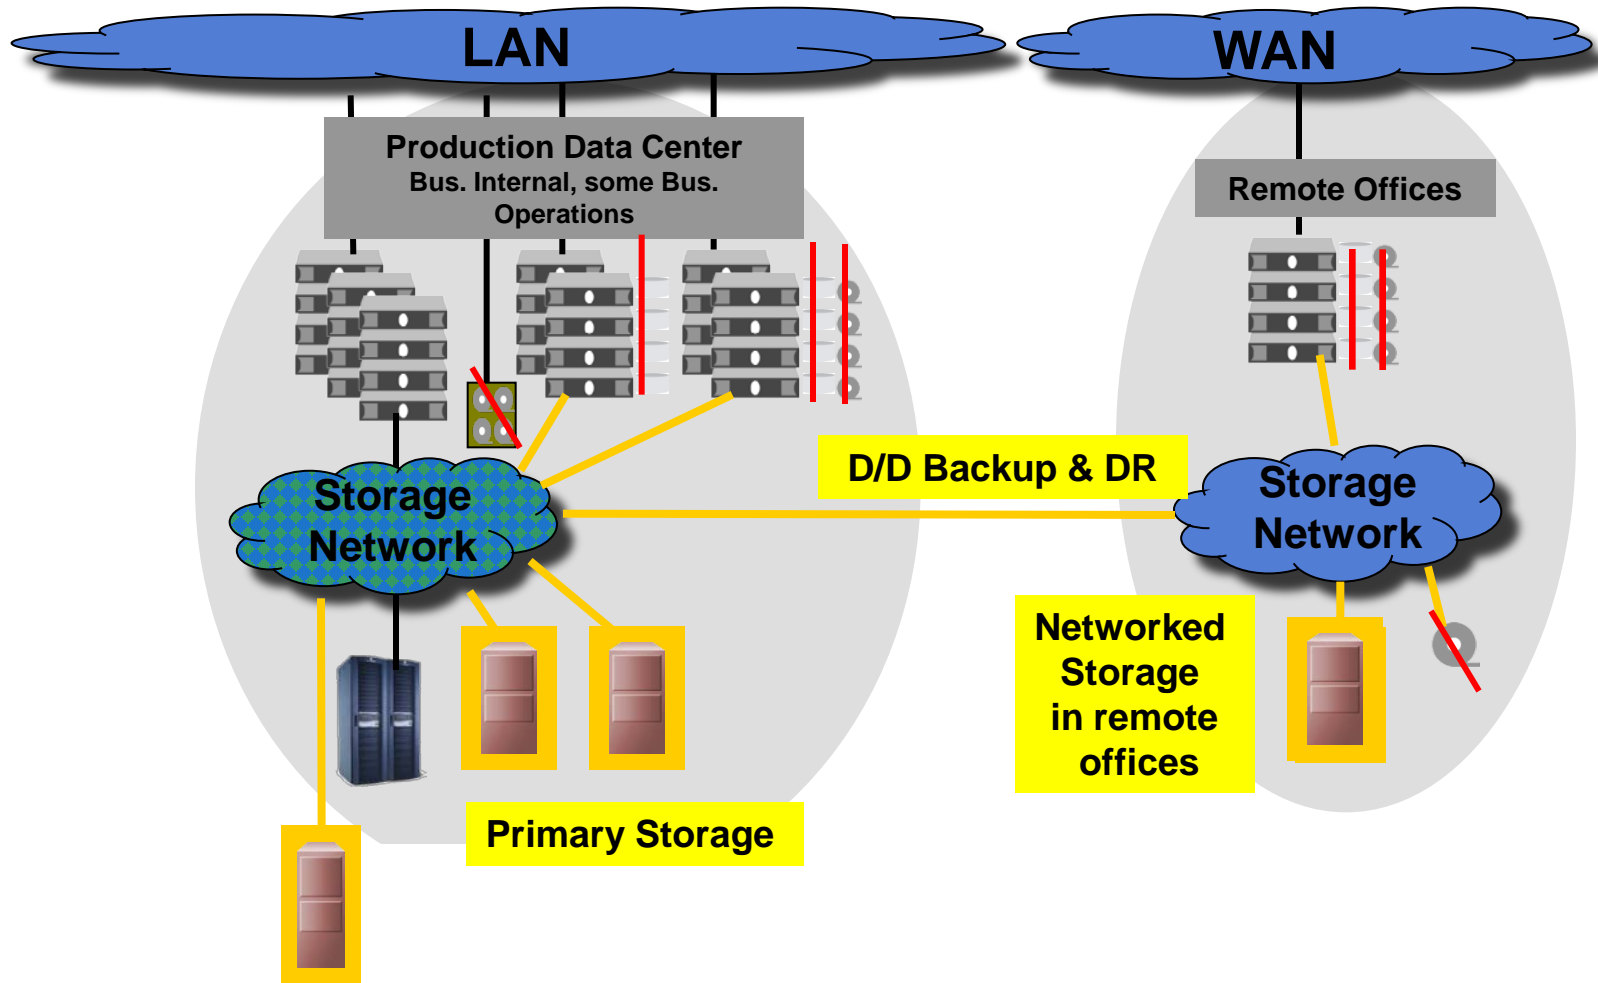
- Less downtime and loss of service for failover with Live Migration.
- Server Virtualization combined with storage virtualization
- Improve availability with less complexity
- Better server utilization due to consolidation
- Seamless backup and recovery
- Management efficiency




# Where iSCSI-based Storage Fits – Large Enterprise



# Where iSCSI-based Storage Fits – Medium/Small Enterprise



# 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



## ➤ 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.


# Futures: 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

- FCoE requires “lossless” Ethernet
  - ◆ Possible with Ethernet plus some extensions
- Data Center Bridging benefits iSCSI and FCoE
  - ◆ Ensures Storage IO receives priority over lower priority traffic
- The IEEE 802.1 DCB WG is defining these extensions
  - ◆ Priority-based Flow Control (PFC): 802.1Qbb
  - ◆ Enhanced Transmission Selection (ETS): 802.1Qaz
    - Including DCBX (DCB eXchange protocol)
  - ◆ Congestion Notification (CN): 802.1Qau
- Standards-compliant products shipping now
- 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 - iSCSI 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  
Suzanne Morgan  
Joe White  
SNIA Ethernet Storage Forum iSCSI SIG**