Creating a High Availability, Zero-Downtime, Storage Area Network Environment

Daniel Cohen
Solutioneer
Brocade Communications

BROCADE

Agenda

- **HP and Brocade – SAN Infrastructure for the Enterprise**
  - Brocade Corporate Overview
  - SAN Benefits
  - Why a Brocade Based HP SAN
  - Brocade Fabric Architecture
  - HP Integrated SAN Solutions
- **Designing a SAN for High Availability**
- **SAN Solutions**
- **Case Studies**
  - **A leading internet infrastructure provider**
  - **Large telecommunications company**
  - **Large North American hospital**
  - **A leading Storage Service Provider (SSP)**

HP and Brocade –SAN Infrastructure for the Enterprise

BROCADE

Brocade Corporate Overview

- San Jose, CA based company founded in 1995
- Brocade has ~90% share of the Fibre Channel switch market*
- Fortune 500 IT managers rated Brocade as one of the top 3 SAN companies
- Comprehensive switch product line - Entry level to Enterprise class switches utilizing a networking model for SAN infrastructure
- Industry exclusive Fabric OS
- Fabric Access Layer (API) integration with SAN management tools
- ~850 employees worldwide
- Major contributor to authoring and editing Fibre Channel standards

*IDC, March 7, 2000

5 Major SAN Benefits

- Improved Availability of Mission-Critical Applications
- More Effective Storage Management
- Improved Storage Utilization
- Improved Availability of Enterprise Information
- Foundation for Disaster Tolerance

Why a Brocade Based HP SAN?

Brocade SANs are the foundation for “always on“ infrastructure

- SilkWorm switch family is the foundation for 99.999% availability
- FabricOS based value added features
- Compatibility and investment protection for switch products

SAN management is simplified – Single Pane of Glass

- Brocade management interfaces integrated into HP Fsam roadmap
- HP Open View products
- Integration with Fabric Access Layer (API)

Brocade networking model for SAN deployment

- Enables shared resources
- SAN infrastructure enables SAN solution deployment
- Highly available and easily scalable
- Auto configuration and self healing

Support for Brocade SAN Fabrics in HP-UX

BROCADE Product Architecture

Value-Added Fabric Software

- **Zoning**
- **WebTools**
  - Fabric Status
  - Events
  - Reports
  - Smart GBICs
- **QuickLoop**
  - Multiple
  - Zonable
  - **Fabric Assist**
- **Fabric Watch**
- **Remote Switch**
  - **Extended Fabrics**
- **Fabric Manager**
  - ***NEW***

API

Application Programming Interface

Base Services

Fabric Services
Management Server
Simple Name Server
Alias Server (multicast)
Cascading

Management Services
SNMP
Telnet
Front panel

Hardware Platform

Silkworm Family ⟷ 6400
2400/2800
20x0 / 22x0

Entry Level to Enterprise Switches

Large scale Backbone
- 128 ports
- Bladed
- 2GB GBICs
- SAN backbone

Integrated Fabric
- 64 ports
- Switch FRU,
- GBICs
- SAN Core

Scalable Solution
- 8 or 16 ports
- FRU power, fan, GBICs
- Build large SANs

Entry Solutions
- 8 or 16 ports, low cost
- GBIC FRU only
- SAN Appliances
- HUBs / Small Fabrics
- Workgroup Fabrics

Fabric Services

BROCADE – Simple Name Server

- Standards based network address assignment
- Distributed across the fabric
- Dynamically scalability
- Automatic legacy device support

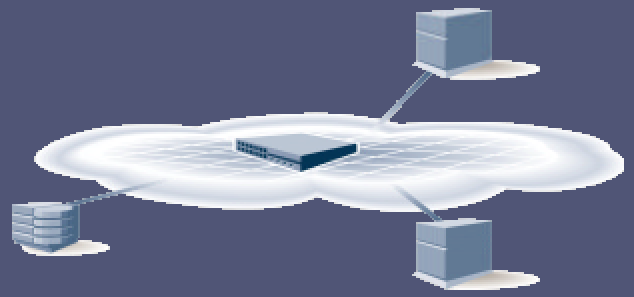

BROCADE Scalability - Cascading

- Any port to any port
- Dynamic recovery / self-healing
- Dynamically configured routes (FSPF)
- Automatic bandwidth scaling
  - Multiple Inter-Switch Links (E_Port)

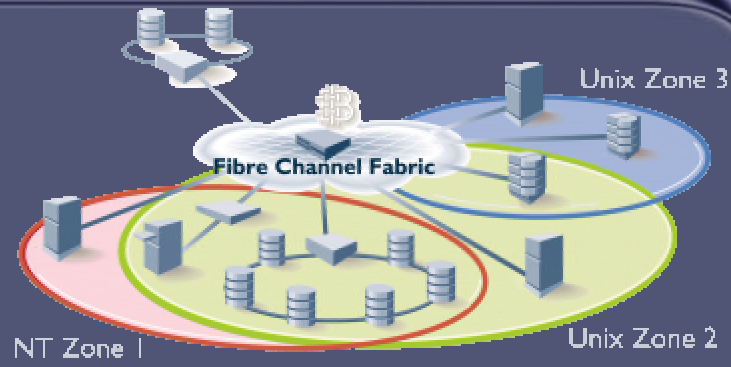

BROCADE Zoning

- One-time configuration
- Updates distributed dynamically
- Port or WWN
- Overlapping zones / temporary zones
- No logical limit on the # of zones

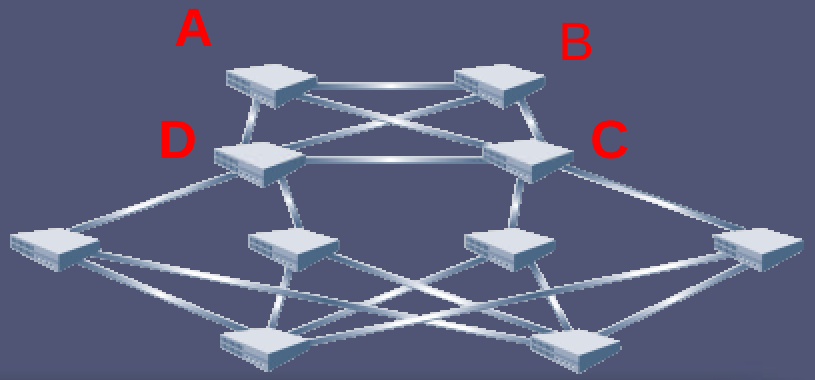

FabricOS Value-Added Software

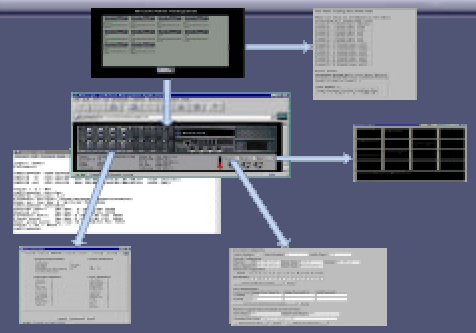

WEB TOOLS

Zoning

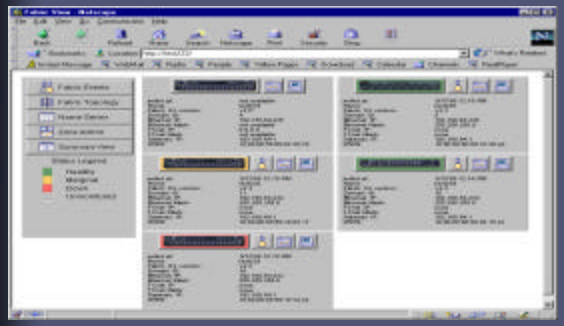

FabricWatch

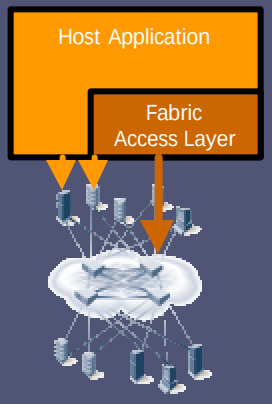

Fabric Access Layer

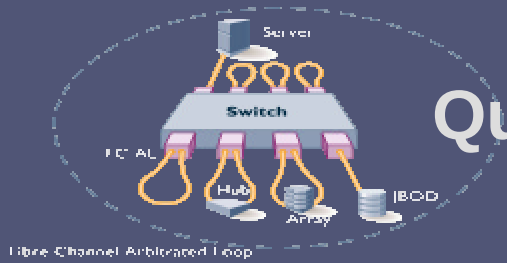

QuickLoop

Remote Switch

Extended Fabrics

Zoning

- Zone by port with zone access controlled by the ASIC
- Zone by WWN
- Zone by QuickLoop
- Overlap zones
- Unlimited zones
- Zones enable fault Isolation

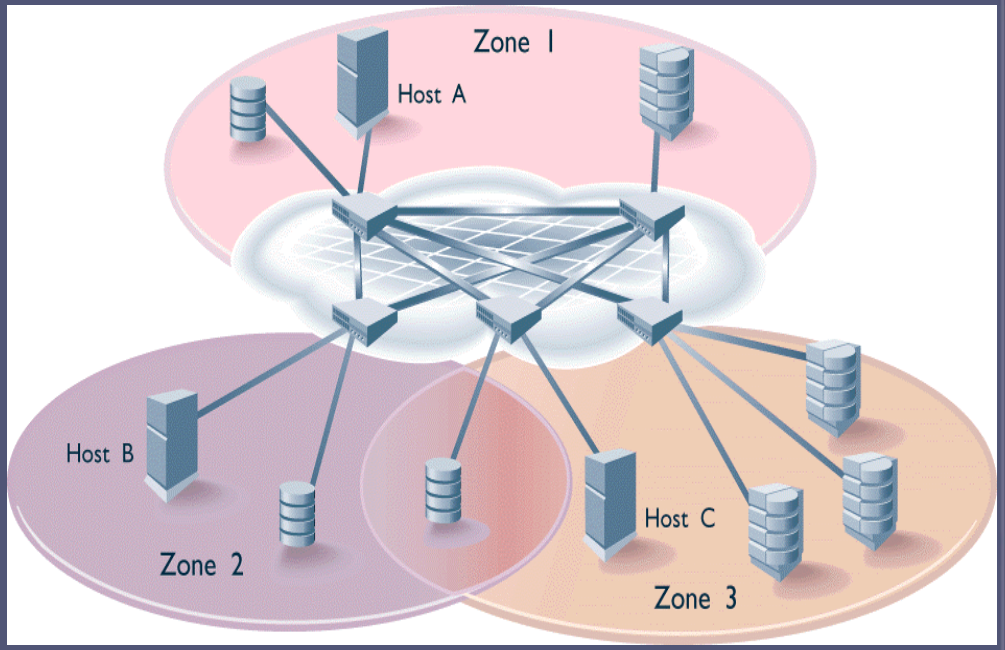

Arbitrated Loop and Fabric Assist

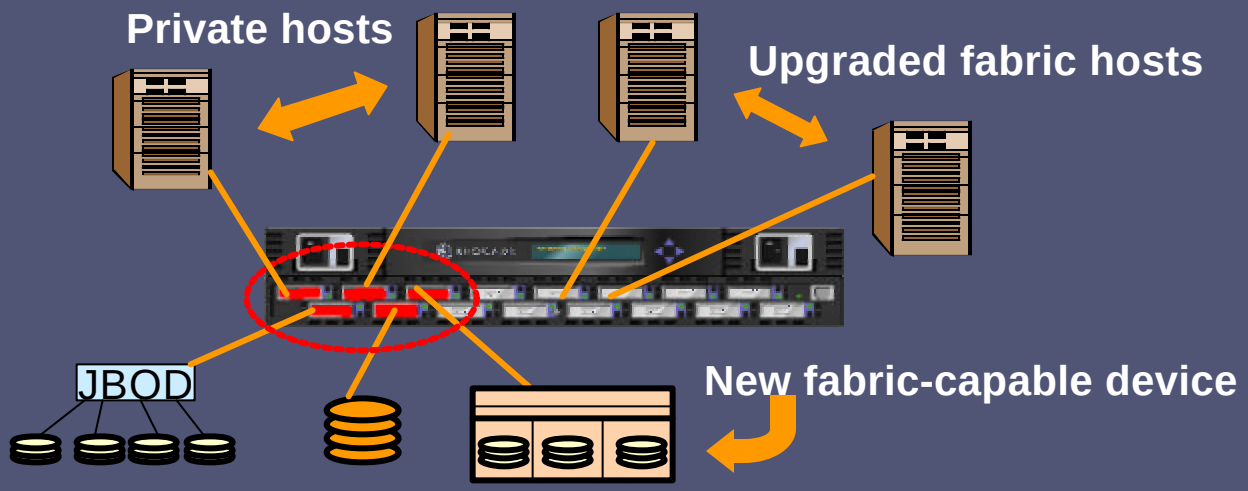

Brocade QuickLoop/Fabric Assist

- Public hosts can access private storage (Translative Mode)
- Private hosts can access public storage (Fabric Assist)
- Attach private loop devices
- Failing devices and the LIP are confined to zoned loops
- Investment Protection

WEB TOOLS

- Monitors switch status
- Zoning configuration
- Fabric topology
- Fabric events
- Requires browser & Java

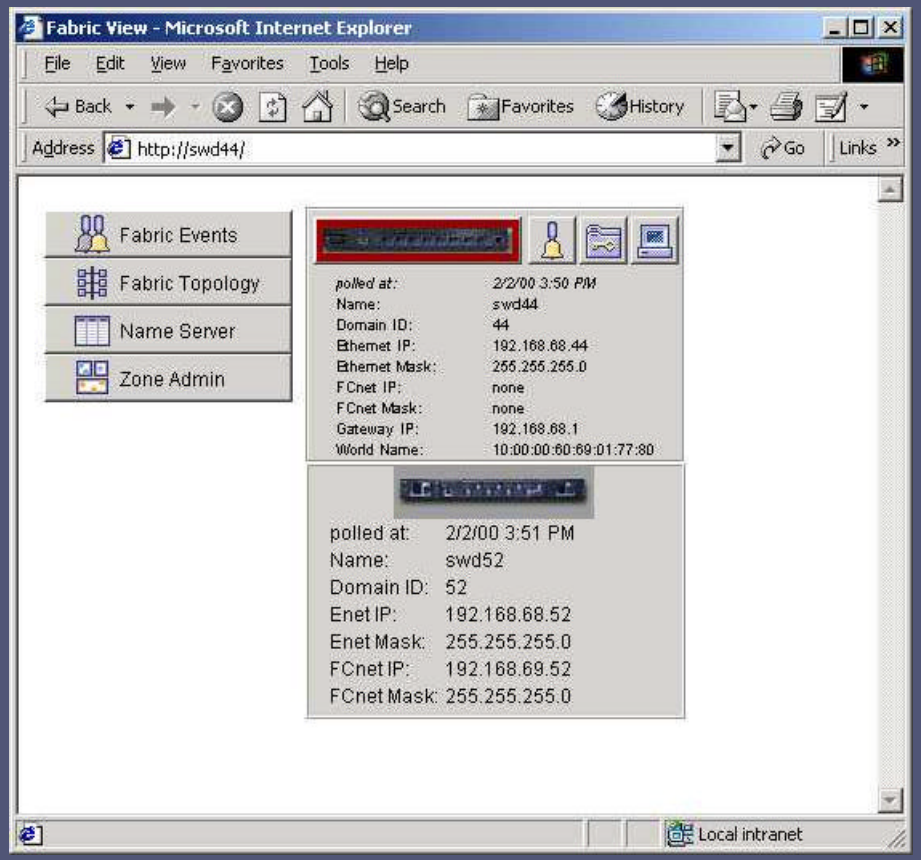

Fabric Watch™ Active Switch Management

Monitors key switch events

- Fabric states
- Errors
- Performance

Escalates via SNMP to:

- Enterprise Managers
- Switch Event Log
- WebTools Event View

Base line functionality

User can set thresholds

Useful for tuning Fabric

i.e. ISL traffic

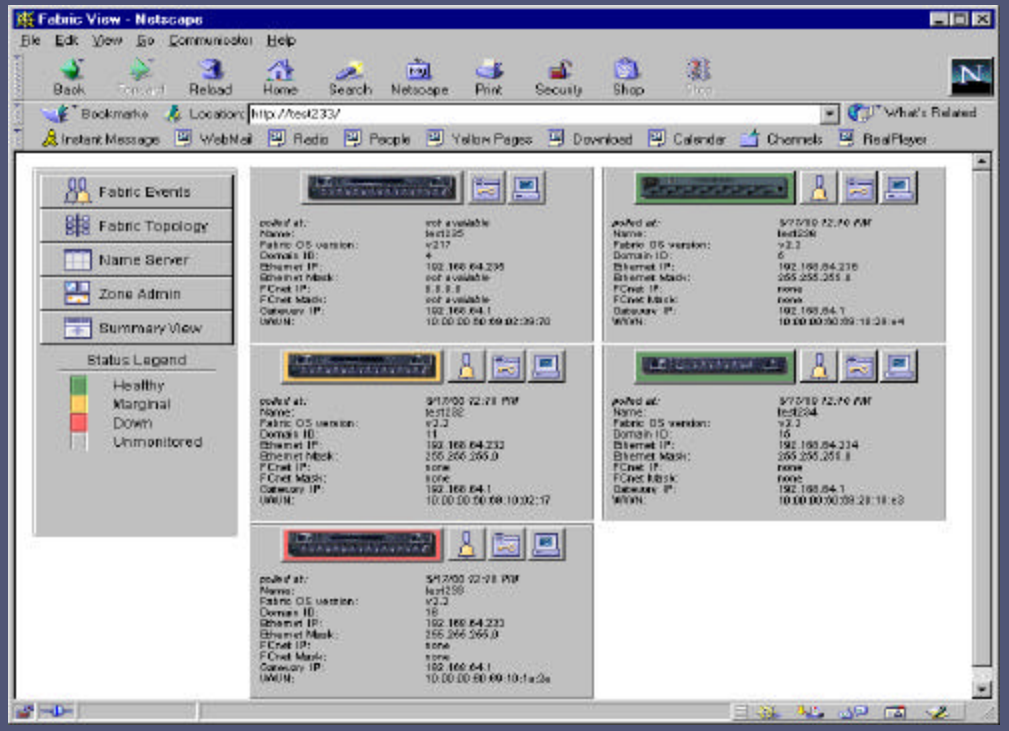

Brocade Extended Fabric

- Brocade exclusive Fabric OS based feature
- Brocade Fabric OS manages buffering between switches
- Managed as a single Fabric
- Full-speed (100 MB/sec) at distances of up to 70 Km
- Requires Extended Long Wave Length GBIC

Fabric Access Layer

Enable external applications to monitor and control the fabric

Enables integrated storage, server, and fabric management

Announced support from ~16 software providers

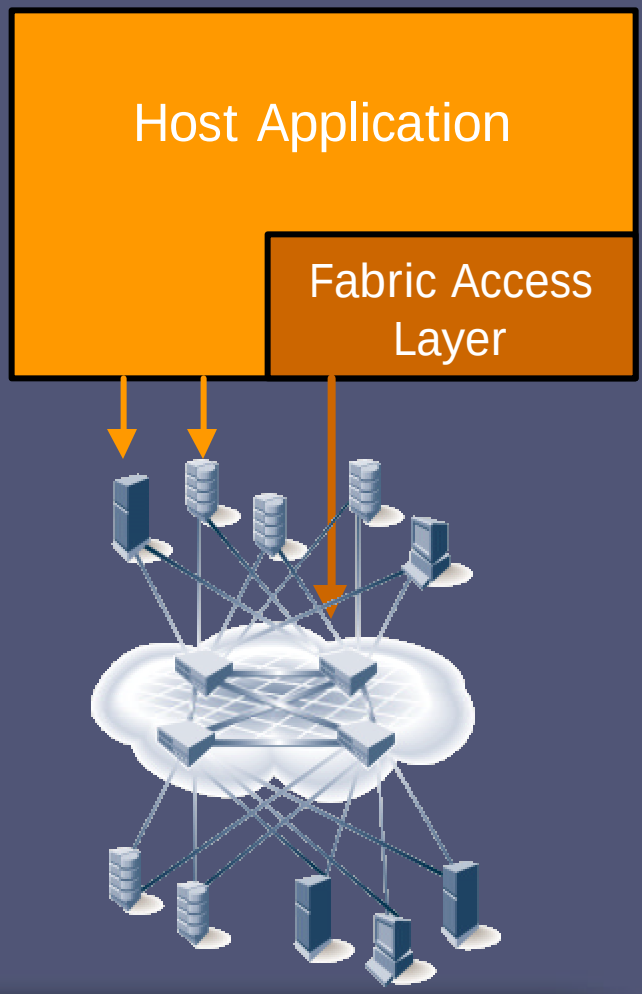

Interface between SANs and SAN Applications

HP Integrated SAN Solutions

HP SAN infrastructure enables SAN based application solutions

- High availability clustering – MC ServiceGuard and AutoPath
- LAN-free and centralized backup – Omniback software with an HP tape library
- Storage consolidation – SCSI storage consolidated to an Enterprise Class Storage Array
- Centralized management – HP Open View products

HP SAN infrastructure simplifies the management of IT resources

- HP-UX has SAN Fabric support “built-in”
- HP Open View products leverage unique SAN capabilities
  - Storage Node Manager
  - Storage Allocater
  - Storage Optimizer
  - Storage Builder

HP Integrated SAN Solutions Cont'd

HP SAN implementations enable non disruptive growth

- Non disruptive server addition
- Non disruptive storage addition
- Non disruptive application deployment

HP SAN infrastructure provides investment protection

- Support for legacy devices
- Compatibility with next generation SAN Fabric switches

HP SAN infrastructure provides security in a heterogeneous environment

Designing A SAN for High Availability

BROCADE

SAN Design Criteria

A Fabric is an extremely flexible architecture

- Small Clusters - 1 to 2 switches, 2 – 4 servers, < 1TB
- Departmental Meshes - 4+ switch meshes, > 2 servers, > 1TB
- Enterprise Backbone - Multiple switch meshes, Backbone, Server/Storage 2 Tier

Design based on …

- Disaster tolerance requirements
- Growth rate
- Performance characteristics
- Future application deployment

High Availability Dual Switch

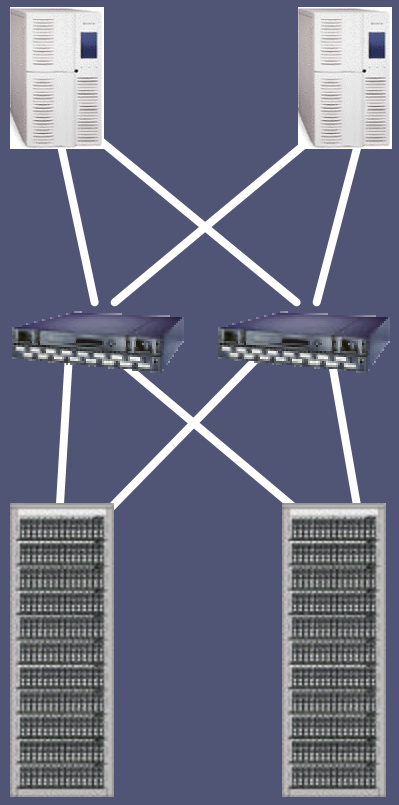

High Availability Mesh

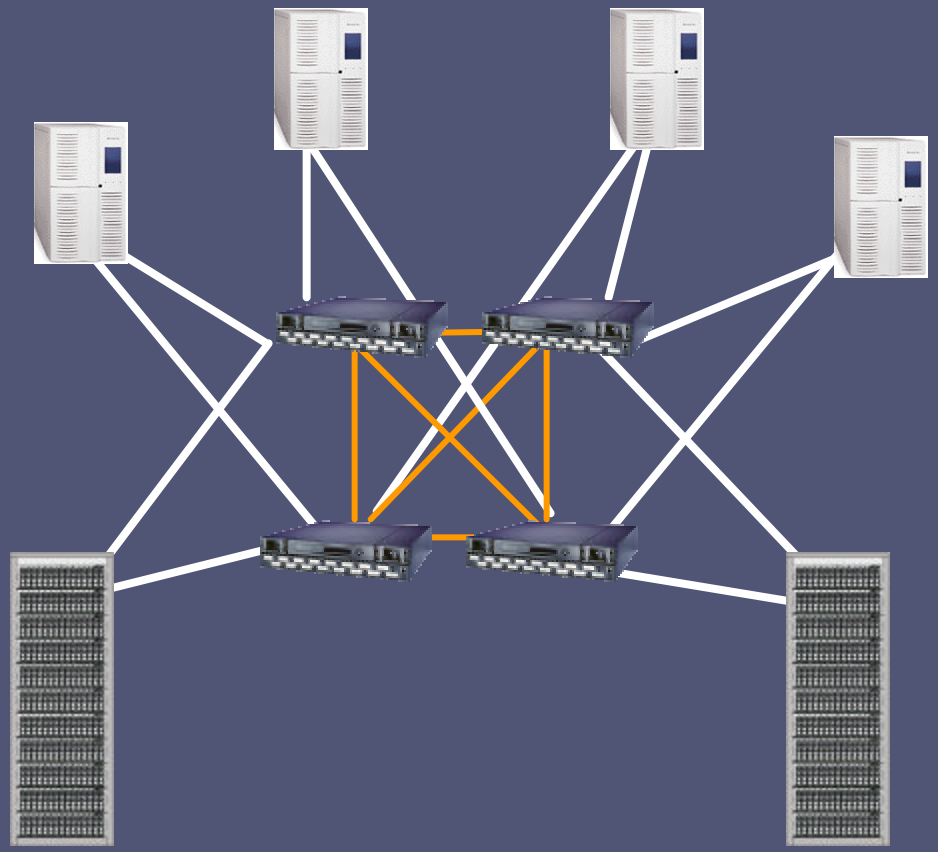

Enterprise SAN Fabrics

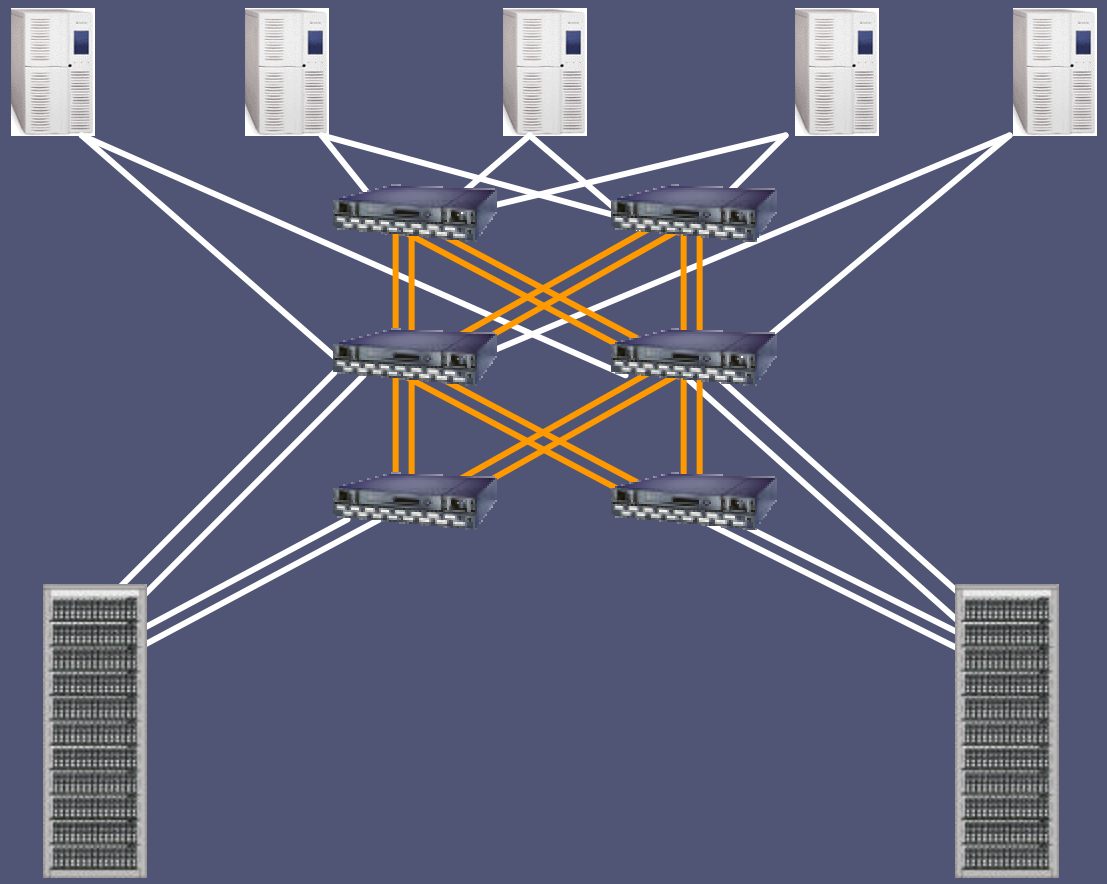

SAN Solutions

BROCADE

SAN Solution for HA-Clusters

Clusters Increase Application Availability and Disaster Tolerance

Better Utilization of Servers – Active/Active

Non disruptive deployment of servers, storage arrays and applications

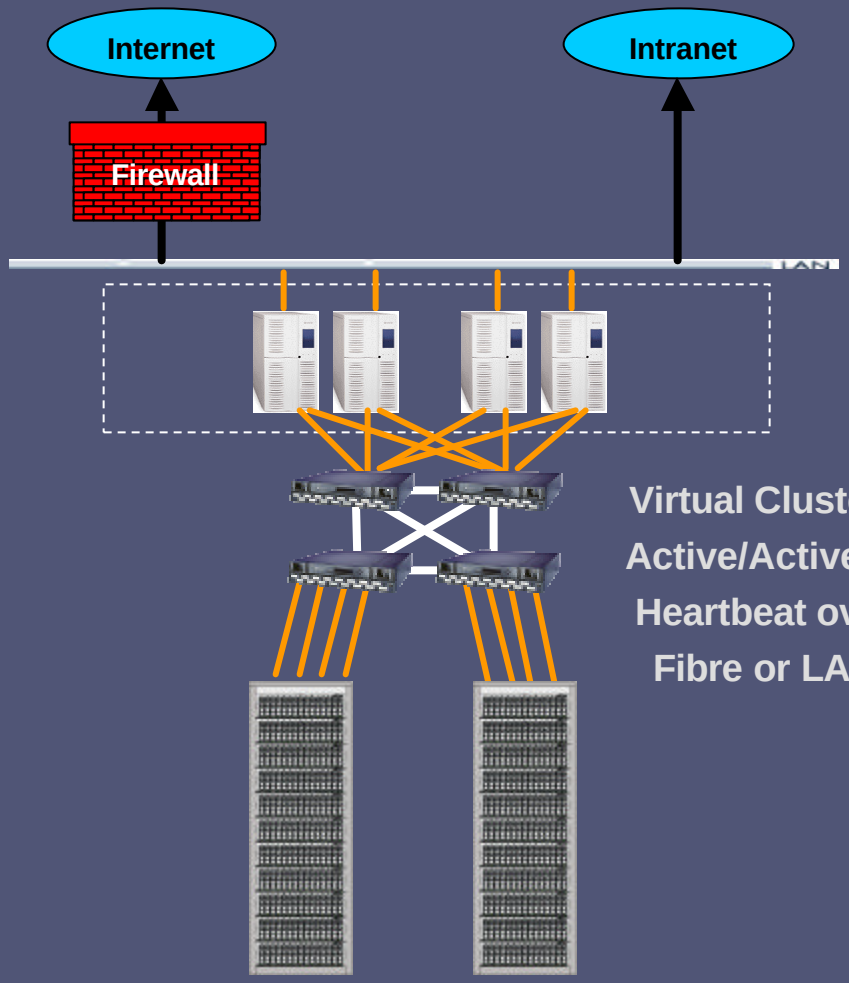

Virtual Clusters Active/Active w/ Heartbeat over Fibre or LAN

SAN Solution for Server Growth

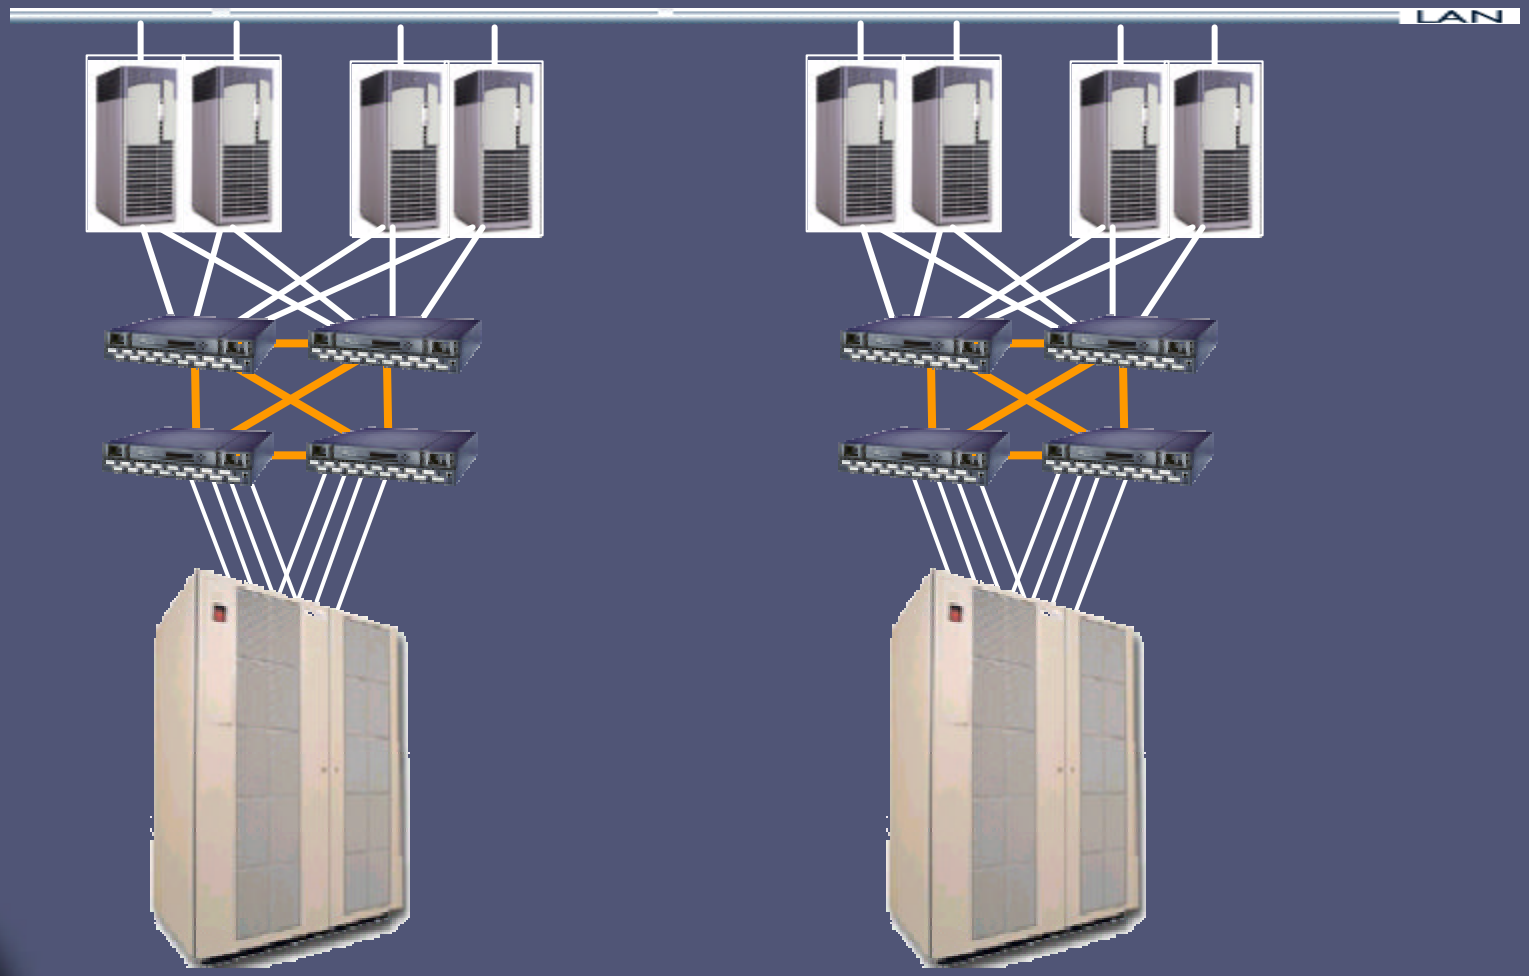

SAN Solution for Heterogeneous Server Growth

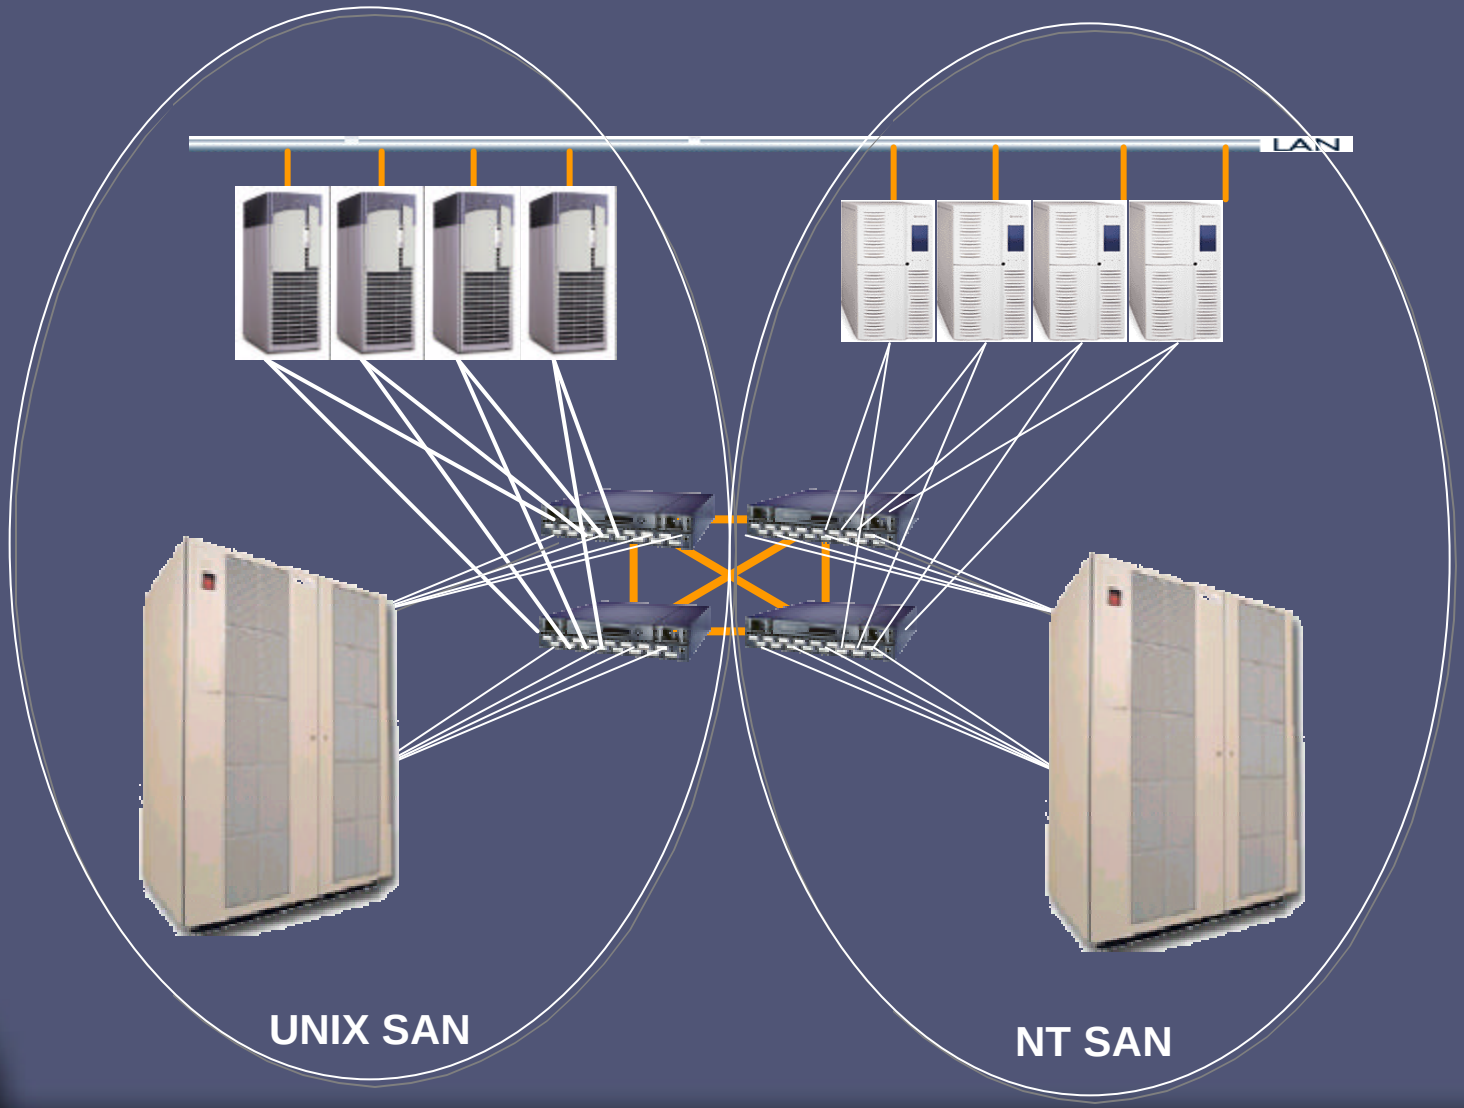

SAN Solution for LAN-Free Backup

- A SAN enables a centralized tape backup and restore implementation
- LAN-Free backup removes backup traffic from the LAN
- Allows for better utilization of valuable server and storage resources

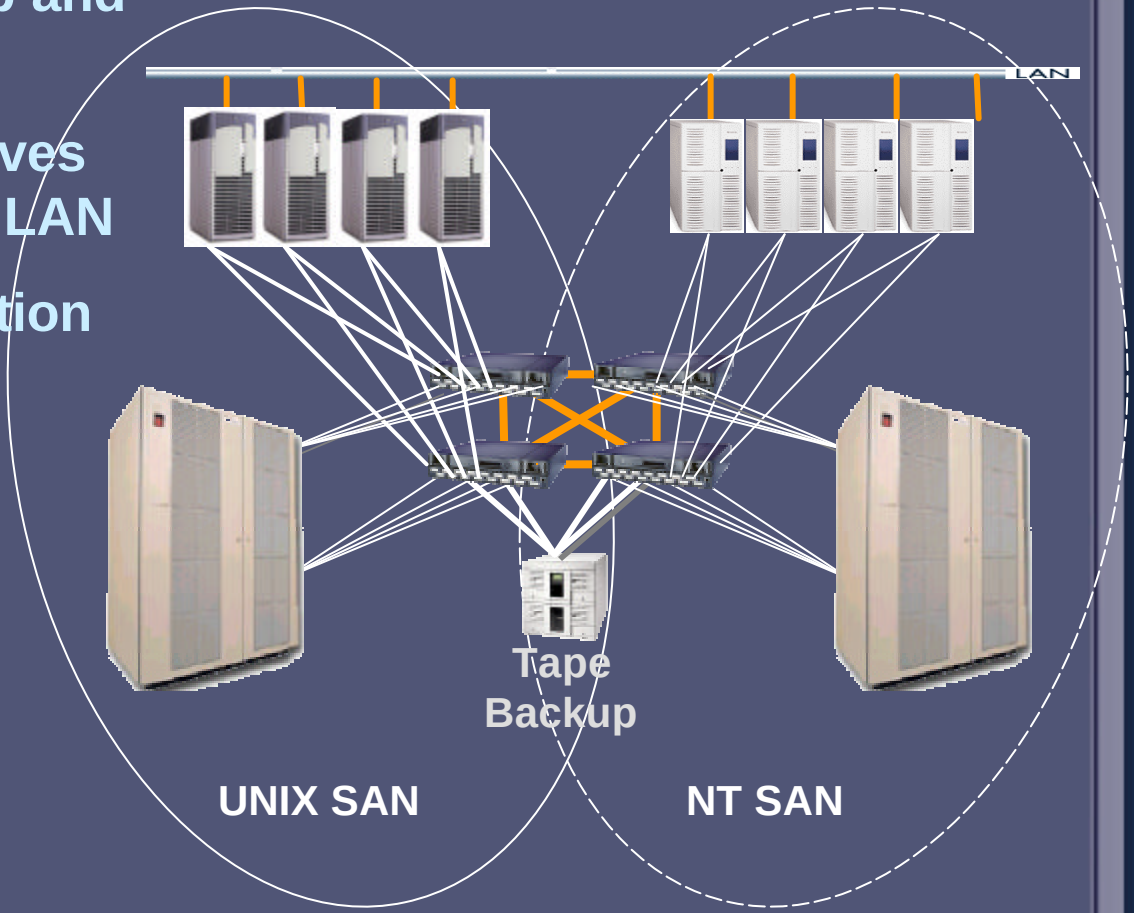

SAN Solution for Storage Consolidation

- NT applications typically grow in server/storage increments
- Storage utilization in NT servers averages ~50%
- SANs enable the addition of a Magnitude as a means of reclaiming under utilized storage capacity
- Enterprise class storage footprint creates additional application opportunities

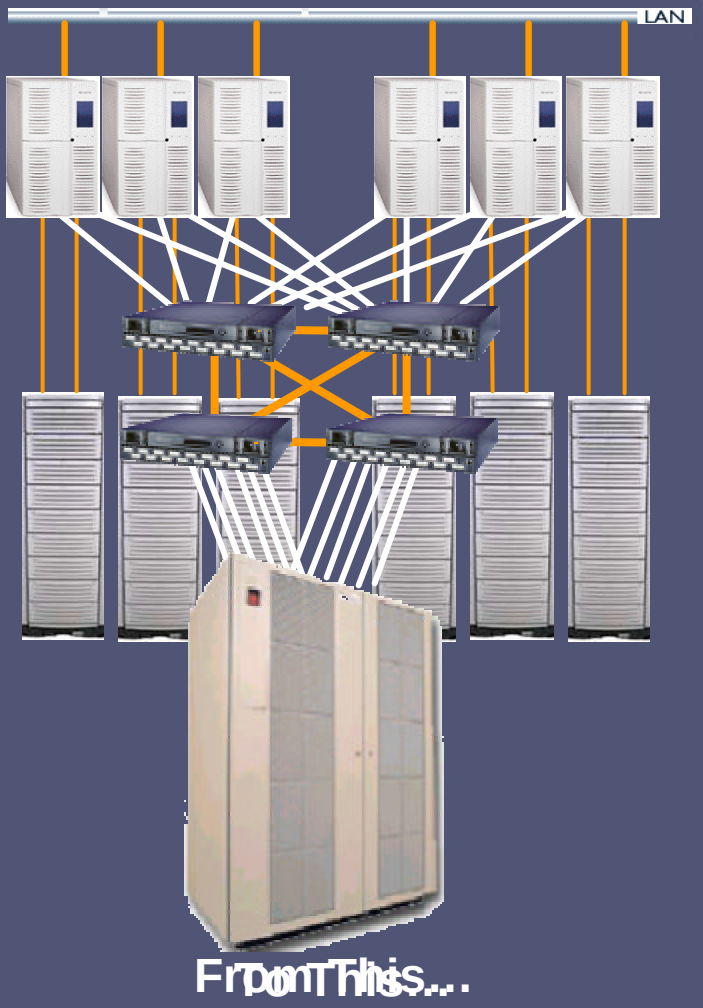

Enterprise SAN Fabrics

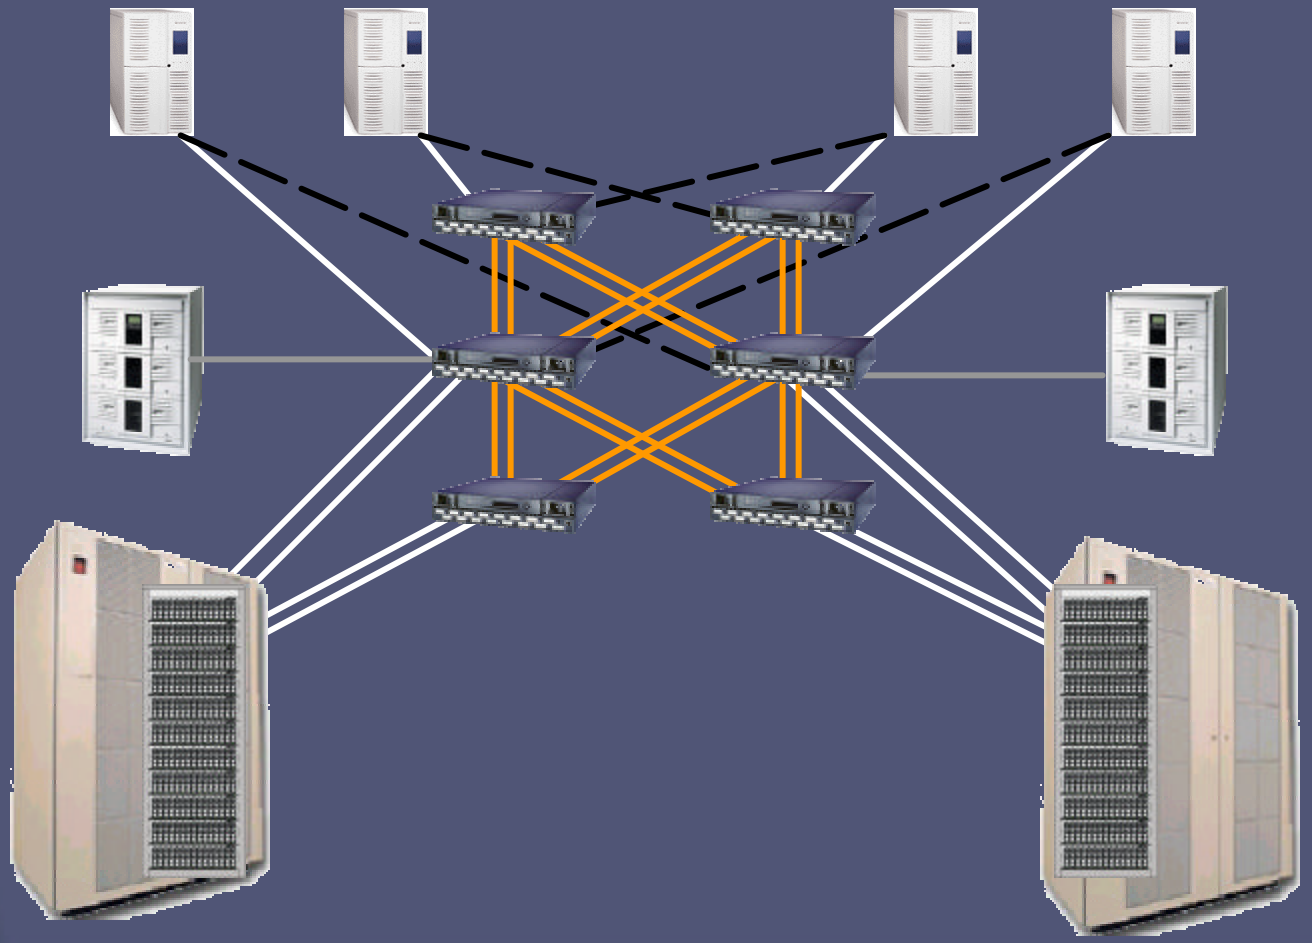

Case Studies

BROCADE

Leading Internet Services Provider

Company Background

The company is a leading provider of Internet infrastructure services in contextual commerce and personalized content markets.

Their services integrate content, commerce, and advertising to deliver personalized information to their customers' users through their Web pages, e-mail and wireless devices.

SAN Project Goals

Provide fast, reliable access to databases from any server

- Avoid the unpredictable behavior of IP-based front-end networks
- Reduce the amount of time required to reconfigure servers
- Overcome the restrictions of SCSI-based storage
  - Cabling distances
  - Number of devices
- Separate client, backup, and shared file system network traffic
- Reduce backup time

Build an infrastructure in which to deploy a high performance shared file system

Use clustering technologies to implement scalable, highly available services

Pre-SAN Configuration

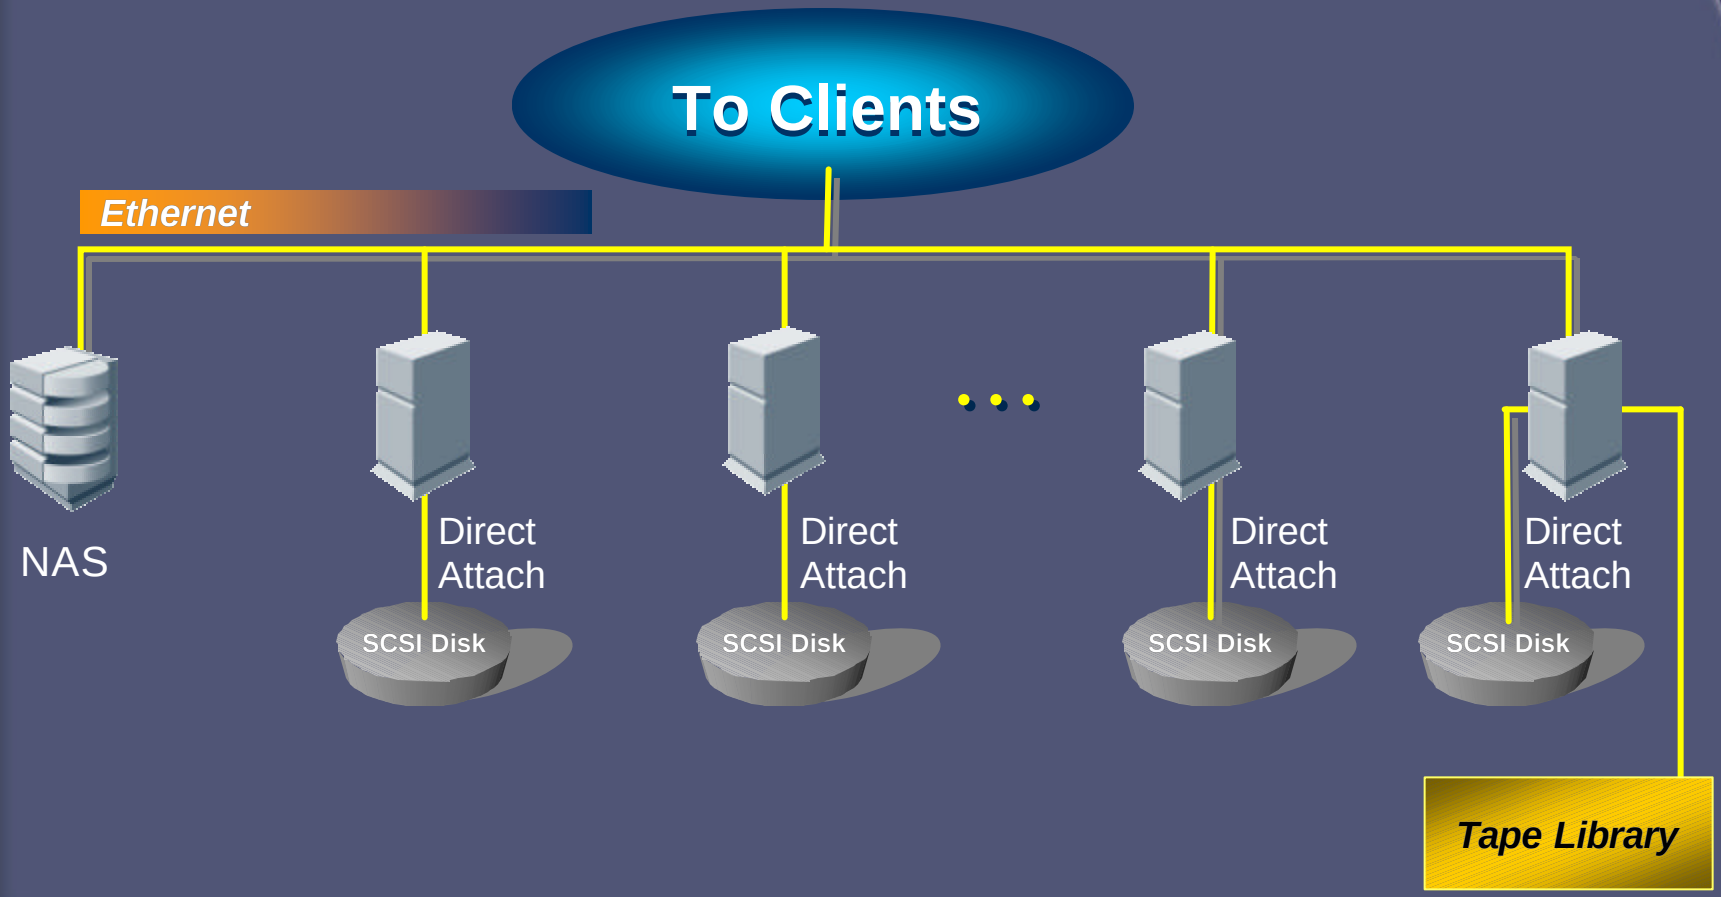

SAN Configuration

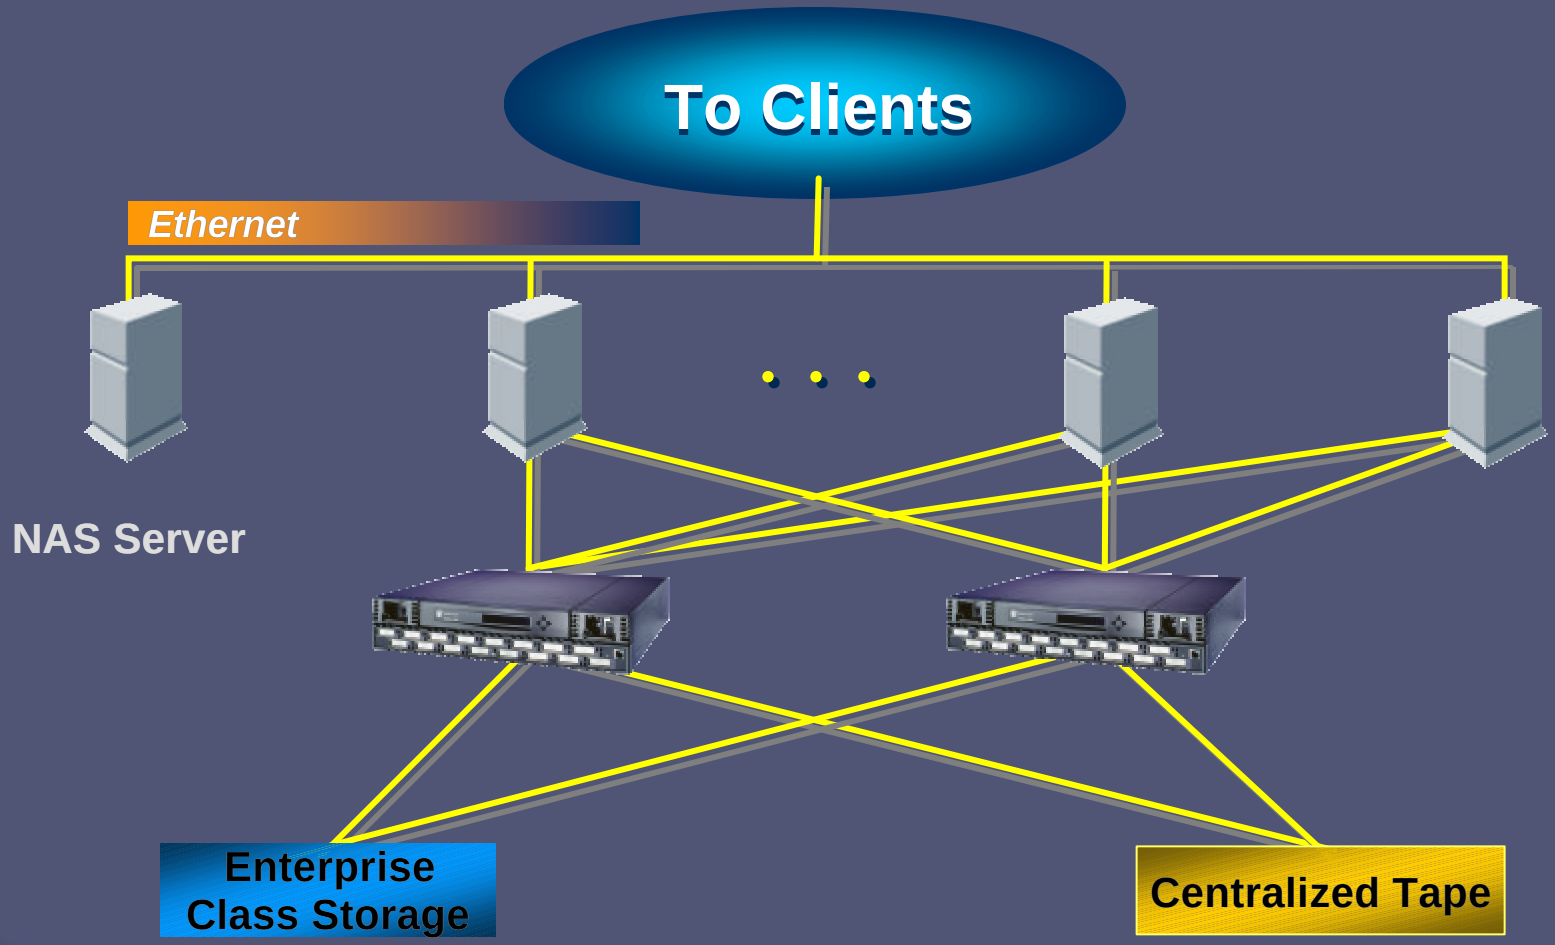

Post SAN Implementation Results

Increased database reliability and performance

Higher service (application) availability

- Immediate access to a failed server's database files
- Reduction of time required to add or move storage from 2-4 days to hours
- Reduced backup time from 10 to 4 hours
- Elimination of re-cabling for configuration changes

Improved performance of client network with the removal of server database traffic

Enhanced data network architecture that provides for future growth and functionality

Improved system administration productivity

SAN Return on Investment (ROI)

SAN Benefit	Savings/year
Reduced HA infrastructure and higher application availability	$266,000
Improved System Administration productivity	$110,000
Improve storage utilization and availability	$11,000
Reduction in backup window	$32,000
Total SAN Savings per year:	$419,000
SAN Costs:	$250,000

7 month SAN ROI

Large Telecommunications Company

"When we put it into our production environment we had a significant increase in the throughput for recovery and backup. Out of a 60 hour window we managed to shave a little bit more than 30 hours off that window."

Brocade-based SANs help avoid as much as $5 million per day in lost opportunity costs!

Large Telecommunication Company

Data center environment

- 10 major data centers with 8 discrete master servers
- 2046 client systems
- Growth to over 46TB in last 6 months
- Primary OS is HP-UX with Sun Solaris and Windows NT
- Primary data servers are HP with Auspex, NetApp, Sun, and Compaq
- Applications are Clearcase, PLS, Oracle, MSQL, Exchange

Large Telecommunication Company

Pre-SAN Issues

- Corporate data had doubled
- 20,000+ designers and developers
  - Business-critical data environment in core networks product division
  - Multi-millions of lines of code managed through centralized build systems
- LAN based backup of corporate data
  - Weekly backup was exceeding the allotted 60 hours
  - Backup ran into production time
  - LAN based backup was congesting the front-end LAN
- Required a more reliable, durable and available environment
- Required a more flexible storage implementation that worked in their heterogeneous environment

Pre-SAN Environment

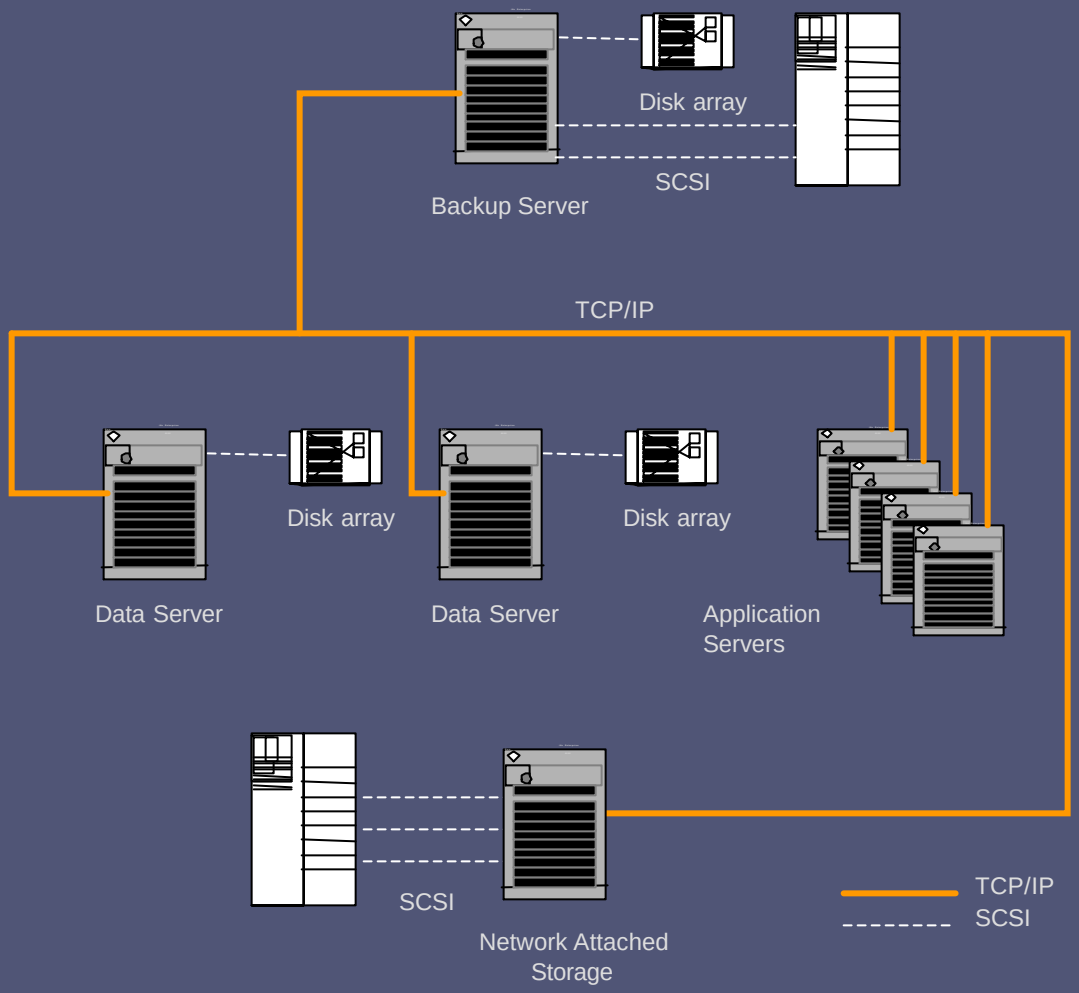

Current SAN Environment

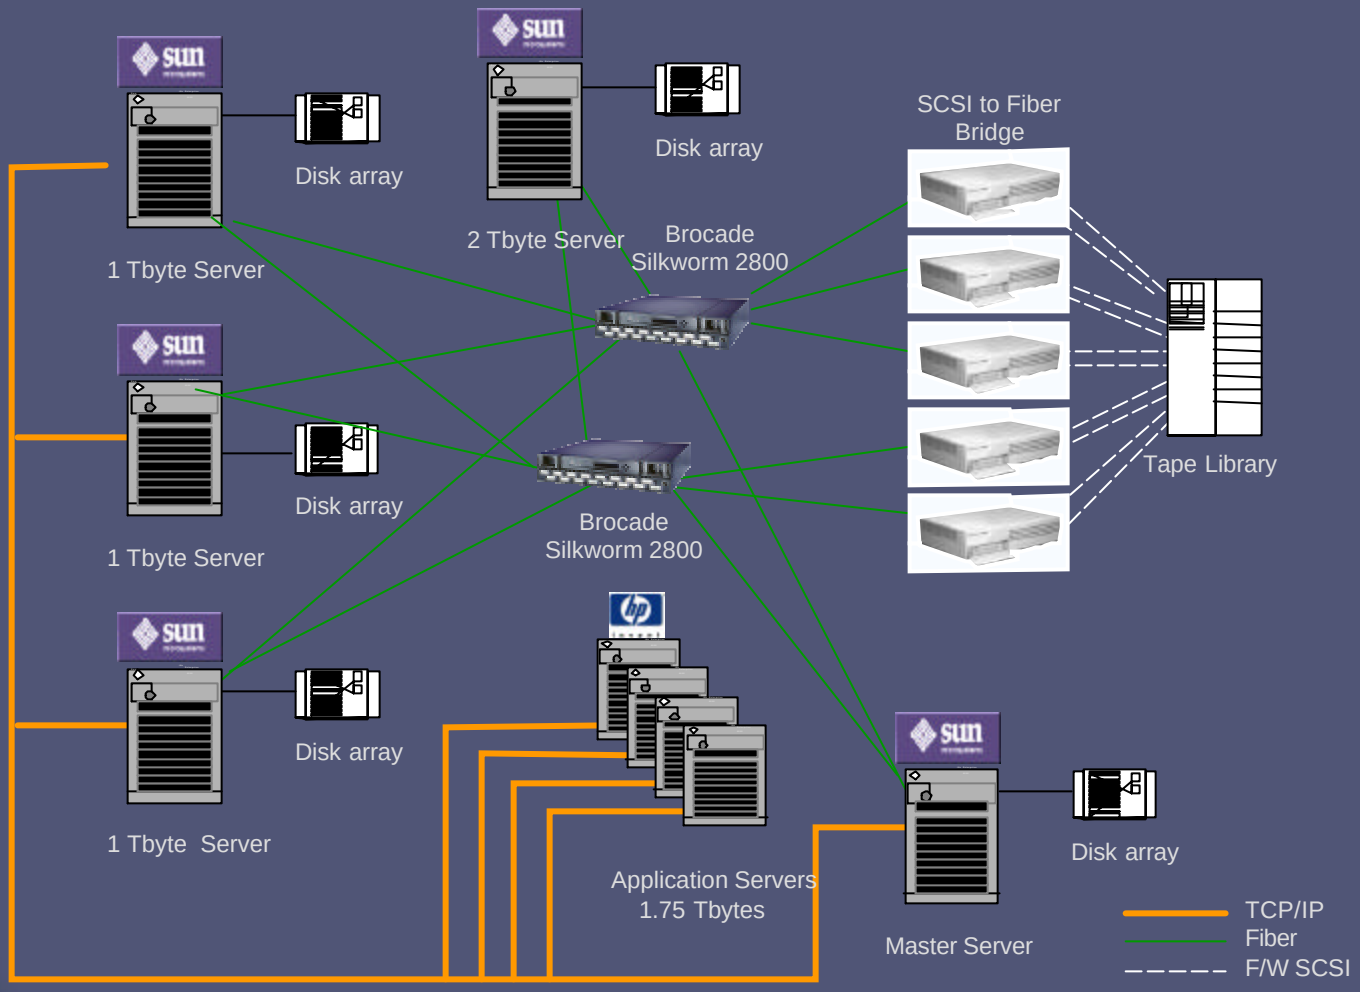

"Instant Payback"

SAN operational in 4 months!

SAN attached storage grew from 2 TBs - 6 TBs in 1 year

Reduced loaded cost per GB from $28/GB to $22/GB

Reduced network traffic by 90%

Reduced backup window from 60 to 30 hours

- Max throughput went from 32MBs/sec to 60 MBs/sec

Improved build recovery time by 100%

- Recovery builds take ½ the time

SAN Internetworking

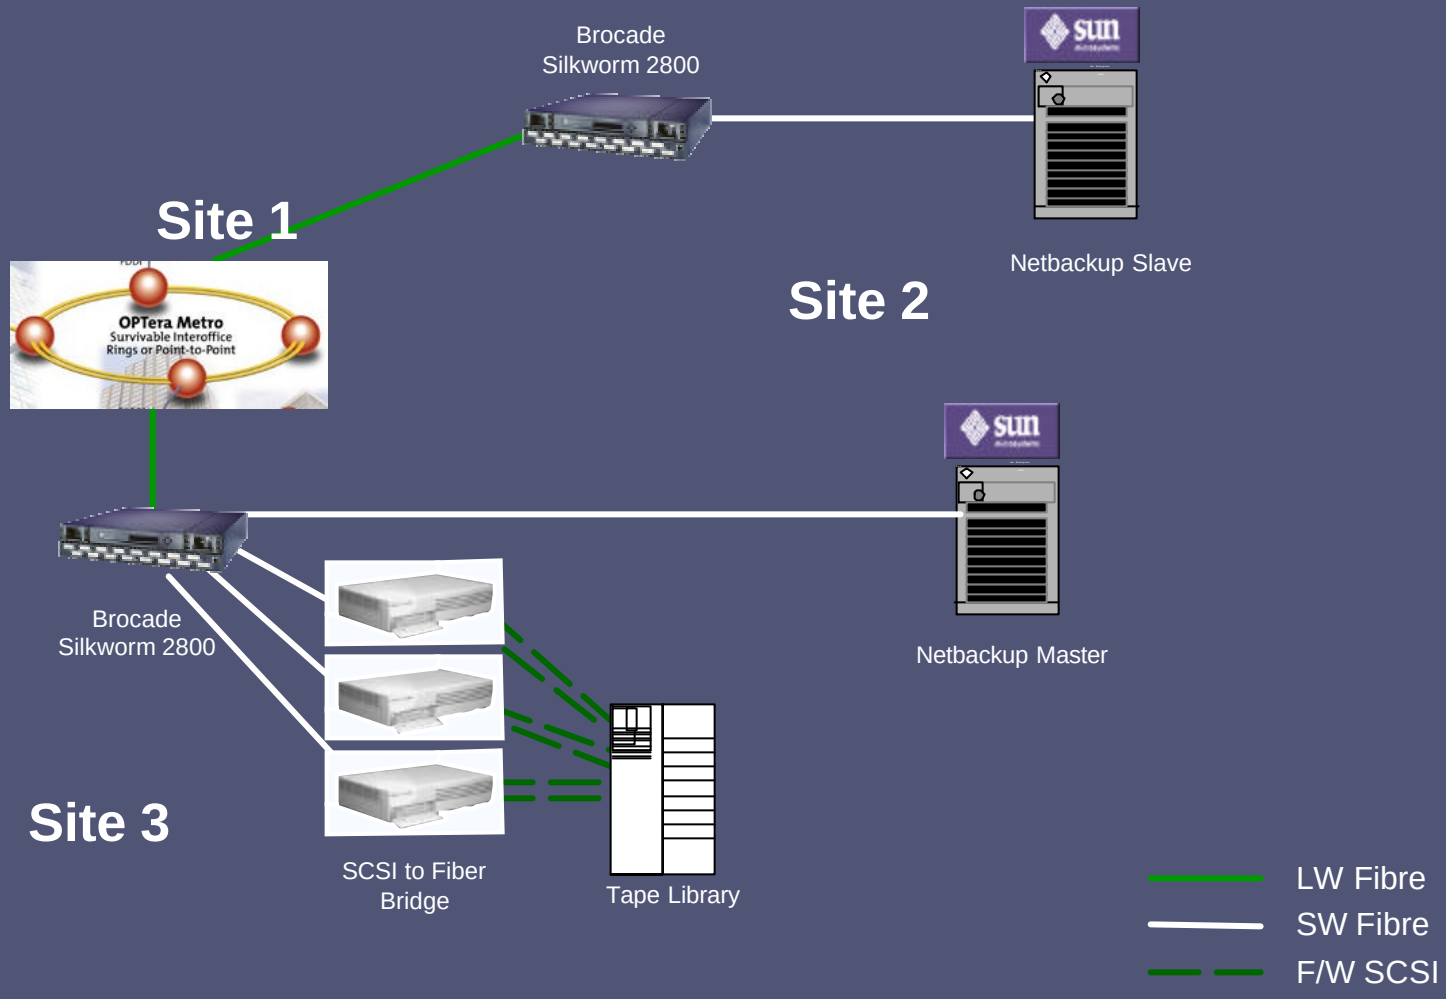

Large North American Hospital

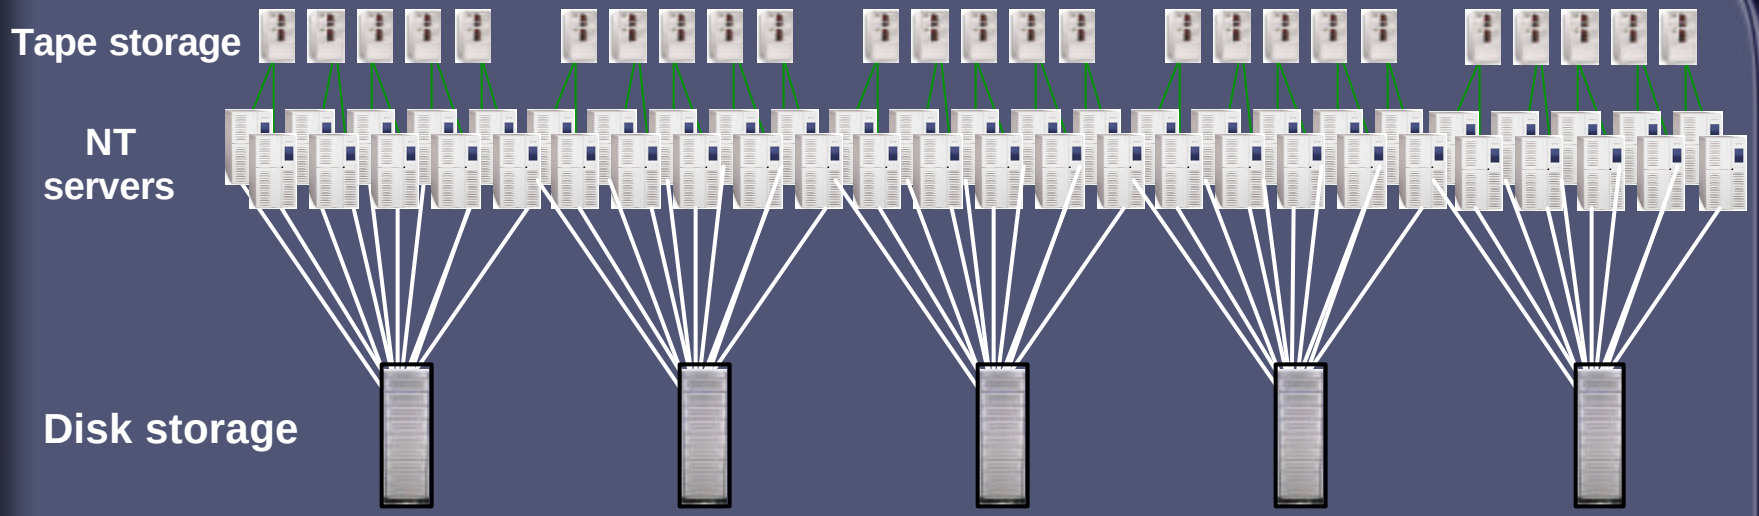

Pre-SAN infrastructure

- 50 NT servers and growing
- Multiple SCSI-attached libraries (wasted capacity)
- One-to-one SCSI architecture provides inadequate backup window
- Tape management problems

Large North American Hospital

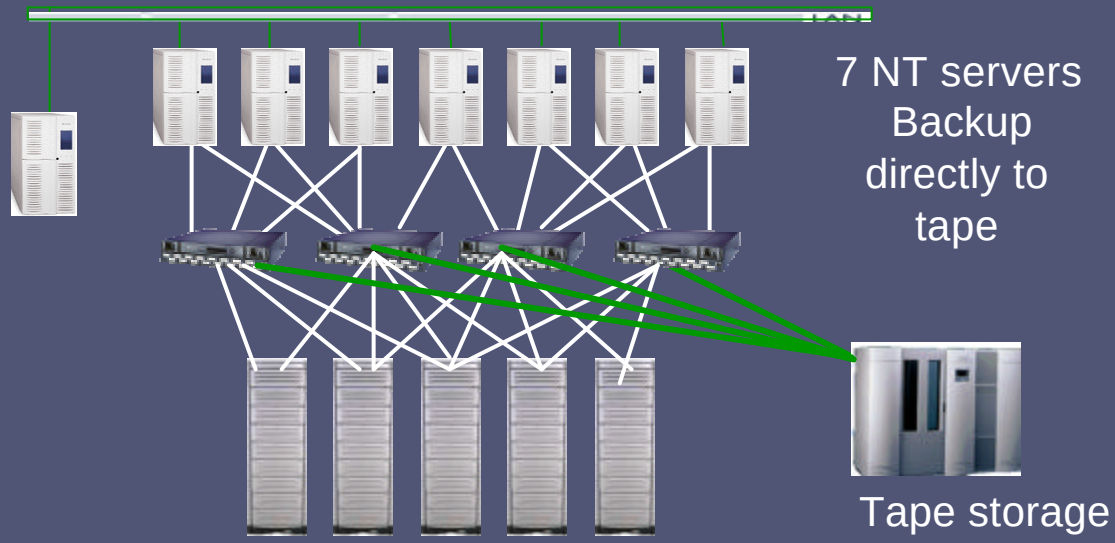

Solution: highly available and scalable SAN

- Reduced from 50 to 7 primary Exchange servers
- Backup window reduced from 18 hours to 4-6 hours
- Improved automation
- Reduced administration

A Leading SSP Case Study

Design goals

- Support multiple customer storage requirements in a major metropolitan area
- Long distances (>10kms) required for disaster protection

SAN backbone connected to DWDM and SONET ring for long distances

- SAN islands isolate customer environments
- SAN backbone (core switches) used for centralized storage

Multiple NAS implementations running NFS

Environment

- Key applications: NFS, Oracle OLTP and Data Warehouse
- Key platforms: HP-UX, NT, Solaris, AIX
- Approx. 100TBs of storage
- Management: CA Unicenter, Veritas NetBackup

A Leading SSP Case Study

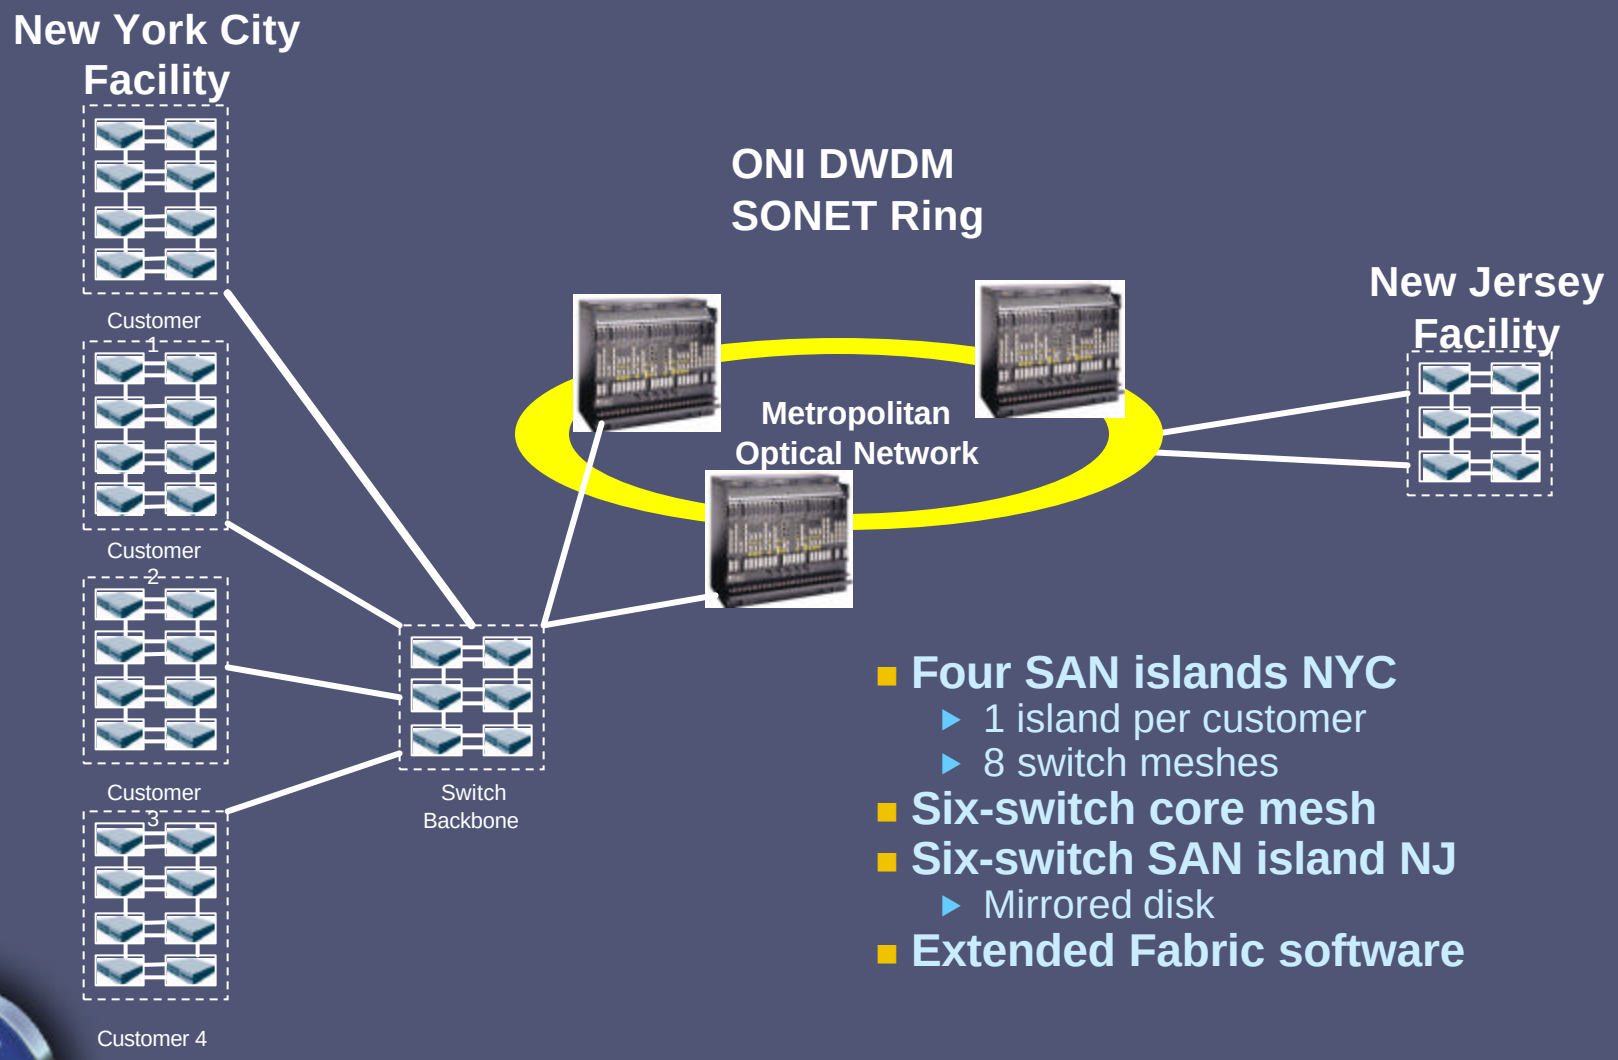

- **Four SAN islands NYC**
  - 1 island per customer
  - 8 switch meshes
- **Six-switch core mesh**
- **Six-switch SAN island NJ**
  - Mirrored disk
- **Extended Fabric software**

Thank you

BROCADE