



Education

iSCSI SANs: Ideal Applications, Large and Small

Jason Blosil, NetApp
Gary Gumanow, Dell

- 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.


- Current IT trends in technologies such as server virtualization, are offering additional reasons to deploy shared storage. Yet, nearly 50% of the storage deployed is still DAS. This presentation will outline reasons to move to shared storage and will focus on how iSCSI offers an ideal solution for SANs both large and small. IP SANs with iSCSI offer a number of advantages to organizations including flexibility, lower cost, and simplified management. Many applications are ideally suited for iSCSI including, server virtualization, infrastructure applications, storage consolidation, and business applications.
- Learning Objectives:
 - What is iSCSI and why it the fastest growing network storage protocol on the market?
 - Identify applications well suited for IP SANs.
 - Identify key features and value proposition of iSCSI, including comparisons to DAS.

- A Case for iSCSI SANs
- Ideal Applications of iSCSI SANs

A Case for iSCSI SANs



The Countervailing Forces



IT Mandates

- Accommodate growth
 - ◆ Data
 - ◆ Additional services
- Do more with less
 - ◆ Utilization (Cap Ex)
 - ◆ Efficiency (Op Ex)
- Ensure 24x7 operations
 - ◆ Availability, DR/BC
- Secure data
- Increase agility

Source: Gartner's 25th Annual Data Center Conference "The Future of Storage" Roger Cox November 2006

Buying Storage with DAS

- Storage purchase includes server hardware with disk capacity
- Capacity requirement is a guess – so is the right server
- Typical use is for a single application
- As capacity requirements increase, so do the number of servers

“I need more storage...

...I'll buy another server.”

Server Components


- Motherboard
- Disk drive
- Network interface
- RAID controller
- I/O cable
- Fan *
- Power supply *



Single Points of Failure

* Commonly redundant in high-end server platforms

DAS Options - JBOD




➤ Advantages

- ◆ Greater expandability / higher density storage
- ◆ Redundant components at the JBOD
- ◆ Possible server cluster & failover

➤ Disadvantages

- ◆ Server clustering is complicated
- ◆ Server & I/O path still not redundant

Additional Challenges with DAS



- Limited capacity in server chassis
- Storage inefficiency - islands
- Software / firmware updates
- Backup and restore
- Site to site failover

➤ Why use DAS if...

- ◆ It's difficult to manage in large quantities
- ◆ Backup / Restore and DR can be very complicated
- ◆ Storage islands prevent redeployment or reallocation – inefficient space utilization
- ◆ Capacity expansion in increments of a full disk
- ◆ Server virtualization features are limited

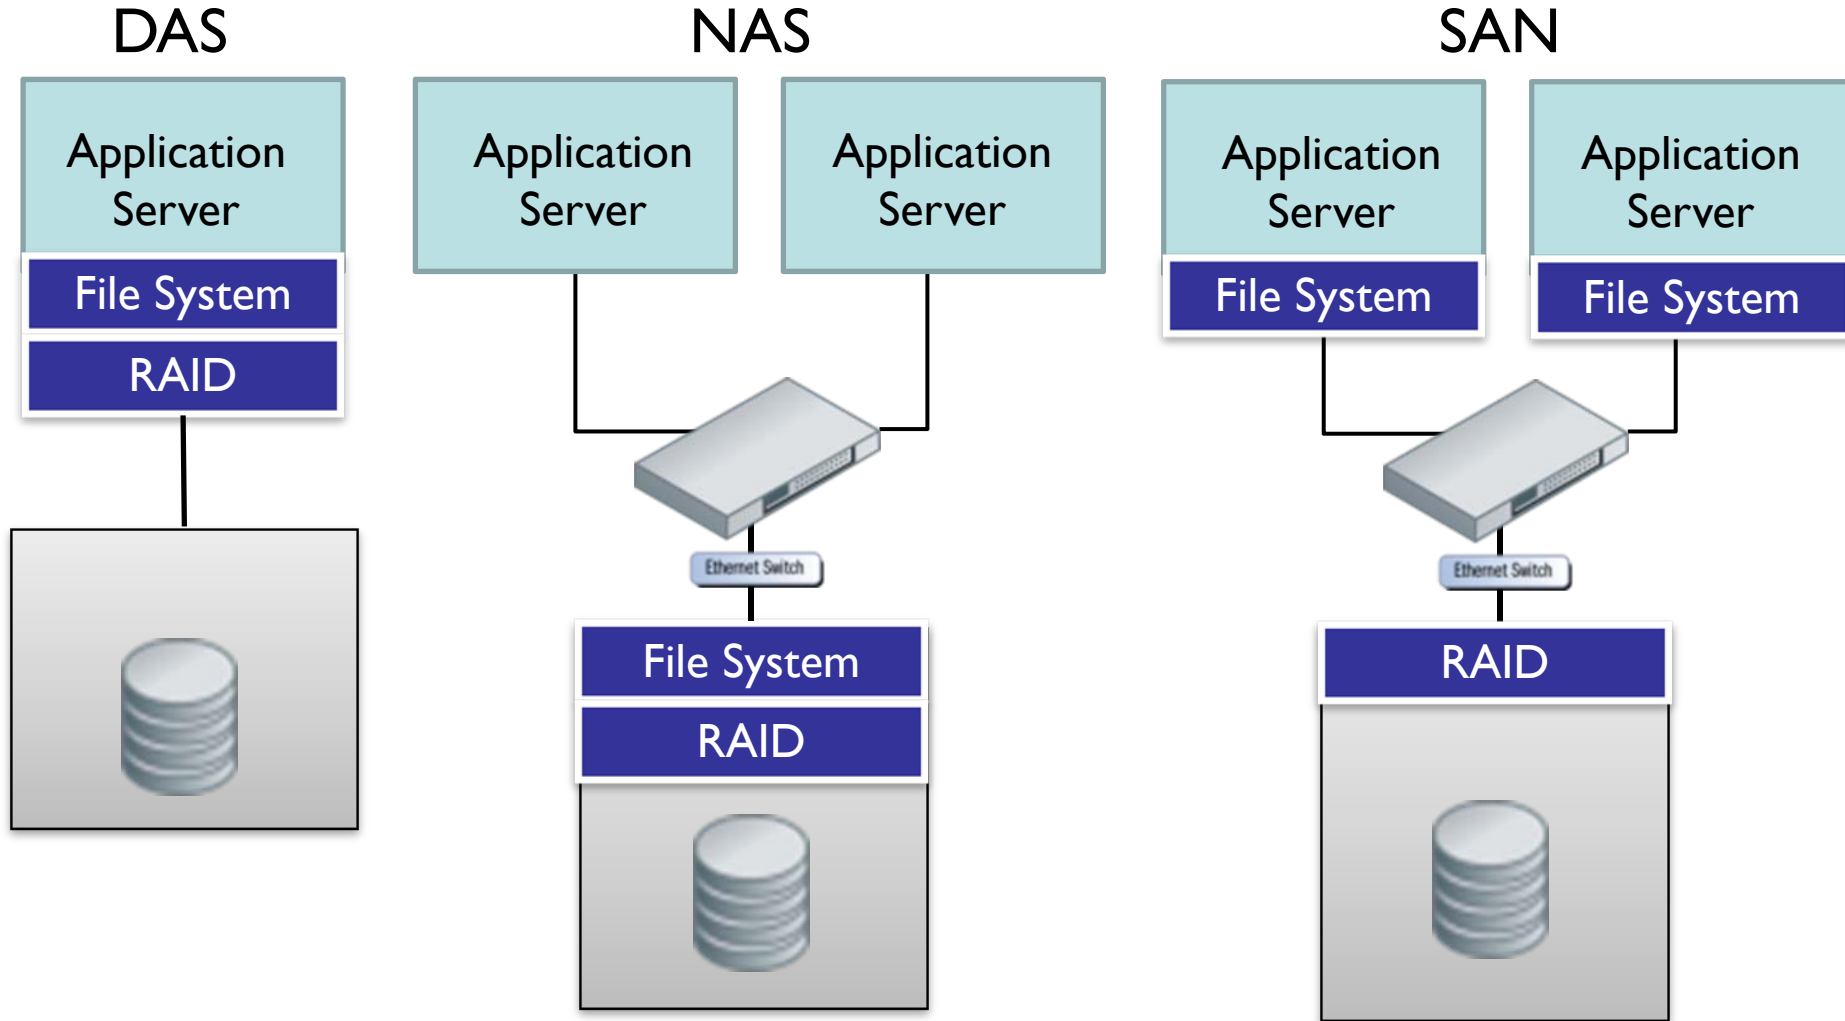
➤ Why we use DAS...

- ◆ Inexpensive to implement
- ◆ Easy to deploy in small quantities
- ◆ Budgets are often departmentalized

In 2008, >49% of
deployed disk storage
was DAS*


* IDC, June 2009

Storage Comparison




Market Share – by Network Protocol

Revenue Share



Source: IDC Storage Tracker, Dec 2009

Capacity Share



➤ Cost

- ◆ All of that redundancy comes with a price

➤ Complexity

- ◆ Difficult to configure and service


➤ Performance

- ◆ iSCSI SANs don't perform well

iSCSI Addresses These Obstacles

What is an iSCSI SAN?

- **Server Adapters**
 - ◆ Standard Ethernet NICs
- **Cabling**
 - ◆ Standard Ethernet cabling
- **Switches**
 - ◆ Gigabit or 10 Gigabit standard switches
- **Storage**
 - ◆ iSCSI Storage Array
- **Skill Set**
 - ◆ Leverages Existing IP Expertise



Options to Increase Storage Capacity

1. Storage Server - \$8,000 est. (tower) / \$13,000 est. (rack)

- ◆ 2 socket motherboard w/ 8GB memory each
- ◆ Redundant power supplies
- ◆ RAID Controller / cables / disks
- ◆ Network Adapter (4 port)
- ◆ OS license and Backup software

2. JBOD - \$7,000 est.

- ◆ Redundant power
- ◆ Disks

3. iSCSI Storage System - \$13,000 est.

- ◆ Redundant controllers & power
- ◆ Disks

Near price
parity...

...Increased
performance,
flexibility,
reliability

DAS vs iSCSI SAN - Configuration


DAS

- Install RAID Driver
- Create drive partition (select RAID level)
- Format partition




iSCSI SAN

- Install iSCSI Driver
- Configure IP addresses on target storage
- Create LUN on storage
- Map drive letter to storage LUN IP address
- Format LUN



iSCSI OS Support

OS	Initiator	Boot
Microsoft	Native S/W, 3rd Party HBA	HBA, S/W
Red Hat	Native S/W, 3rd Party HBA	HBA, S/W
Oracle	Native S/W, 3rd Party HBA	HBA, S/W
SUSE	Native S/W, 3rd Party HBA	HBA, S/W
IBM (AIX)	Native S/W, Native HBA	HBA, S/W
Sun (Solaris)	Native S/W, 3rd Party HBA	HBA
HP (HPUX)	Native S/W	No
VMware	Native S/W, 3rd Party HBA	HBA, S/W
Novell (Netware)	Native S/W	No



➤ Modular design

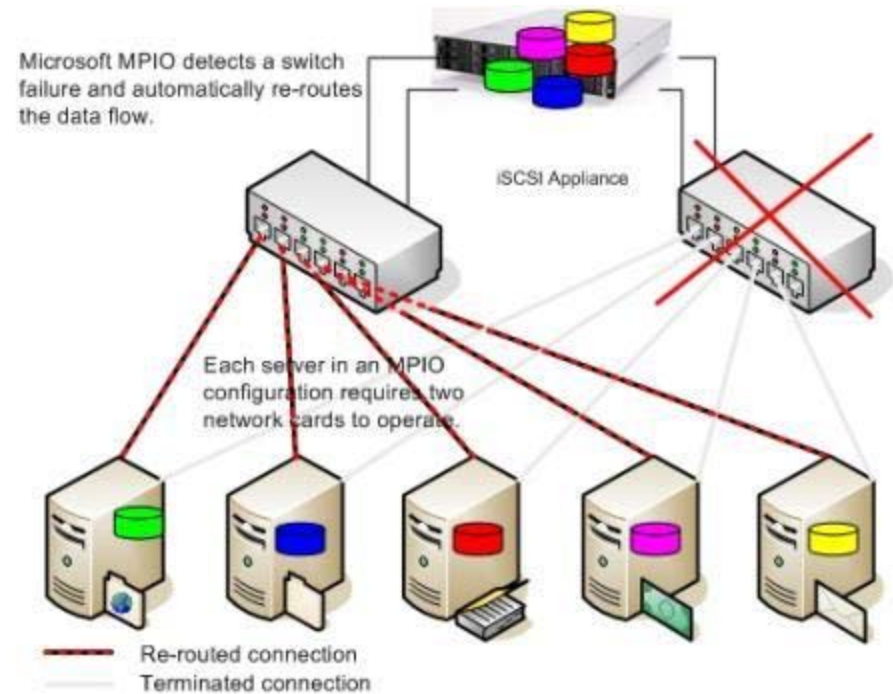
- ◆ Redundant hot-swap controllers
- ◆ Redundant hot-swap power & cooling units
- ◆ Expandable I/O ports

➤ Software updates are non-disruptive

➤ No active components on passive midplane


iSCSI SAN Built for Reliability and Accessibility

- Redundant components
 - ◆ No single point of failure
- Improved data protection
 - ◆ Advanced RAID protection
 - ◆ Advanced software features such as:
 - › High availability failover
 - › Snapshot backup
 - › Server-less backup
 - › Site to site backup
- Multipathing (MPIO)
 - ◆ Fault-tolerance
 - ◆ Load balancing



iSCSI SAN Offers Increased Performance & Scalability

- Service multiple applications with single storage pool
- More spindles for better IOP performance
- High performance networks with multiple paths
 - ◆ Port bonding for IP traffic – increase overall bandwidth
- Easily add capacity online
- Capacity allocation and resizing of LUNs in increments less than a drive



➤ Software Initiator + Standard NIC

- ◆ Software initiator runs on host CPU
- ◆ Low cost (Free download)
- ◆ May offer highest performance, but highest CPU overhead
- ◆ About 85% of iSCSI deployments


➤ TCP/IP Offload Adapter

- ◆ Lower host CPU overhead
- ◆ Added cost for adapter, but uses OS iSCSI initiator

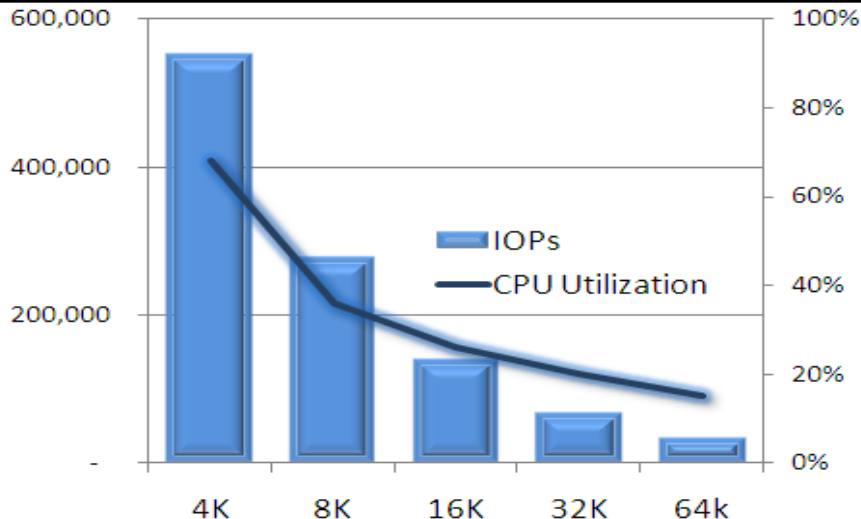
➤ iSCSI HBA

- ◆ Ships with proprietary initiator
- ◆ Highest cost solution

Read/Write IOPs and Throughput Test



Read/Write IOPs and CPU Test



1,030,000 IOPs

Single Port

10GbE line rate

10k IOPs per CPU point


Performance for real world apps

Future ready: Performance Scales

552k IOPs at 4k represents


- 3,100 Hard Disk Drives
- 400x a demanding database workload
- 1.7m Exchange mailboxes
- 9x transactions of large eTailers
- Jumbo frames: >30% CPU decrease is common for larger IO size (jumbo frames not used here)

Ideal Applications for iSCSI SANs




iSCSI SANs for Storage Consolidation

Current Environment



- Inefficient storage utilization
- Management complexity
- Unsatisfactory data availability

Consolidation Solution



- Increased data availability - failover
- Simplified data management
- Efficient storage utilization
- Highly scalable

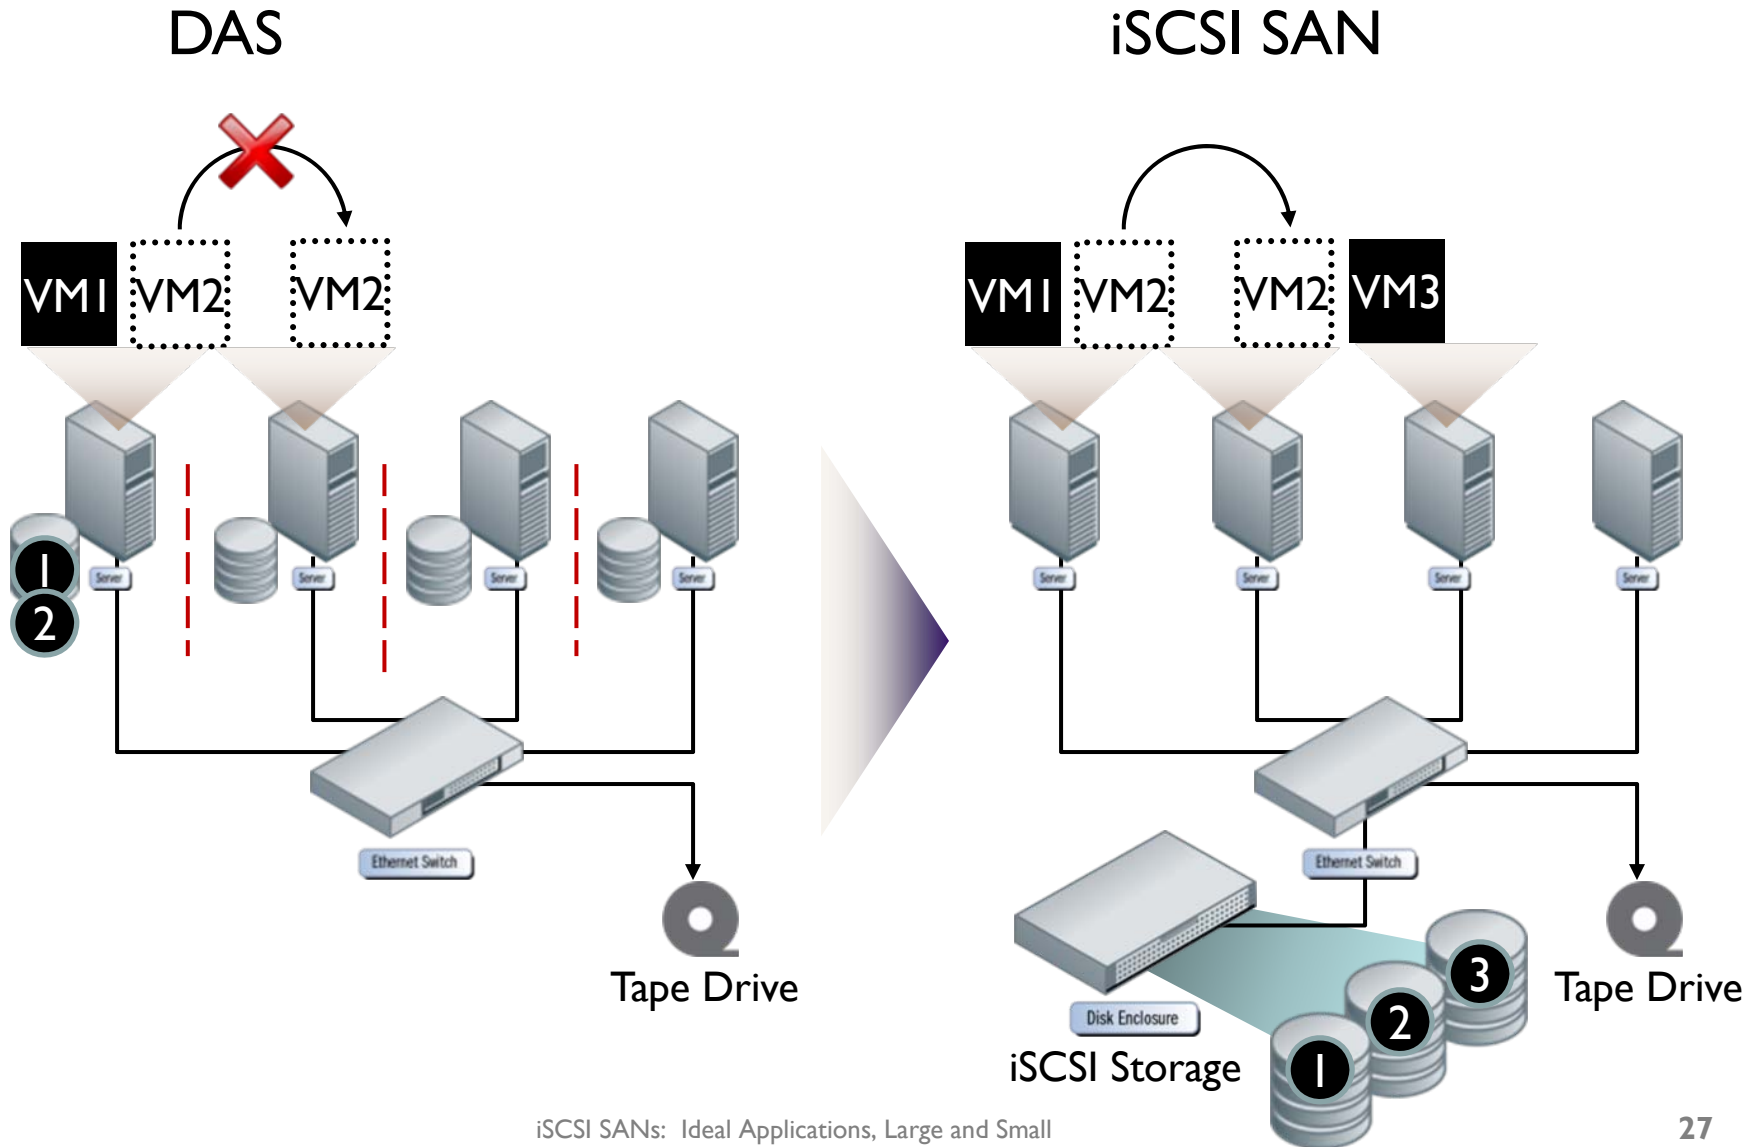
➤ Why iSCSI for Blades

- ◆ Economics of onboard or backplane Ethernet
- ◆ Leverage infrastructure as a Resource Pool
- ◆ Unified storage and data network
- ◆ Server admins familiar with Ethernet
- ◆ Affordable hardware
- ◆ Lower TCO


➤ Why iSCSI for Virtualization:

- ◆ Lower cost solution consistent with economic benefits of virtualization
- ◆ Server administrators are primary implementers of virtualization – familiar with Ethernet
- ◆ Virtual IP addressing aids mobility and scaling

iSCSI SANs Enable Virtualization




iSCSI SAN – Network Boot w/ Rapid Cloning




- **Rapid server deployment**
 - ◆ Clones created in seconds / minutes
 - ◆ Easy to scale
- **Improved management**
 - ◆ Simplified OS updates
 - ◆ Boot volumes protected
- **Lowers server costs**
 - ◆ No low use local disks
- **Improved disk performance and efficiency**
 - ◆ Boot volumes spread across shared disks


iSCSI SANs for Backup and Restore **SNIA**



- Reduce complexity and increase efficiency
 - ◆ Important in virtual server environments
- Single point backup
- Reduces network congestion to host
 - ◆ Snapshot and disk to disk replication

iSCSI SANs for Disaster Recovery





- Historically, IT infrastructure is an application-centric silo
- Virtualization technology enables a shared and service-oriented infrastructure
- Service providers can now offer a lower-cost business model
- All three models are likely to coexist at the enterprise data center

Ethernet / iSCSI is ideal technology for Cloud

IT as a Service (ITaaS)

IaaS

Infrastructure
as a service

IT Services:

- Servers
- Network
- Storage
- Management
- Reporting

Examples:

BT
Telstra
T-Systems (ITaaS)

“PaaS”

Platform
as a service

Application
building
blocks and
standards

Examples:

Amazon EC2
Force.com
Navitaire

“SaaS”

Software
as a service

Applications

Examples:

Yahoo! E-mail
SalesForce.com
Google apps

“StaaS”

Storage
as a service

Storage

Services:

- Primary
- Backup
- Archive
- DR

Examples:

Amazon S3
Nirvanix



What Drives Cloud Adoption

Business Benefits

- Turn capex into opex
- Faster business innovation
- Risk sharing with vendors
- Increased productivity

Enterprise IT Benefits

- Elastic scalability
- Pay-as-you-go efficiency
- Data access any time, any where
- Predictable cost structure
- Operational efficiency



Scalable



Cost efficient



Flexible

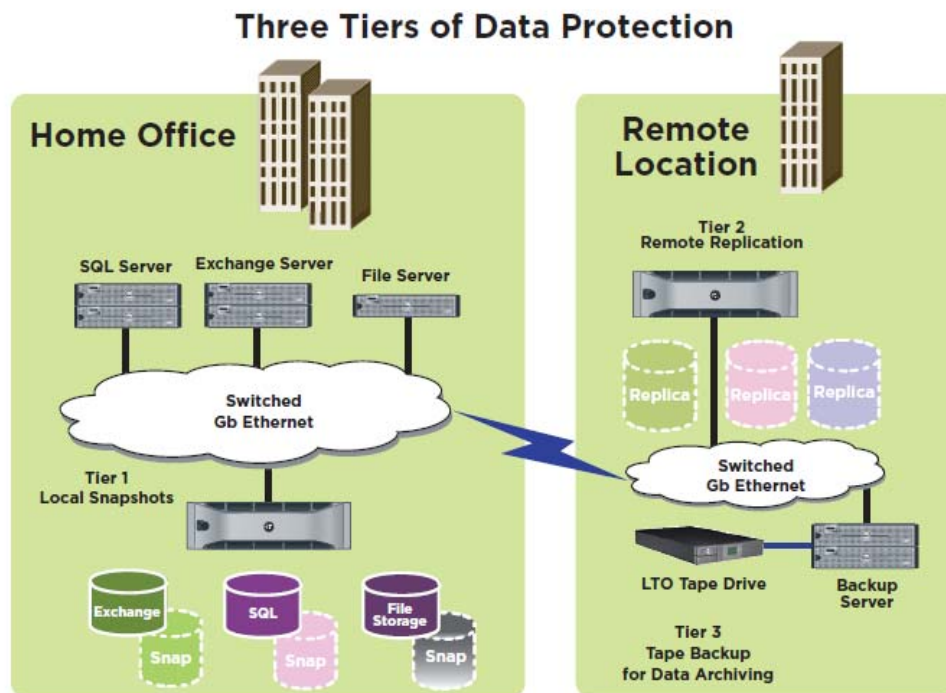
Case Study: Network Boot

- Test & development lab with over 1500 blade servers operating near 100% utilization
- Over 1200 blades boot off iSCSI SAN
- Can reconfigure all or subset of grid in minutes with host-specific personalization
- Native SAN boot OS: Red Hat Linux, SUSE Linux, Windows, VMware ESX, Solaris 10
- Lab needs only 252 disk drives instead of typical 3000 disks to boot 1500 blade servers



Engineering Data Center Lab

Case Study: Disaster Recovery



Tier 1: Local Snapshots

Tier 2: Remote Replication

Tier 3: Tape Backup for Data Archiving

Asset Management Company

Requirements

◆ Resilience

- ◆ “No human single point of failure”
- ◆ Component failure / multiple scenarios
- ◆ Detailed risk analysis; servers, network and applications

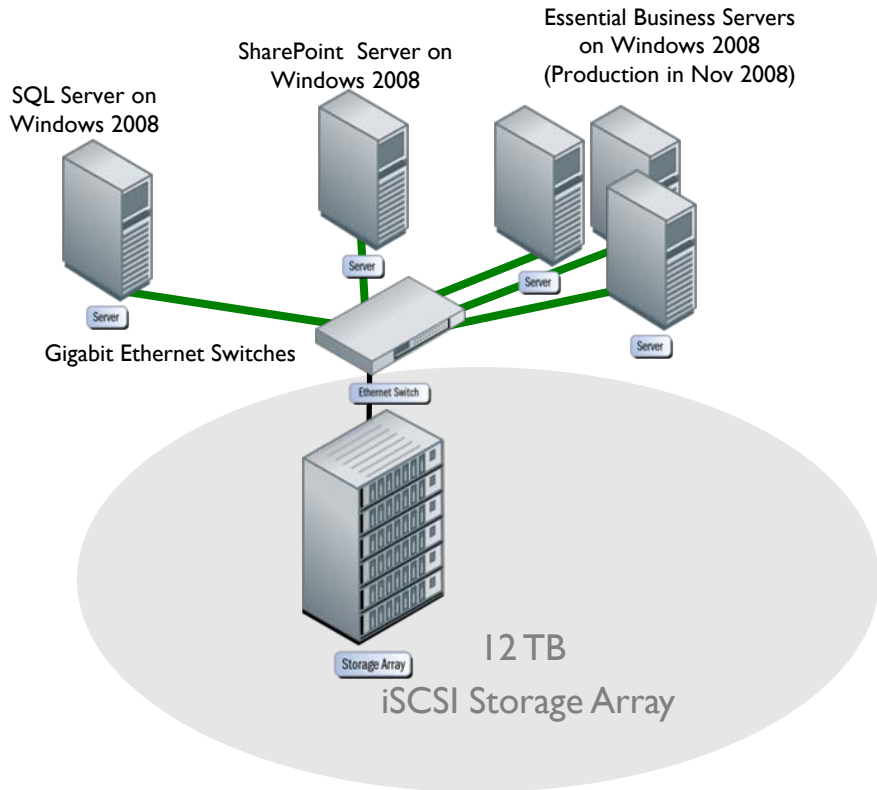
◆ Automation

- ◆ Recovery procedures
- ◆ Relevance
- ◆ Declaration must be manual

◆ Recovery

- ◆ 95% of DR servers are virtual. DR relevance is automated when combined with iSCSI
- ◆ Can provision new, fully patched server in less than 20 minutes (from request to completion)
- ◆ Physical capacity expansion is very simple
- ◆ Expanded file server volume in less than 5 minutes with ZERO impact to users (SAN and server volume)

Case Study: Windows Consolidation



“The [iSCSI] SAN is easy to setup and manage; and it scales effortlessly. I can quickly backup and restore my applications using snapshots.” Logistic and IT Manager

- Applications:
 - ◆ Current: SQL Server, SharePoint Server, Essential Business Server BETA, WennSoft Project Managers Portal, Microsoft Dynamics Great Plains
 - ◆ Planned: Exchange Server 2007
- Pain Points:
 - ◆ DAS storage difficult to scale and manage with growing need to keep more data online
 - ◆ Insufficient backup window and restores from tape takes too long
 - ◆ FC SAN too complex and costly
- Solution:
 - ◆ Updated to iSCSI SAN with Gigabit NICs and software initiator
 - ◆ Consolidated storage onto a 12-TB iSCSI SAN with array-based snapshots
- Benefits:
 - ◆ Simplified server and storage management with boot from SAN and transparent SAN expansion
 - ◆ Fast backup and restore using SAN-based snapshots

Large European telecom operates information and communications technology for multinational corporations and public institutions.

- Since 2004: flexible IT over the cloud
- SAP and other applications “as a service” to multinational enterprise customers
- More than 300 large customers now using their dynamic IT services
 - ◆ One petroleum enterprise outsourced **all** global IT services
- Storage infrastructure deployed with IP SAN (iSCSI)

Other Applications

- **Accessibility of stranded servers**
 - ◆ Lower workload servers without FC access but Ethernet onboard
- **Web services**
 - ◆ ftp
 - ◆ http
- **Infrastructure services**
 - ◆ Active Directory
 - ◆ DNS and iSNS

What Lies Ahead



- Robust roadmap
 - ◆ 10Gb today, 40Gb, 100Gb future

- Enhanced Ethernet with Data Center Bridging (DCB)
 - ◆ Improved QOS

- IPv6

➤ iSCSI SANs Offer:

- ◆ Easy to use / cost effective storage network
- ◆ Supported by all major operating systems
- ◆ Shared storage pool improves storage efficiency
- ◆ Redundancy improves data availability
- ◆ Flexibility and scalability for server virtualization
- ◆ Dynamic network deployments for Cloud Services

Questions?



- Please send any questions or comments on this presentation to SNIA: trackstoragemgmt@snia.org

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

- SNIA Education Committee

**Jason Blossil
Gary Gumanow
Jordan Plawner**