

Concepts In Storage Area Network Design

- Jim Gursha, President
- High Performance System Solutions, Inc.
 - 1735 York Avenue, 32H
 - New York, NY 10128
- HP Interex West Symposium
Session # 3148 March 2004

HPSS

1735 York Avenue, 32H

New York, NY 10128

212-831-0291/917-359-2087

jimgursha@sandisks.com



Topics

- Storage Area Network Rationale On A Departmental and Enterprise Basis.
- Consolidation Efficiencies of A Storage Area Network.
- Server Effectiveness
- Storage Savings
- Fiber Channel SAN's
- SAN Fabric
- Fabric Benefits
- Basic Fabric Design
- Switch Interconnectivity
- SAN Based Backup/Restore.
- Network Considerations For SAN's.
- High Availability and Disaster Tolerant SAN Considerations.
- HP SAN's Virtualization Concepts, Differences and Considerations.

SAN Rationale

Departmental Vs Enterprise

Consolidation Efficiencies

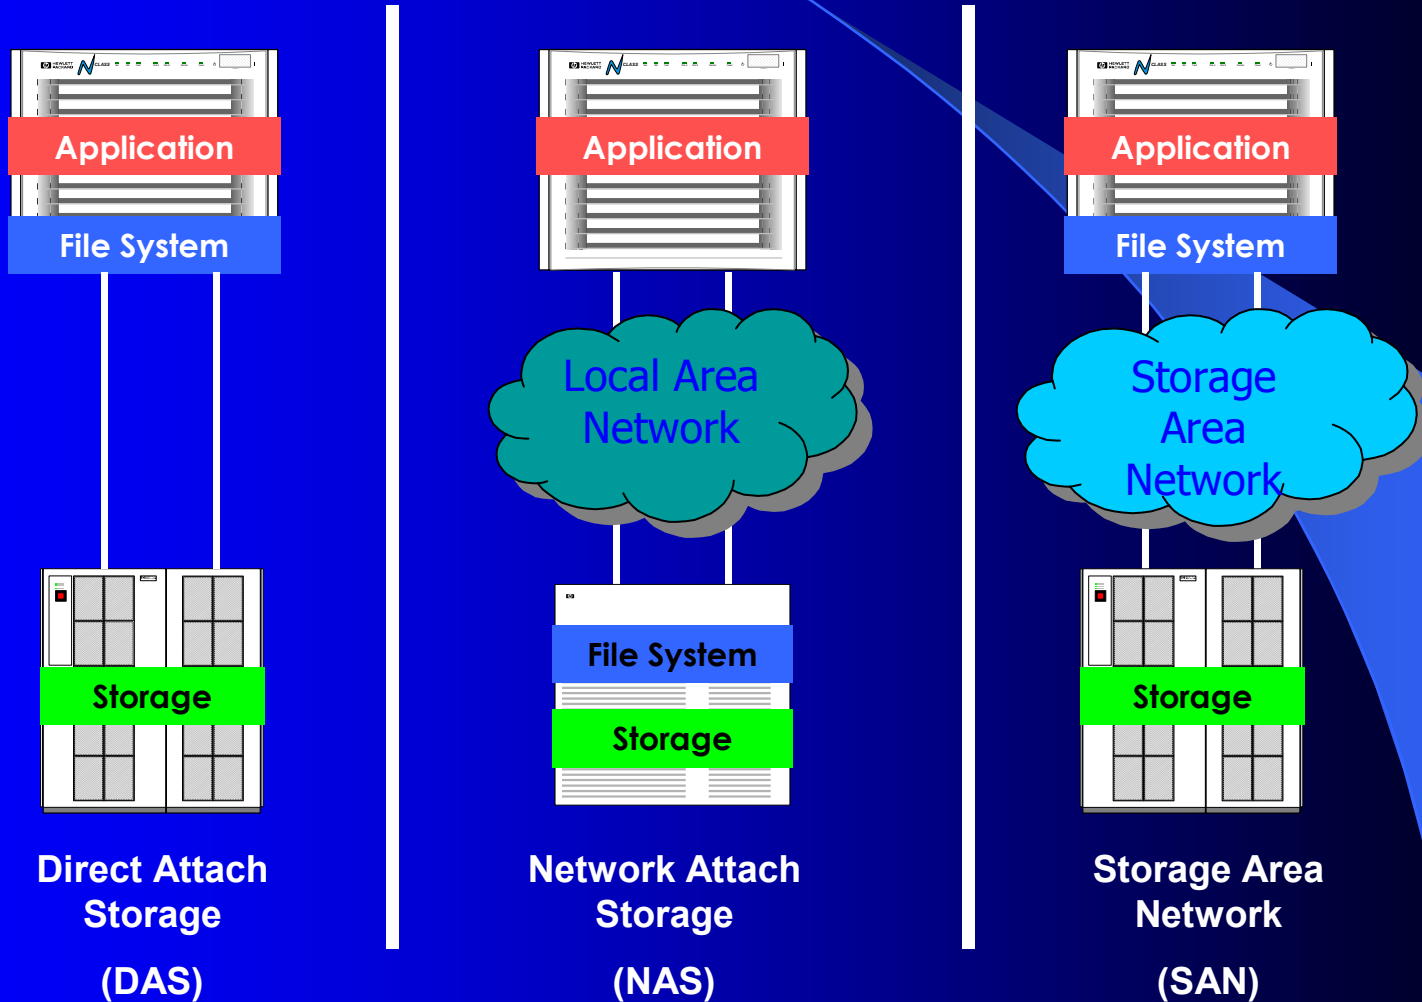
What is Consolidation?

- Server and storage consolidation is a process gaining greater efficiencies and simplicity in your IT infrastructure. This process involves evaluating your IT environment and implementing a strategy of combining servers and storage to address your business requirements for:
 - Reducing costs
 - Increasing manageability
 - Increasing availability
 - Increasing performance
 - Increasing flexibility

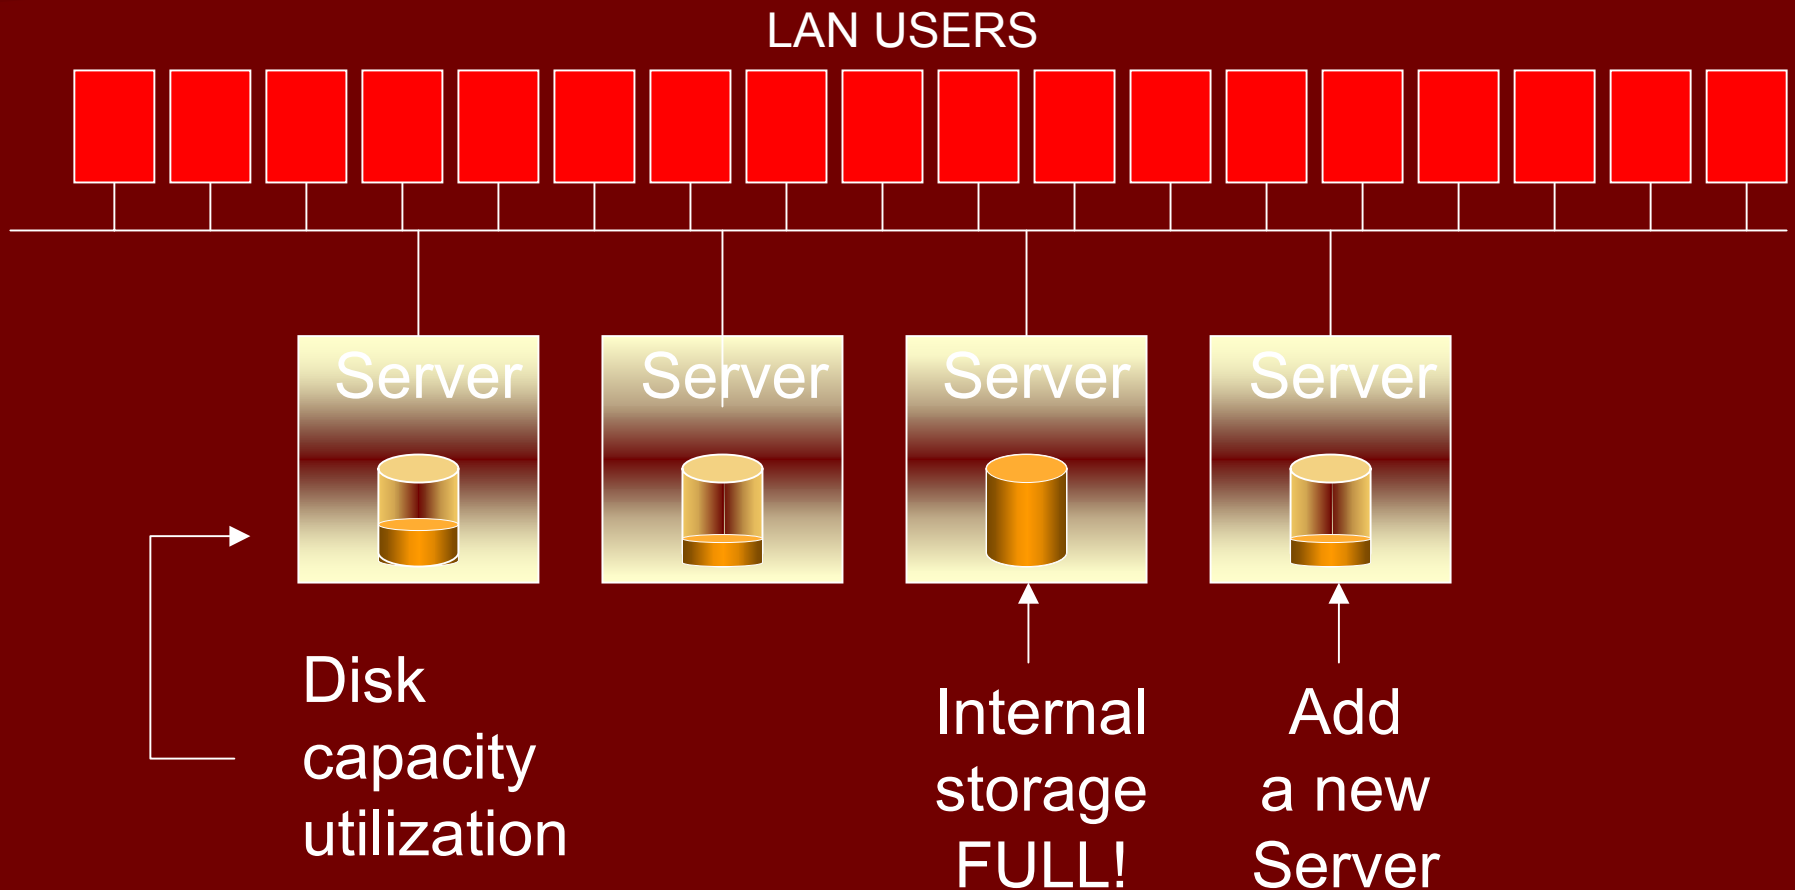
Consolidation Candidates

- Candidates for server and storage consolidation are companies which are experiencing:
 - Recent mergers or acquisitions and need to integrate their IT infrastructure
 - High management costs for large numbers of servers and storage systems
 - Loss of or reduced support staff and high IT staff training costs
 - Older servers and storage systems which need to be replaced or come off lease
 - Need to replace outdated and non-standard platforms with a corporate standards

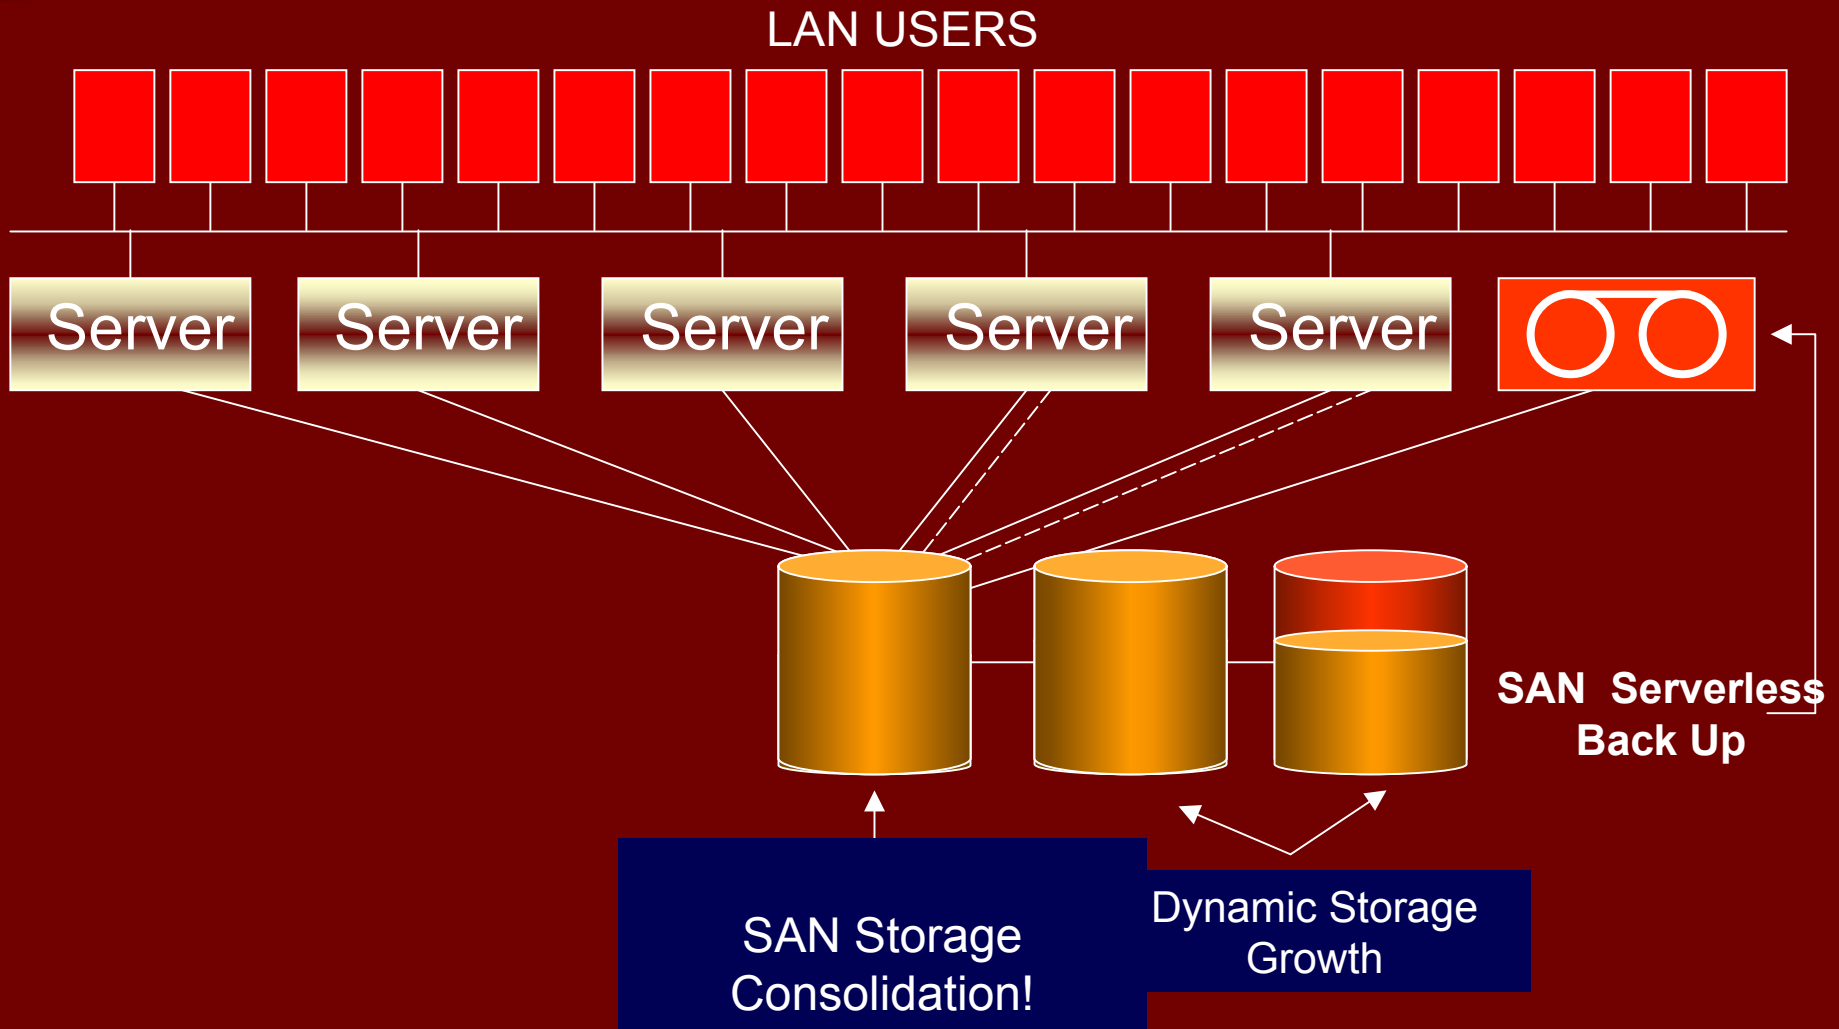
Storage Connect Architectures



Example DAS Server Storage (No Consolidation)




SAN Consolidation Strategy



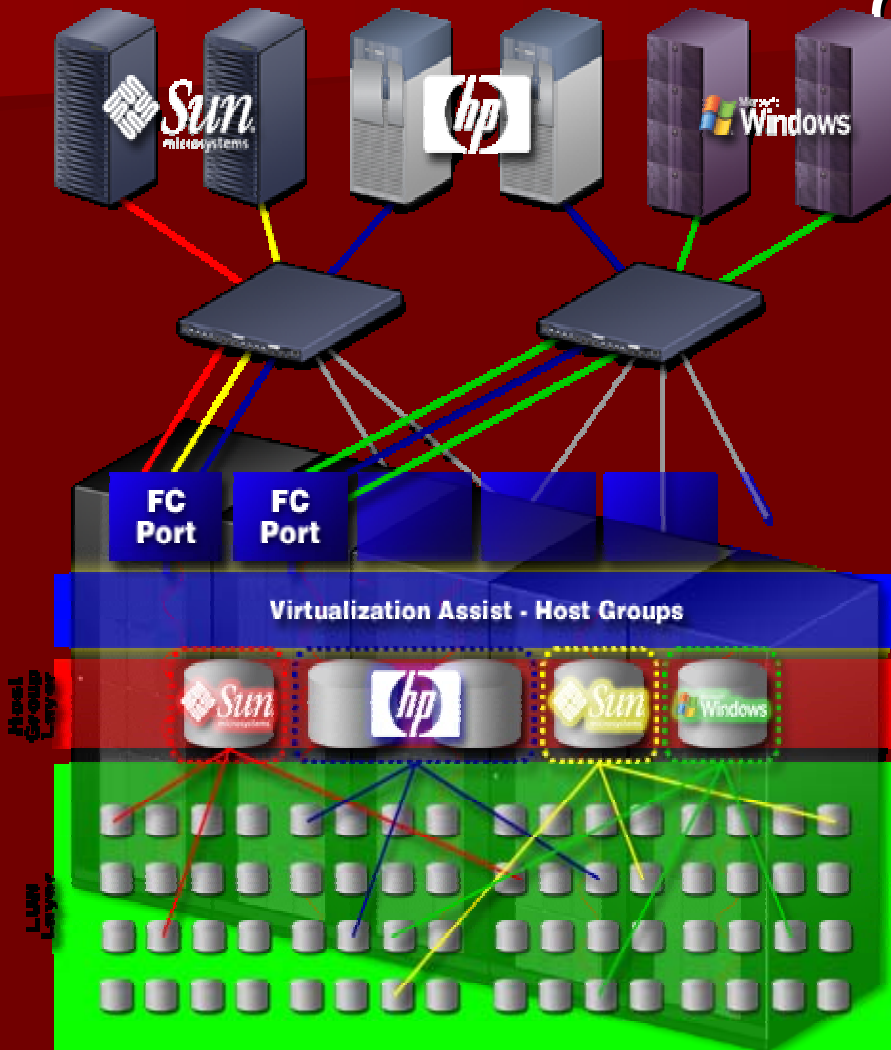
Ultimate Consolidation Including NAS (Where You Want To Be)

Benefits

- A single storage pool to manage
- Fewest servers to manage
- SAN investment optimized
- Unlimited NAS scalability
- Maximum storage flexibility and performance




HP XP Series



Consolidation
requires

- Broad connectivity
- Very high throughput
- Large capacity
- Capable management tools
 - Security, performance, allocation, availability

Storage Consolidation with SAN



Server Effectiveness

HP ProLiant server from 4P to 8P

Enhanced enterprise performance



ProLiant BL40p, DL560, DL580, ML570 - up to 4 processors











ProLiant DL740, DL760 - up to 8 processors



NEW:

A complete family of Itanium-based servers: HP's Integrity Servers




High end	<p>Launch: 6/30/03</p>  <p>Superdome 16,32,64-way</p>	<p>Fall 2003</p>  <p>HPTC Cluster</p>	<p>2004</p> <p>Itanium</p>  <p>NonStop</p>
	<p>Launch: Fall 2003</p>		
Midrange	 <p>8-way</p>	 <p>16-way</p>	
	<p>Launch: 6/30/03</p>		<p>Fall 2003</p>
Entry	 <p>rx2600 (2-way)</p>	 <p>rx5670 (4-way)</p>	 <p>Carrier grade (2-Way)</p>



Storage Savings

Driving Trends

Storage Growth Is Growing *Very Fast!*



IT staffing budgets increasing 5% a year

Corporate data growth increasing 100% year

Storage Administrator Implications



SAN Customer Benefits

(Storage Savings)

- Better disk capacity utilization (average unused storage in a DAS environment typically >50%).
- Manage with less people (typically one storage administrator can manage 4 or more times the storage).
- significantly shorter backup windows.
- Drives and tape on the same storage network.
- Reduced LAN traffic.
- Flexible environment for future growth.
- Usually should be able to add storage, switches, tape to the SAN while applications run.

Fiber Channel SAN's

Host Bus Adapters

- Individual Component Utilization Is A Necessary Part Of SAN Architecture and Implementation.
- Selecting the Right HBA Will Lower Overall SAN Costs.

Host Bus Adapters

- Connect the Server to the SAN.
- Alleviate the Server From Some I/O Processing.
- Typically, Assist in the Execution of Parts of Communications Protocol.
- Compatibility Across HBA's.

Fiber Channel Switches



➤ Hubs

- 7 port and 12 port models
- Much lower cost than switches
- Used in smaller SANs

➤ Switches

- 8, 16, 64 and 128 port models
- Better isolation than hubs
- Switches are more \$ than hubs.
- Scalable to large SANS
- Cascading, Zoning, Quickloop
- Higher MTBF, easier repair



➤ GBICs (GigaBaud Interface Cards)

- Short and long wave variants
- LW to 100Km
- SW to 500m



FC Switch Products

**2GB Products
Are Replacing 1Gb!**

SAN Core 2/64

SAN Director 2/64

SAN Edge 2/32
SAN Edge 2/16

SAN Switch 2/16

SAN Switch 2/16 EL
SAN Switch 2/8 EL

SAN Director 64

MSA Fabric
Switch 6




SAN Switch
Integrated 32/64

SAN Switch 16-EL
SAN Switch 16

**Current
1Gb Products**

MSA1000 only

SAN Switch 8
SAN Switch 8-EL



Fiber Channel Implementation

Switches implement FC Fabric

- Switches act like network routers
- *FC Switches provide non-blocking Dynamic Paths Through The Fabric*
- Multiple communication paths
- Bandwidth is aggregate of paths
- Switches provide logical isolation

Departmental SAN's

Modular SAN Array 1000

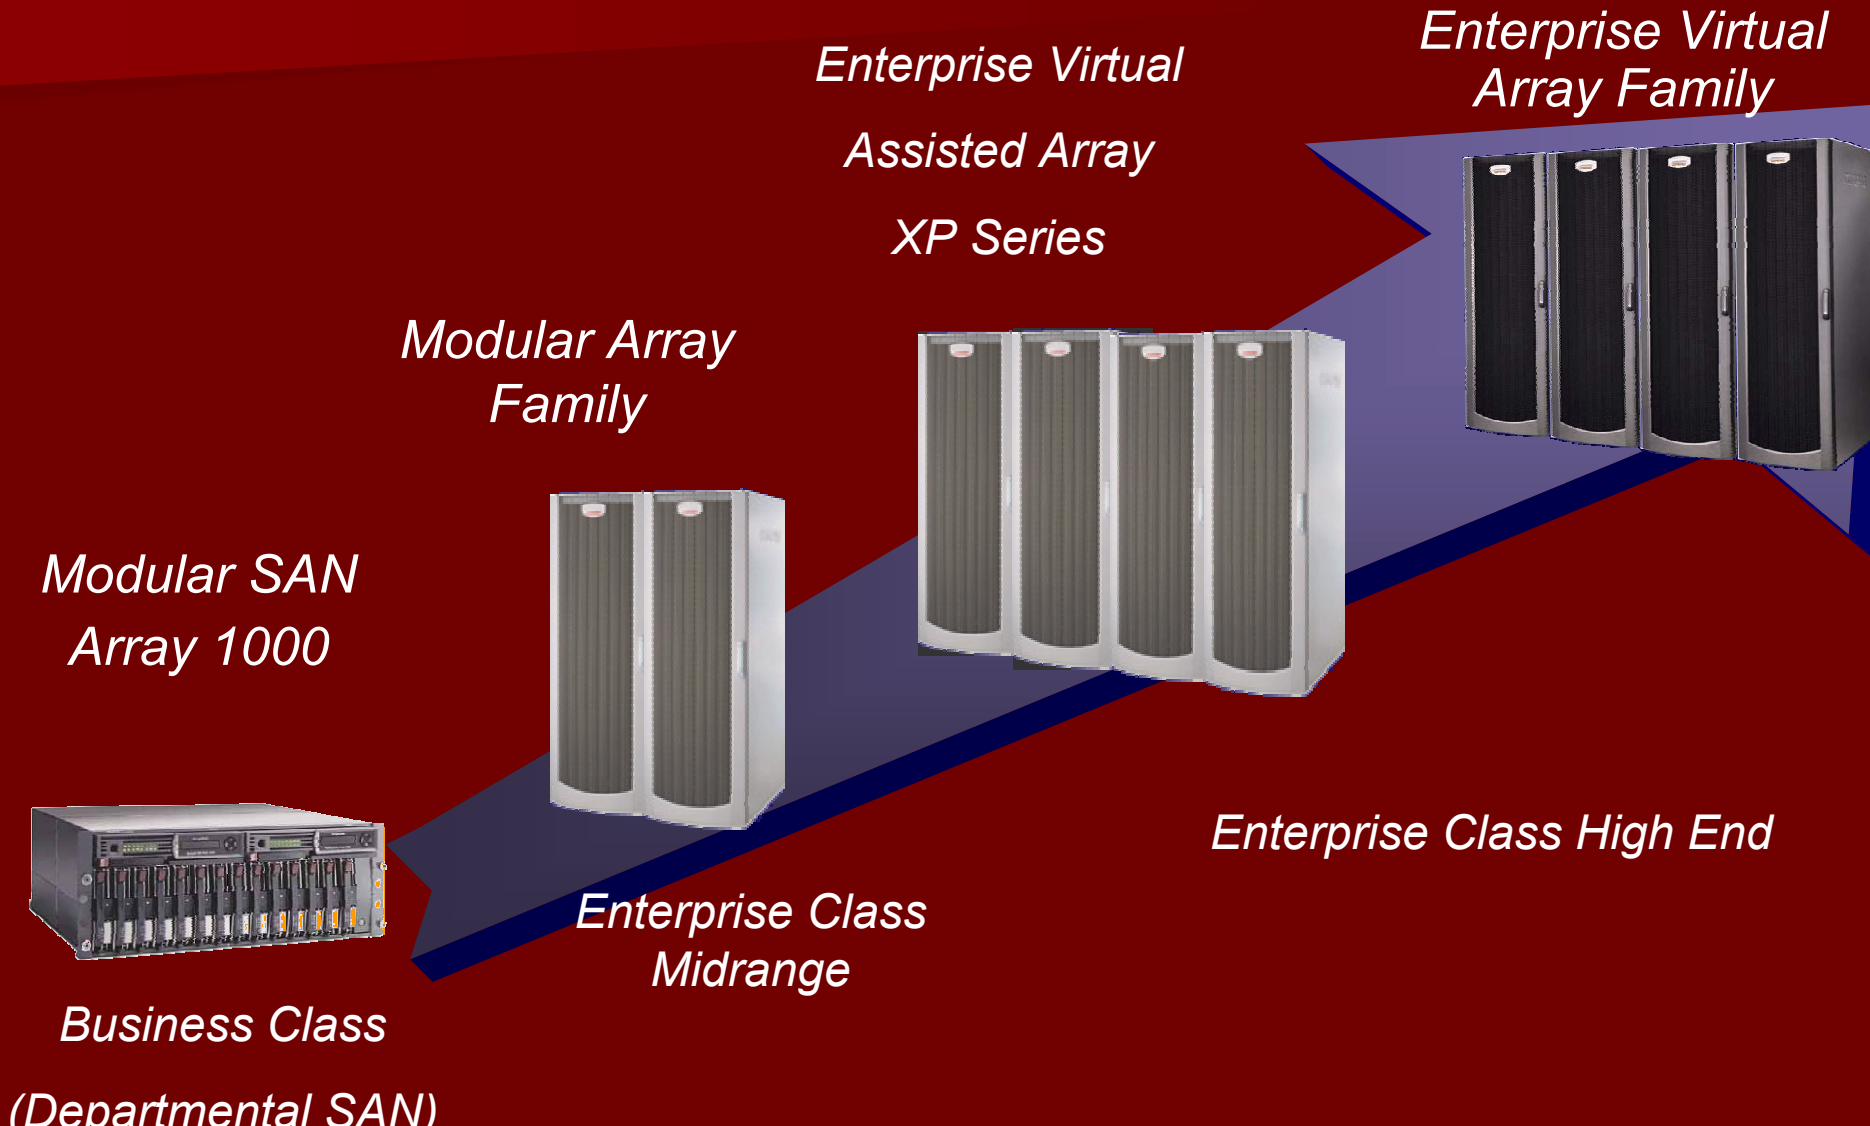
Technical Overview

- Key Elements
 - Scalable to 3TB
 - Fully redundant components
 - 2Gb FC host connections
 - Ultra3 SCSI -- 1" universal drives
 - High performance controller
 - Boot from SAN capability
 - ProLiant & x86 servers
- Competitive Advantages
 - *Built on Compaq's exclusive DtS Architecture*
 - Embedded 2Gb fabric switch option
 - Advanced Data Guarding



Enterprise SAN's

Scalable Storage



StorageWorks Enterprise Virtual Array

- Modular, scalable and highly available design
 - Redundant Power
 - Redundant Cooling
 - Distributed hot spare disk drives
 - Mirrored Cache with battery backup
- Multi-level virtual RAID architecture
 - **RAID V0**
 - **RAID V1**
 - **RAID V5**
- Full 2Gb fibre channel front to back
 - 1Gb fibre channel SAN compatible
- Co-exist with HSG80 class Modular Array's
- Compaq SANworks SecurePath™ compatible
- High performance HSV110 Array Controllers
 - **Support up to 240 disks**
 - **Capacity support up to 12.3TB in single cabinet**
- Centralized, unobtrusive manageability



Significantly Higher “Utilization” of Purchased Capacity

- Up to twice the typical 40-50% Open Systems utilization . . . Based upon dynamic pool/LUN expansion, etc.

- Importance to Business/IT:

- Customer minimizes purchase of unusable capacity
- Just-In-Time capacity increments for application growth
 - Even to the point of adding one disk-at-a-time
 - Dynamic Pool/LUN Expansion (w/Server support)
- No “droop” effect in performance
- No intensive storage administration “gyrations”

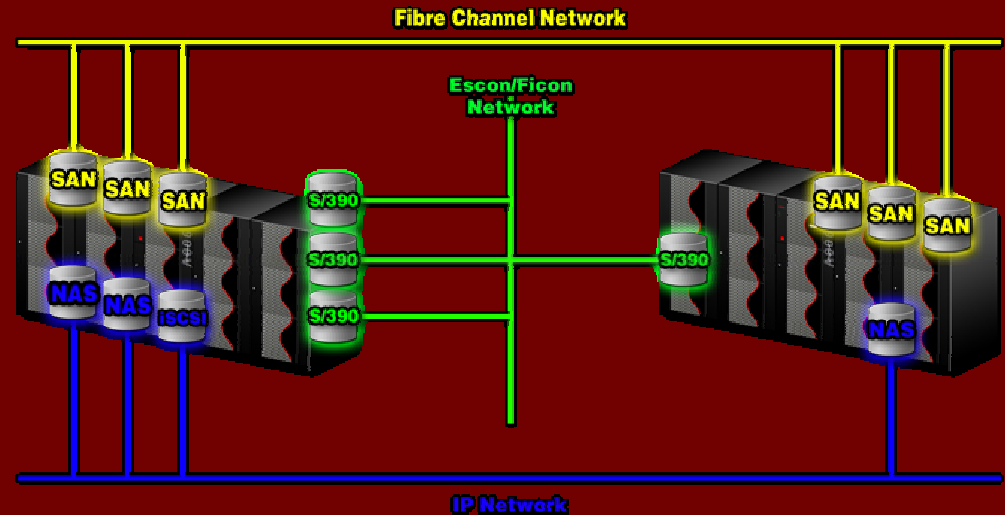
- BOTTOM LINE:

Much lower “effective” price/MB . . . Easily “justifiable”

Simplify with XP Series

■ Storage that is easy:


- Who would think a single device could do so much?
 - Heterogeneous connectivity
 - Convergence of data types
 - Multi-Protocol flexibility
- Multi-dimension scalability
 - Capacity, performance, connectivity, workload mix
- Eliminates hardware and software complexity
 - Less stress, reduced expense
- Self-managing
 - Automated quality of service



XP Consolidation


- **Host Storage Domains**
 - **Eliminates SAN Islands (Host types)**
 - **Less physical complexity & network cost**
 - **Enables much larger TB/unit**
- **FICON, PAV and Mod 27**
 - **Eliminates ESCON Bottlenecks**
 - **Larger S/390 systems**
- **Performance Maximizer (QoS)**
 - **Allows consolidation without compromise**
 - **Reduces management**
- **Multi-Protocol support for ubiquitous access**
 - **Fibre Channel, FICON, ESCON, others in the future**
 - **Provides consolidation point for all storage**

Fibre Channel High Performance Disk Arrays

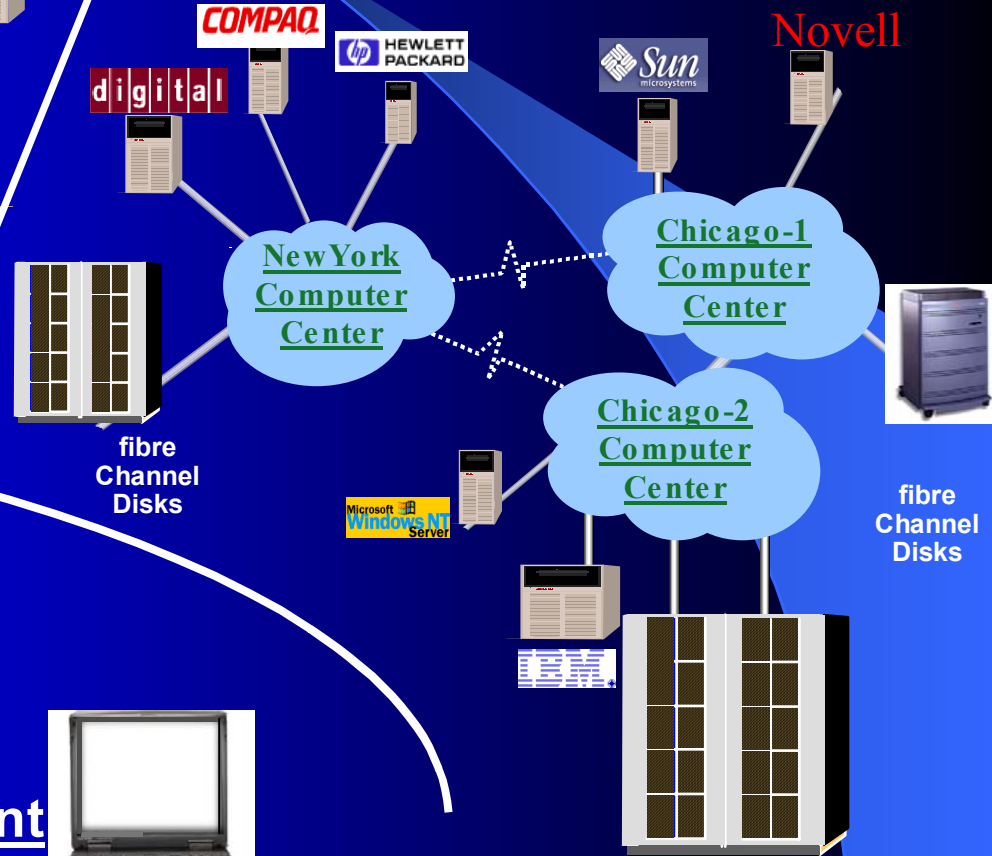


The Future - Open Virtual Storage

SAN Wide Virtual View



Actual Geography



Individual System View

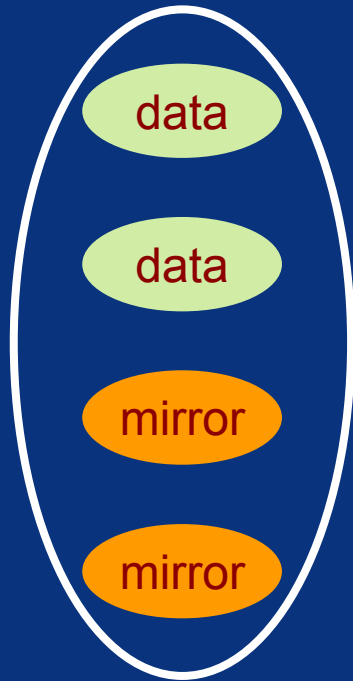


Global Virtual Management



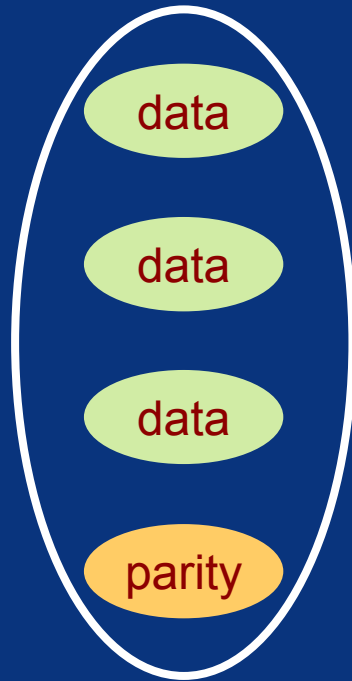
RAID Implementations

RAID types




RAID 1 2D+2D

50%




RAID 5 3D+1P

75%



RAID 1 4D+4D

50%




RAID 5 7D+1P

87.5%


storage efficiency, usable/raw

RAID 0+1

- Striping First plus Mirroring
 - Dual Pair RAID group configuration
 - Higher performance in very random environments
 - 32 Slots pre-fetched for sequential access




advanced data guarding



- similar to RAID 5, data and parity is distributed across all drives
- the capacity equal to two drives is reserved for two sets of parity data (RAID 5 distributes one set of parity data)
- RAID ADG can withstand two simultaneous drive failures without downtime or data loss
- supports online RAID level migration from RAID 1 or RAID 5


SAN RAID Design Before Analysis






Most Servers Dual Homed HBA's (Connections Not Shown)




Linux In-Band Virtualization
1 HBA to Clients, 1 HBA to Disks, 1 HBA Spare



Linux In-Band Virtualization
1 HBA to Clients, 1 HBA to Disks, 1 HBA Spare



Shelf L0 -- Dual Fiber Loop With RAID-5 Sets Using 72GB Disks (Note - All RAID Set Disks Physically Adjacent)




LUN L2 Three Oracle Databases w Redo Logs on 4 Disk RAID-5 Set

LUN L4 MUMPS IDX Database (CISR50 128GB) 4 Disk RAID-5 Set

LUN L6 Backup EMAIL Database (86GB) 4 Disk RAID-5 Set

Shelf L1 -- Dual Fiber Loop With RAID-5 Sets Using 72GB Disks (Note - All RAID Set Disks Physically Adjacent)




LUN L3 MUMPS IDX Database (CISR50 128GB) 4 Disk RAID-5 Set

LUN L5 One Oracle Database (64 GB) w Redo Logs & Prod EMAIL Database (120GB) on 4 Disk RAID-5 Set

LUN L7 Not Used

SAN RAID Design After Analysis

Most Servers Dual Homed HBA Connections. (Servers and connections not shown)




LUN L2 Three Oracle Databases w Separate Redo Logs on RAID-1+0 Sets

LUN L4 MUMPS IDX Database (CISR50 128GB) RAID-1+0 Sets

LUN L6 Backup EMAIL Database (86GB) RAID-1+0 Sets


Shelf L1 -- Dual Fiber Loop With RAID-1+0 Sets Using 72GB Disks




LUN L3 MUMPS IDX Database (CISR50 128GB) RAID-1+0 Sets

LUN L5 One Oracle Database (64 GB) w Separate Redo Logs & Prod EMAIL Database (120GB) on RAID-1+0 Sets

Shelf L2 -- Dual Fiber Loop With RAID-1+0 Sets Using 72GB Disks



Shelf L3-- Dual Fiber Loop With RAID-1+0 Sets Using 72GB Disks



When ***NOT*** To Use RAID-5

- In general, RAID-5 disk configuration is not recommended for use with most relational databases (Do Not USE with Oracle databases)!
- Write performance suffers very significantly compared to a RAID 1+0 (mirroring and striping) configuration (checksum Calculation & Update Degrades I/O Performance!)
- Time window availability notwithstanding significant increases in volume of data will probably cause total elapsed time to exceed acceptable time limits.

RAID 1+0 Or RAID 0+1

In

RAID 1+0 Or RAID 0+1?

- RAID 0+1 configuration disks striped together into sets, followed by the mirroring of the sets.
- RAID 1+0 configuration drives are mirrored followed by the striping together of the resultant mirror sets.
- RAID 1+0 better I/O (read and write)
- Performance increase because of simultaneous I/O across various spindles.
- RAID 1+0 **MORE** expensive (you get what you pay for !)

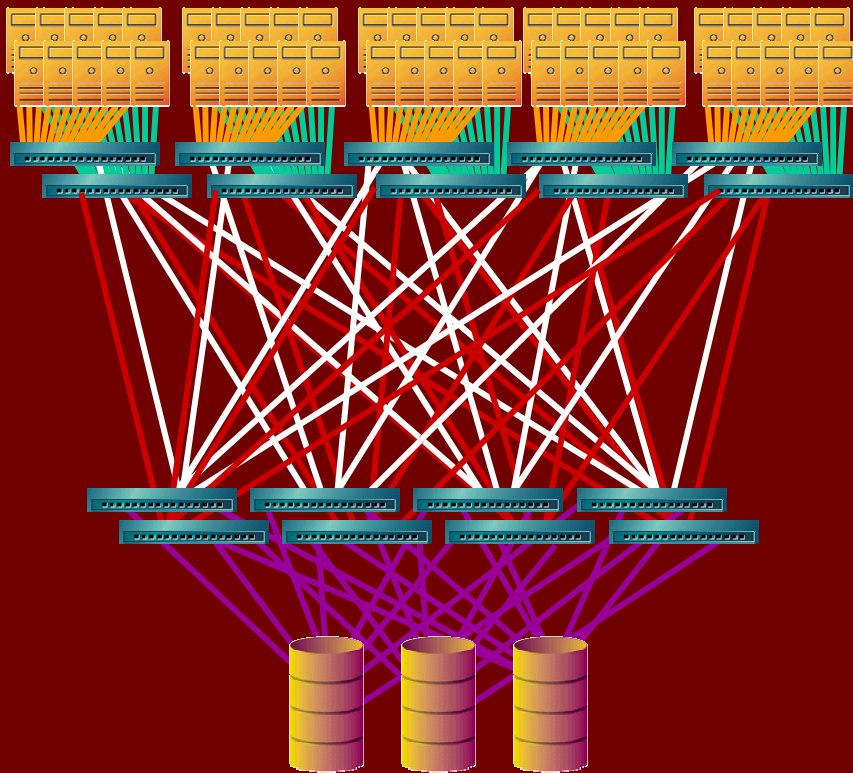
SAN Fabric

Fabrics

- Scalability means more than adding ports
 - Supporting data growth while sustaining infrastructure performance, availability, and management capability
- Enable scalability
 - Design fabric that supports key environment attributes
 - Examples: Data Store Size and location, Backup, Distance, Applications
- Fabric design drives product selection
 - Switched Infrastructure
 - Building block size
 - Availability
 - Link performance
 - Management
 - Device to Fabric-wide


Large Fabric Design

Switch Fabric



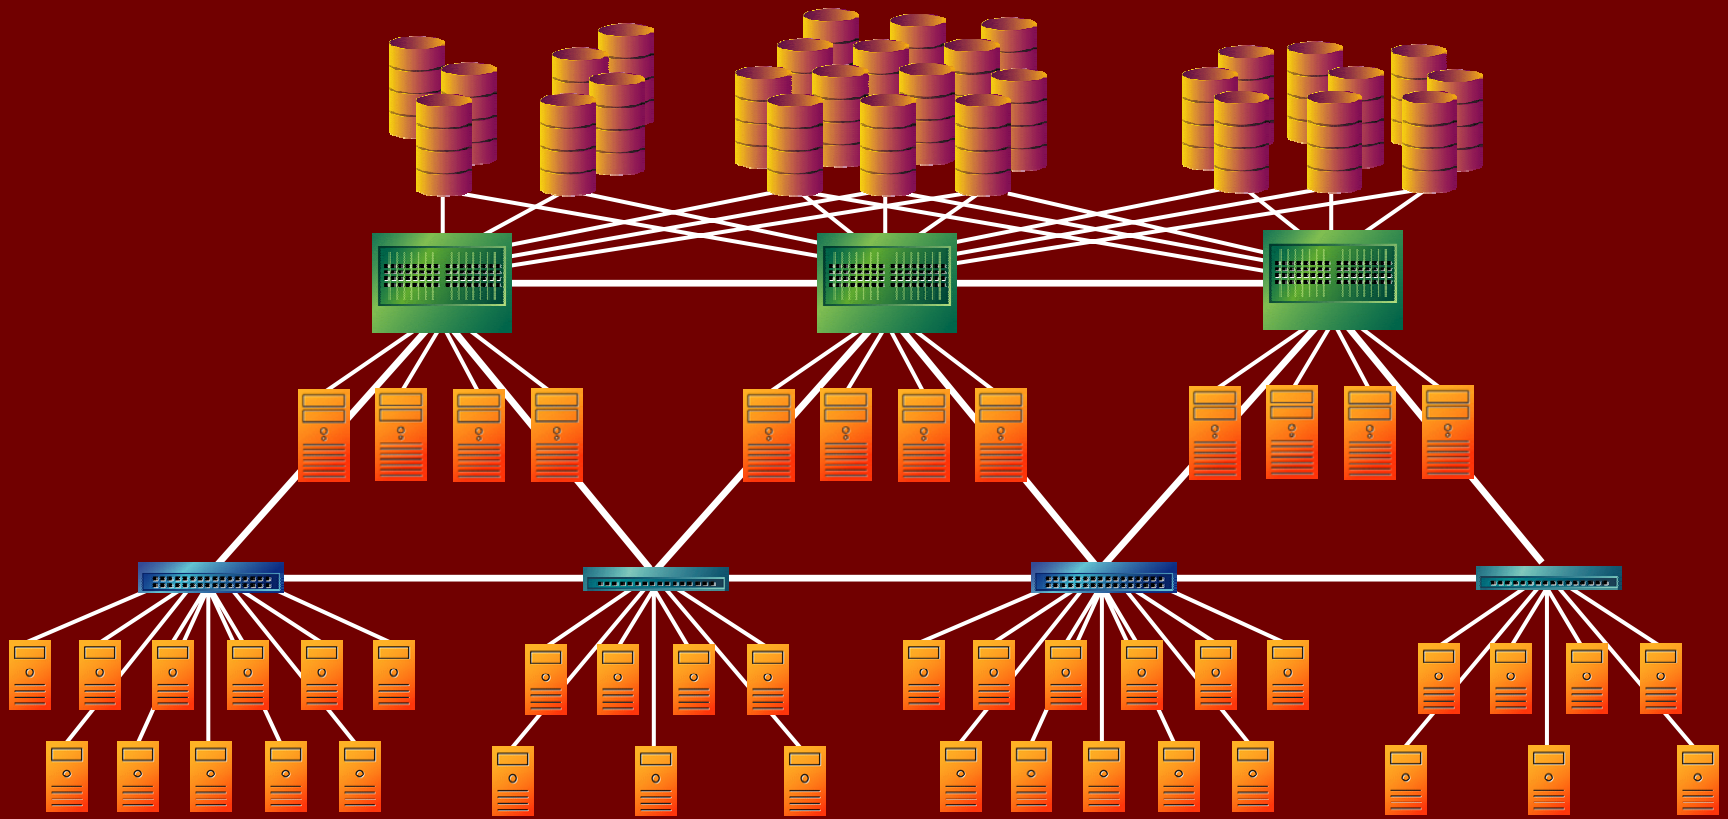
Enterprise Storage

Director Fabric



Enterprise Storage


Fabric Architecture Design



Directors and Switches

- **Availability *Depends On Vendor***

- Director: ~99.999% downtime < 5 minutes / year
- Fabric Switch: ~99.9%
 - downtime 8.8 hours / year




- **Any-to-Any Connectivity**

- Director: large port count
- Fabric Switch: small to medium port count



- **Serviceability**

- Director: non-disruptive
- Fabric Switch: potentially disruptive




- **Scalability**

- Director: better fabric scaling
- Fabric Switch: Use Multiple Switches


Improve Storage Utilization

DAS



- Limited scalability and bandwidth
- >50% of storage maybe unused
- Admin costs 5-6 times acquisition cost

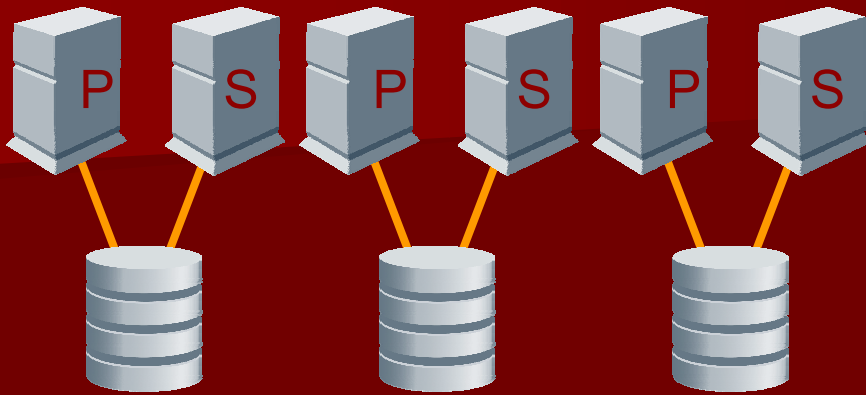
SAN



- Reduce storage headroom
- Enable just-in-time provisioning
- Combine UNIX and NT storage growth needs


Improved Application Availability

DAS



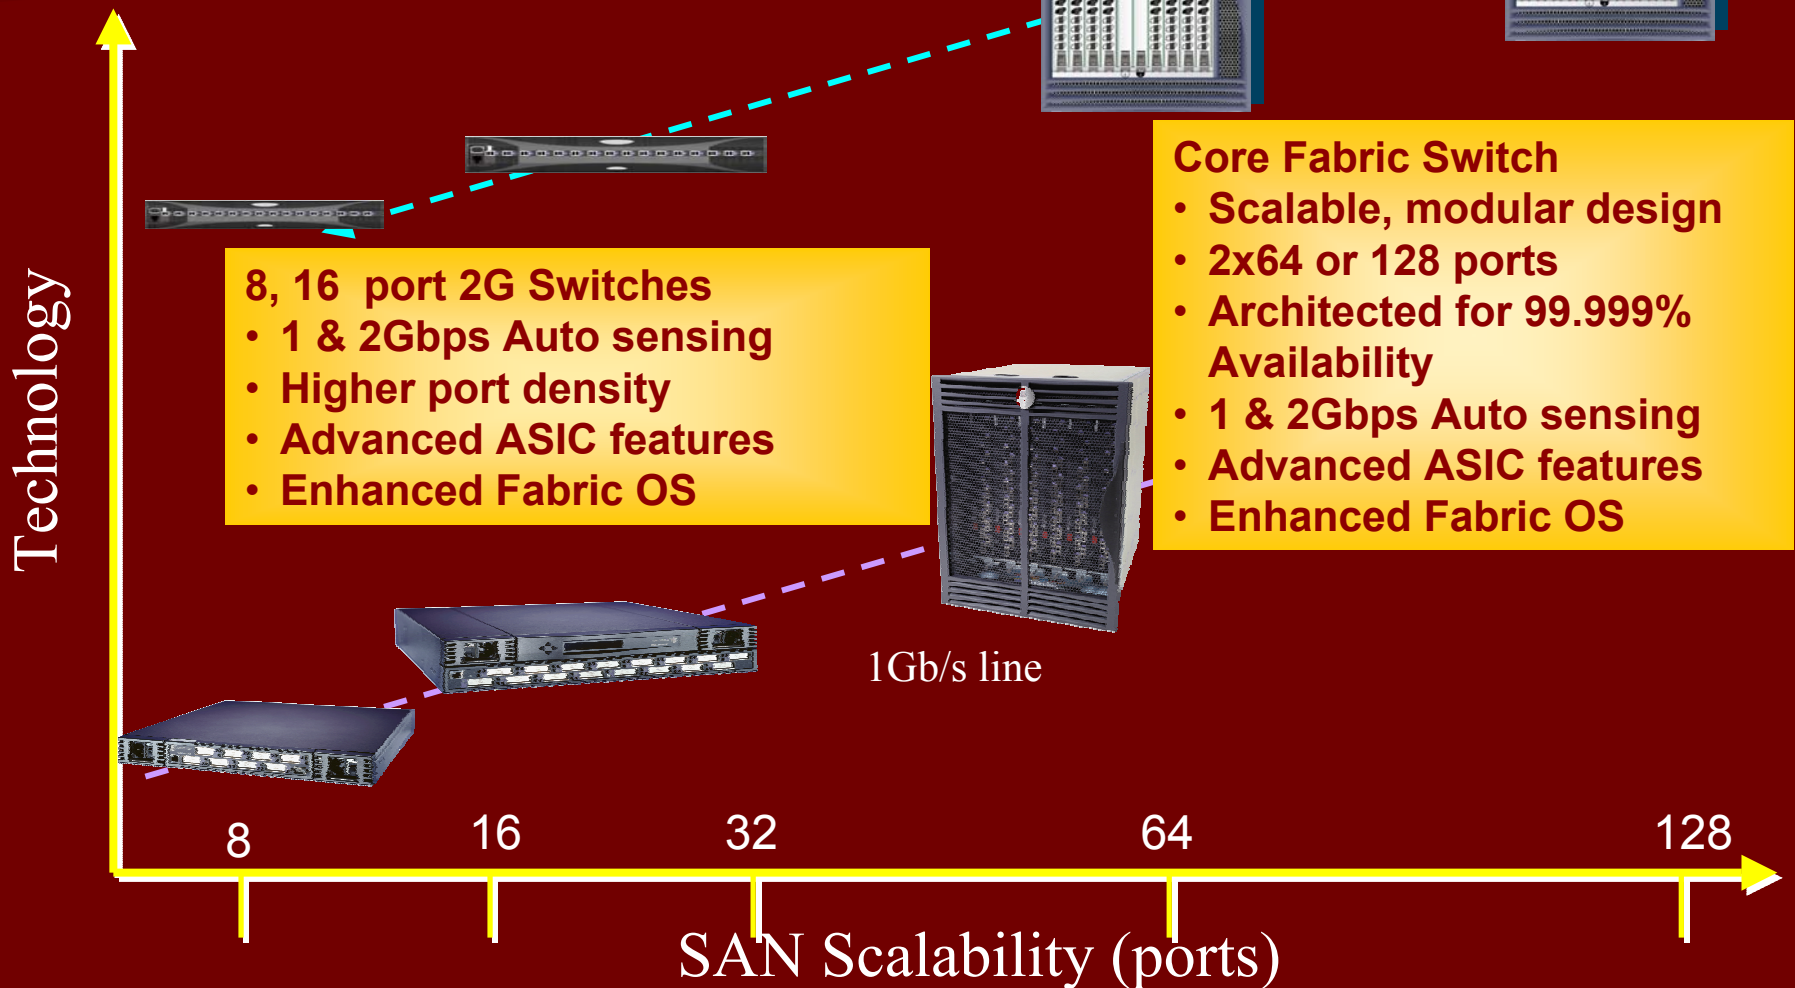
- 3 High Availability Clusters
- One standby server for every primary server
- Data isolated from each cluster

SAN



- Much less expensive to implement high-availability
- Storage can be consolidated
- One standby server (N+1)
- Data available to all servers

2Gb/s Product Line

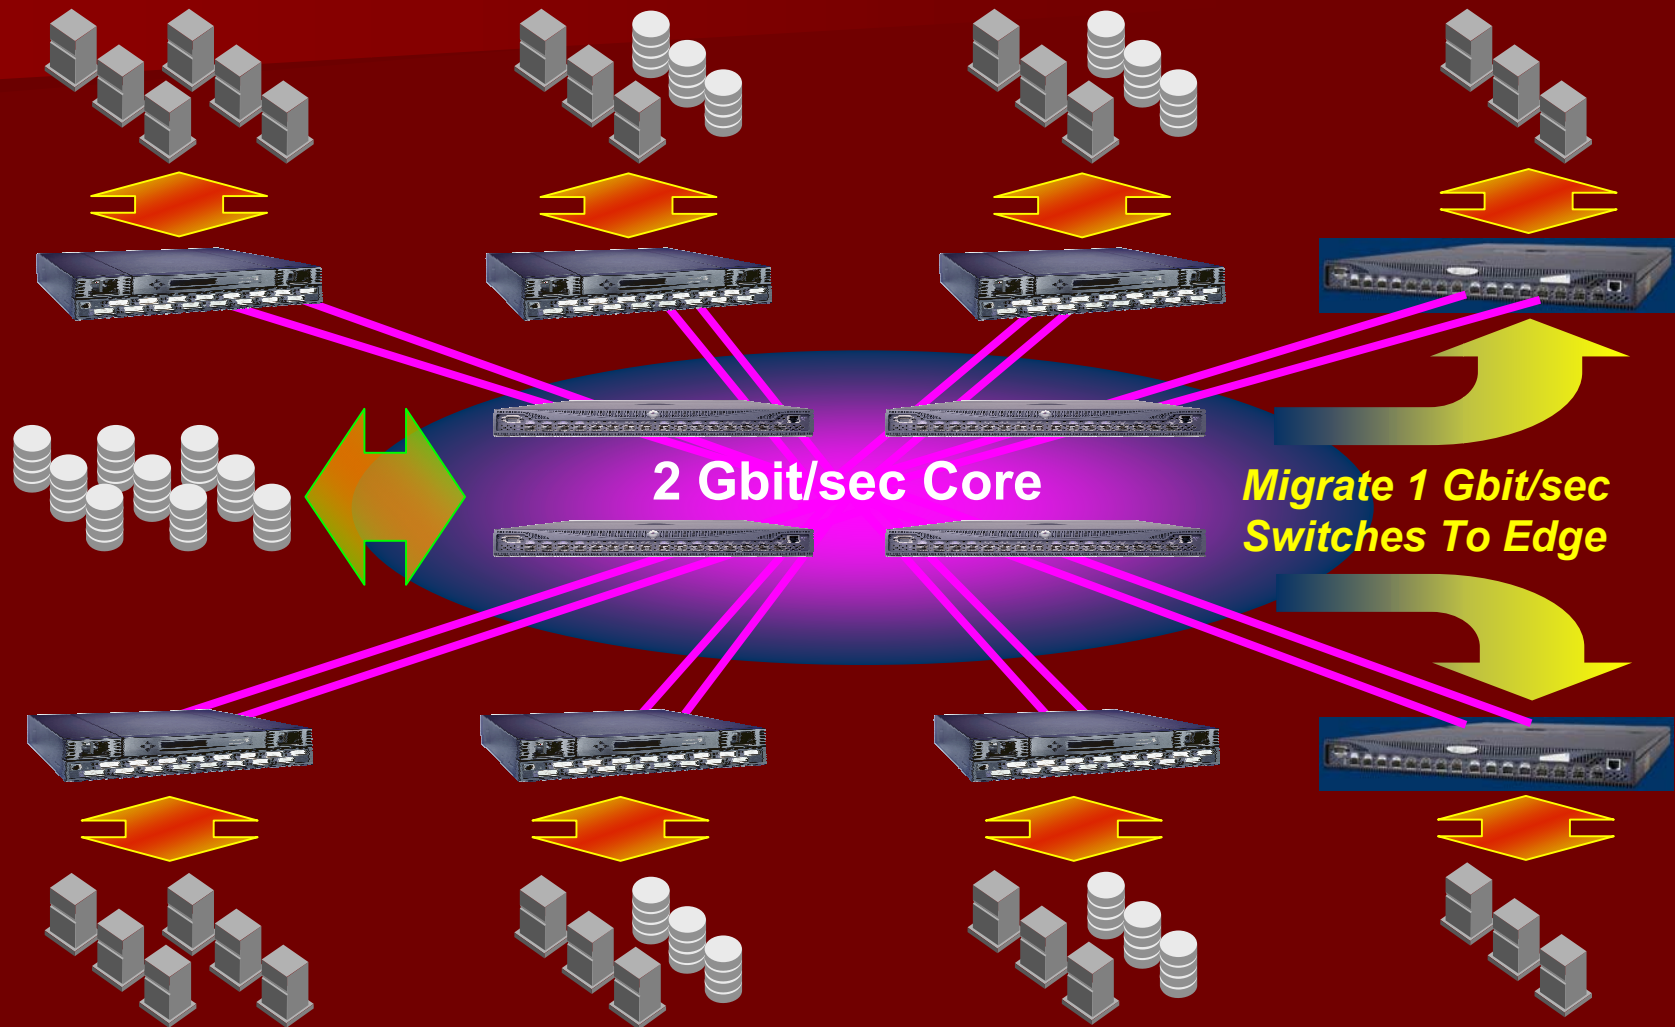


SilkWorm 12000 Core Fabric Switch

- Flexible, modular architecture
 - Scalable 64/128 port design
 - 2 Gbit/sec ports; Autosensing
 - 3rd generation Brocade ASIC
- 99.999% availability
 - Redundant, hot-swap elements
 - Non-disruptive software updates
 - Redundant 64 port switch config
- Intelligent fabric services
 - Interswitch link trunking
 - Frame filtering
 - Global performance analysis
- Multi-protocol architecture
 - 10 Gbit/sec fibre channel
 - IP storage interconnect
 - InfiniBand



Compatibility: Ensuring Investment Protection



Fabric Services

Current fabric services:

- Foundation fabric services
 - Device discovery, routing, and name services
- Advanced fabric services
 - Loop emulation, SAN internetworking, hardware-enforced zoning, health monitoring

New fabric services:

- Inter-switch link trunking
- Advanced zoning
 - HW WWN
- Global performance analysis

Benefits

- ✓ Simplifies network design and administration
- ✓ Improves performance
- ✓ Improves security
- ✓ Lower total cost of ownership
- ✓ Enhanced high availability


Other Fabric Services

- Inter-switch link (ISL) trunking
- Advanced zoning
 - HW WWN
- Global performance analysis


Fabric Services

Current Services

- Distributed name server → Auto-discovery → Zero administration
- FSPF routing → Scalability/cascading → No single points of failure
- Zoning → Hardware and software enforced → Increased security
- Quickloop/
translative mode → Legacy device and loop support → Investment protection
- Extended fabrics → Extended fabric support → Spans up to 120Km




Inter-Switch Link Trunking



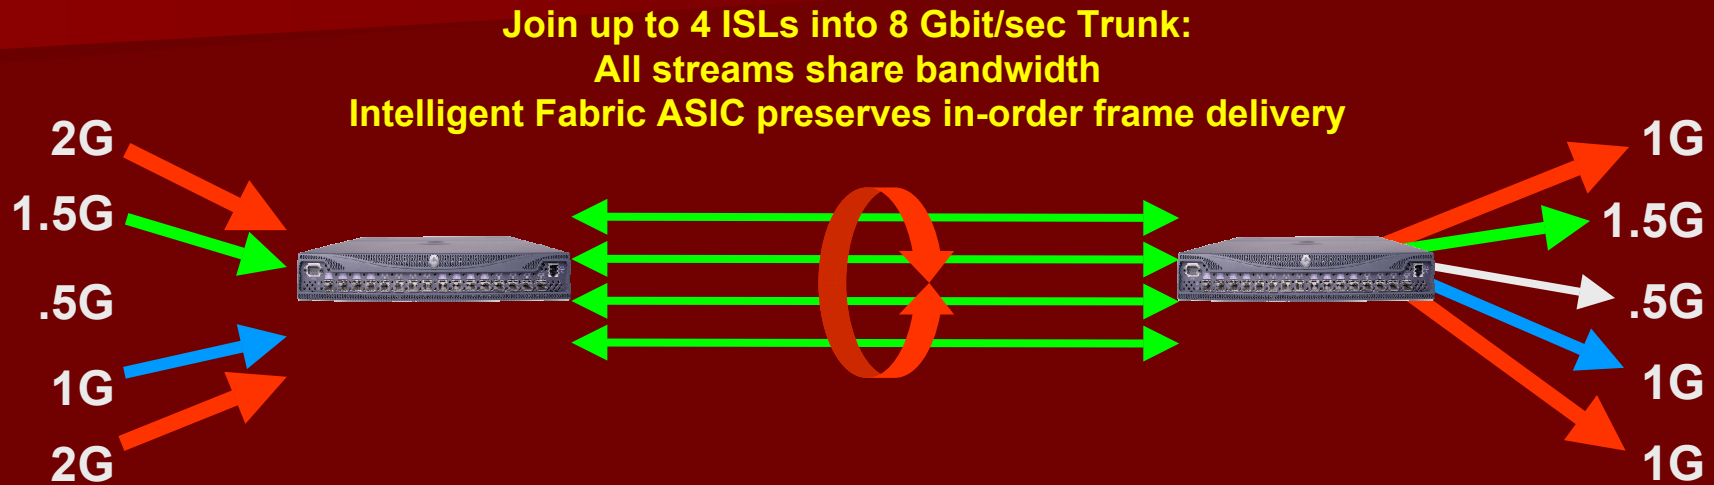
- ISL Trunking
 - 8 Gbit/sec logical links
 - Aggregate edge traffic
 - Zero management
- Simplify network design and management
- Simplifies ongoing administration (manage one link versus four links)
- Maximizes fabric performance
- Provides increased high availability in case of link failures

Intelligent Bandwidth Utilization: Dynamic Load Sharing



- Load sharing across multiple ISL links
- Round robin assignment
- Can get “unlucky” with multiple high utilization traffic assigned to same link
- In our example, theoretical maximum is 8 Gbit/sec, but effective throughput is 5 Gbit/sec

Intelligent Bandwidth Utilization: Inter-switch Link Trunking



- Aggregate traffic onto fewer logical links
- Automatically created when switches are connected
- Managed as a single logical 8 Gbit/sec ISL
- Fault-tolerant – will withstand failure of individual links
- Supports redundant trunks between switches

Multi-Fabric SANs

Overview

- What are multi-fabric SANs?
- When should they be used?
- How should they be implemented?
- How can they be managed?
- Is there a long-term multi-fabric SAN strategy?
- Question & Answer; Session Wrap-up

What are Multi-Fabric SANs?

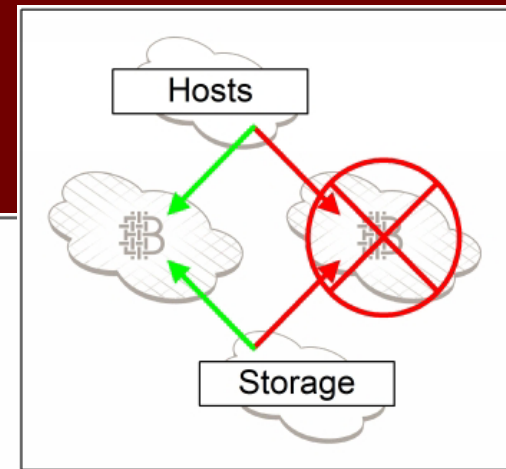
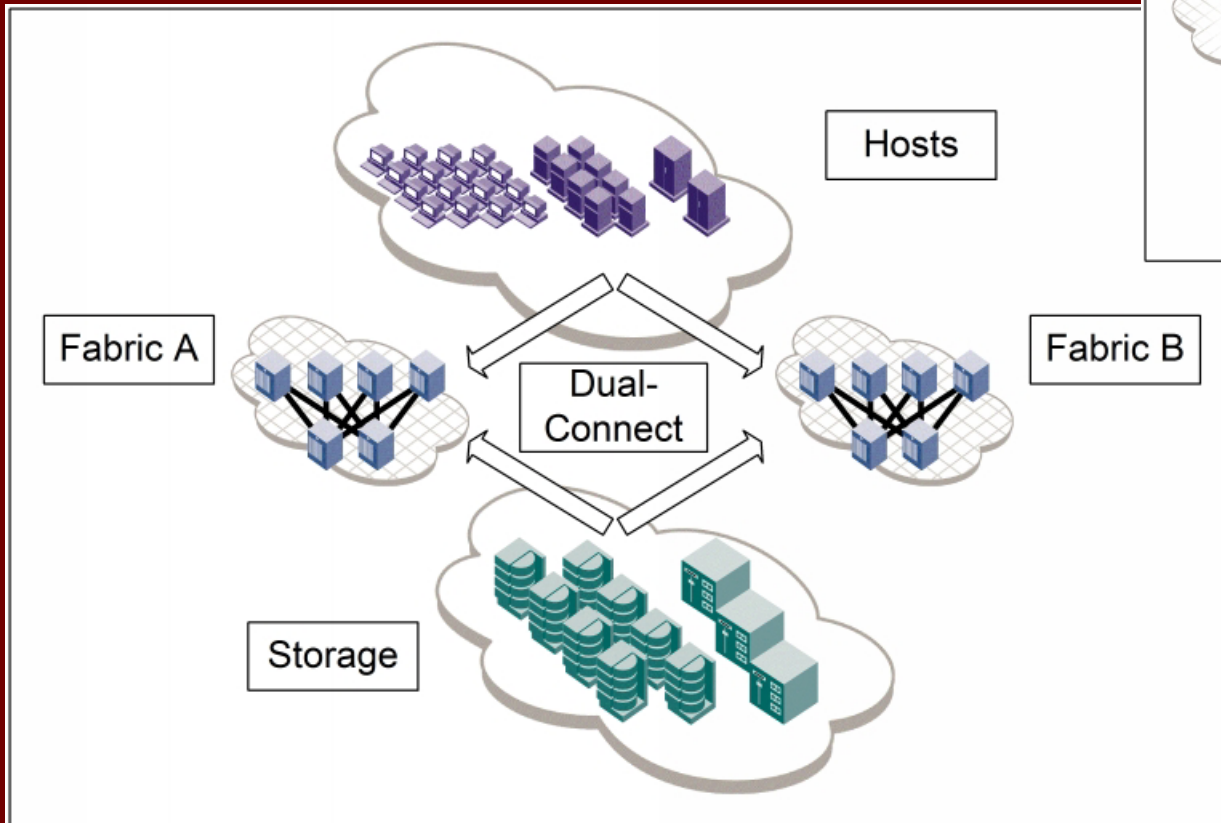
SAN: A Storage Area Network can consist of one or more related fabrics and the connected nodes.

Multi-Fabric SAN: A SAN consisting of more than one fabric that provides one connectivity solution


- If there is only one fabric and one connectivity solution, “multi-” doesn’t apply
- If there is more than one fabric, and *many* connectivity solutions, use multi-SAN, not multi-fabric
- **NOTE:** MF SANs cost the *same* per-port as single fabric SANs. Cost is *not* a factor in the design process.
2x96=1x192

Dual Fabric SAN (HA)

- Used for redundancy
- “One of *anything* is not HA”




Fabric Extension




*Used to specify the speed of fiber optic networks. The base rate (OC-1) is 51.84 **Mbps**. OC-2 runs at twice the base rate, **OC-3** at three times the base rate (155.52 Mbps), etc. Planned rates are: OC-1, OC-3, OC-12 (622.08 Mbps), OC-24 (1.244 Gbps), and OC-48 (2.488 Gbps)

Zoning

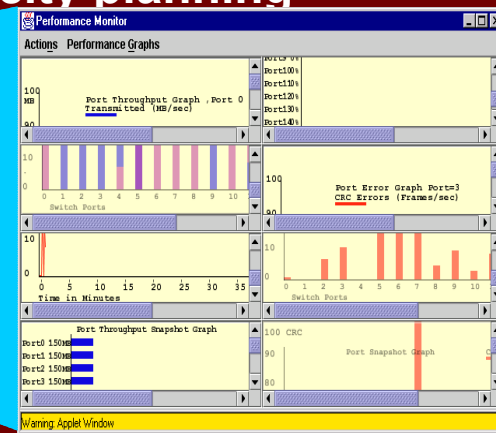
- Provides OS & storage isolation
- Store multiple zone configurations
- Zones based on port # or device WWN
- Updates distributed dynamically across the fabric
- Overlapping zones allowed
- No logical limit on the # of zones



Performance Analysis



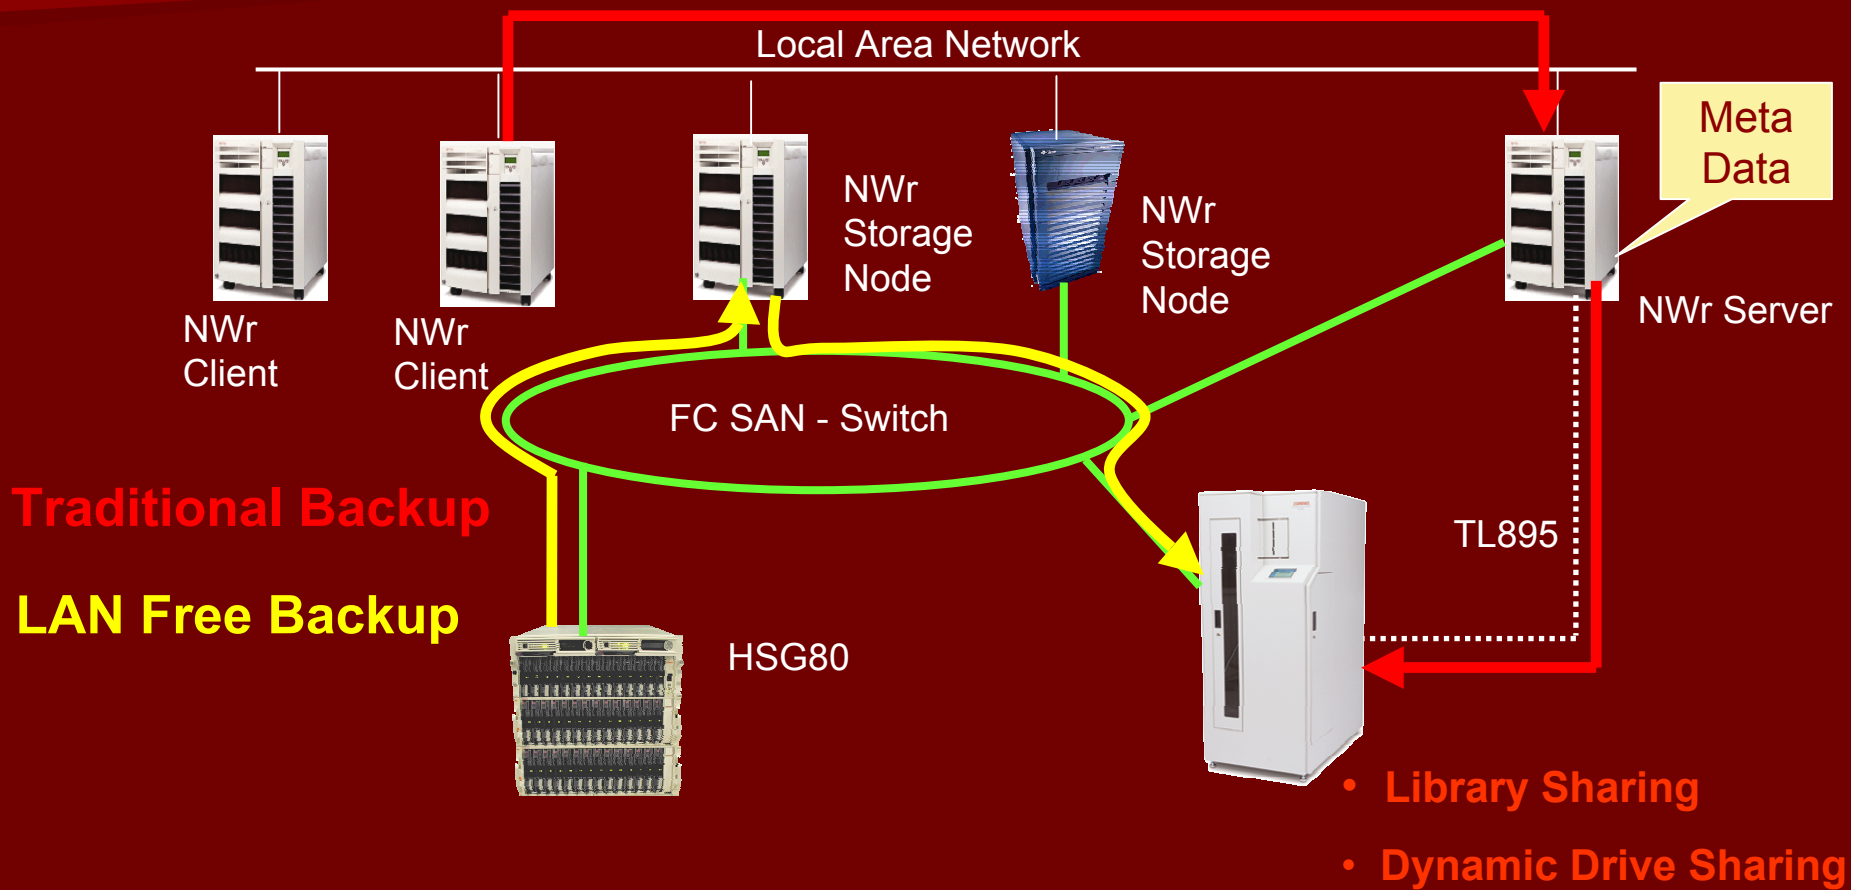
- Performance analysis
- Proactive SAN management
- Provides end-to-end performance measurement from the source to the destination target
- Optimize fabric resource allocation
- Maximizing performance tuning
- Reducing trouble-shooting time
- Improve capacity planning



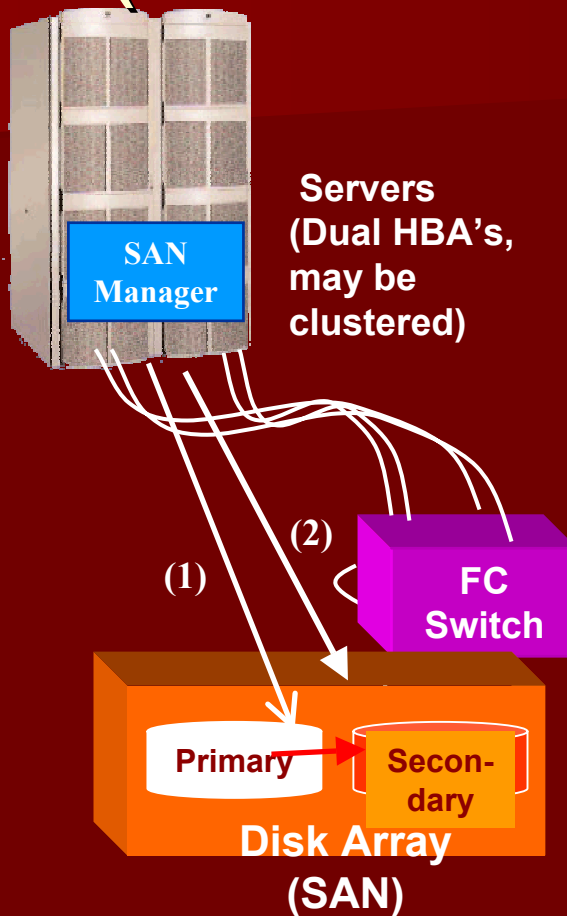
SAN Based Backup/Restore and Business Continuity

High Availability With Storage
Area Networks

LAN Free Backup Storage Node / SAN Storage Node




Back Up Critical Data (Clone/Shadow)



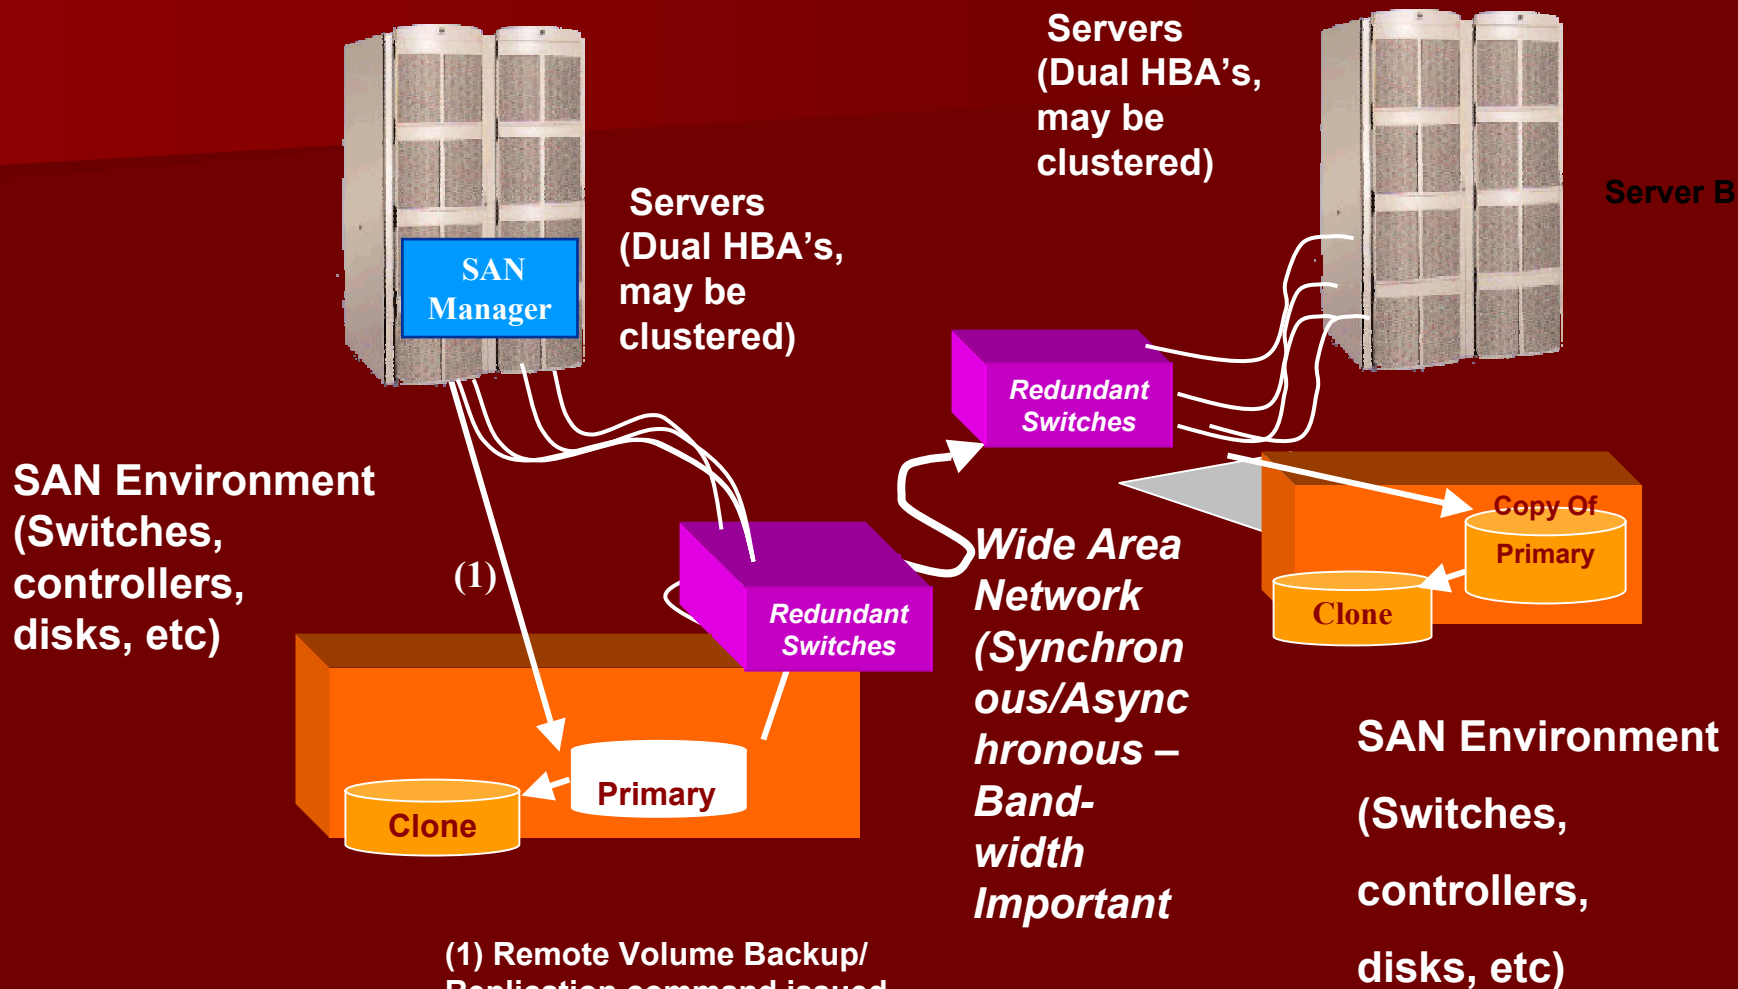
Servers
(Dual HBA's,
may be
clustered)

- (1) Primary Disk is Cloned(SAN) or Shadowed(Servers). Primary FULLY copied to Secondary. If Shadowing New write data stored into both primary and secondary volumes.
- (2) Secondary Split off from primary and then backed up to tape

Point-in-Time Copy Enterprise Volume Manager



Wide Area Back Up Of Data



(1) Remote Volume Backup/
Replication command issued.
Executing remote backup
should be completely SAN
based (no server involvement)

What Does Downtime Cost

Downtime costs major Internet players about \$8K per hour

Forrester Research Inc.

Downtime costs \$1,400 per minute on average

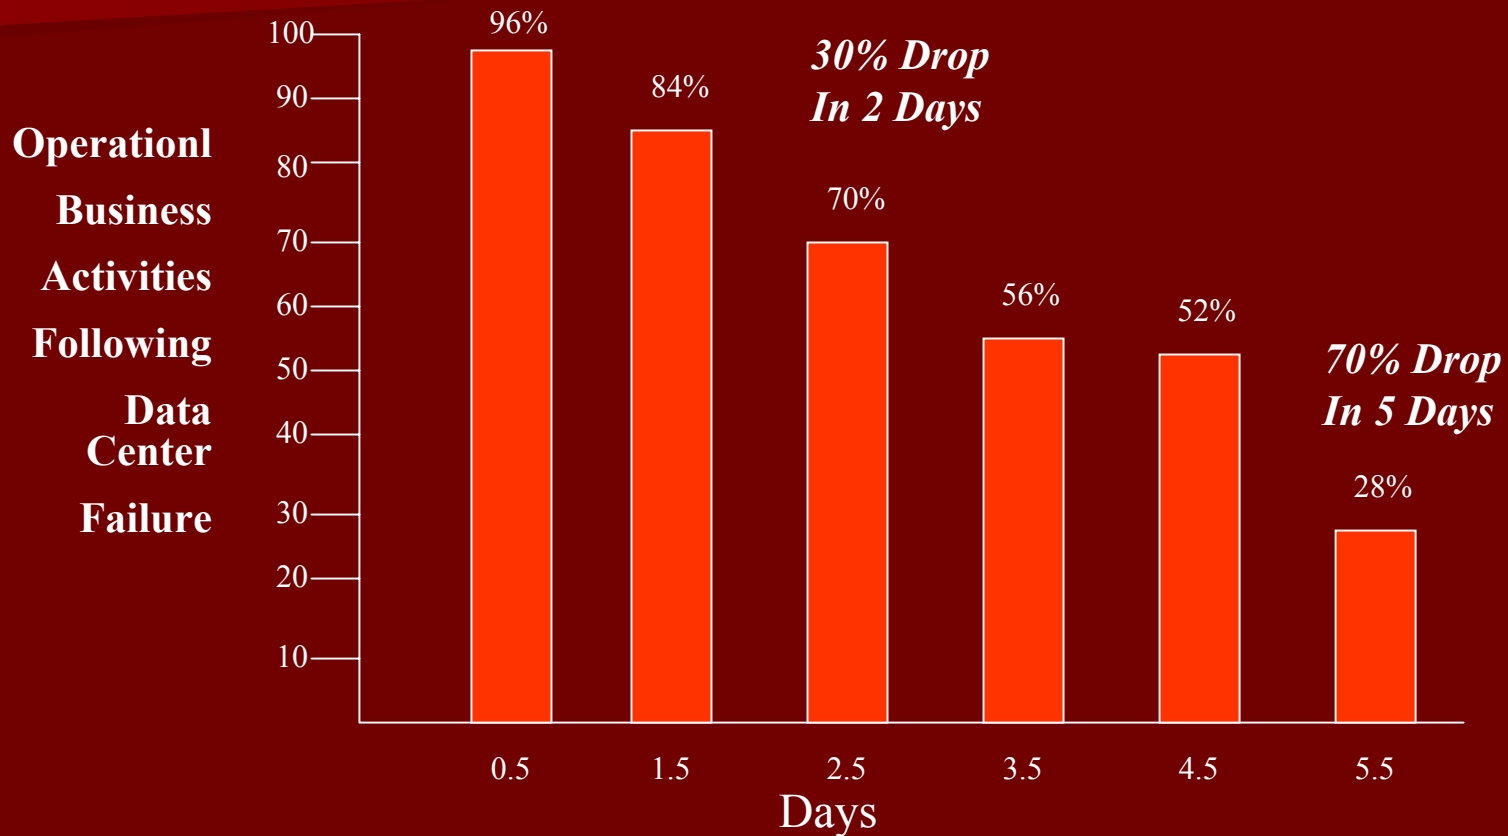
Oracle

Downtime costs medium-sized businesses > \$100K/hour
which is > \$1 Million annually

IDC,
Corp.

<u>Business</u>	<u>Average Loss</u>
Retail Brokerage	<i>\$6.45M/Hr</i>
Credit Card Sales Authoriz	<i>\$ 2.6M/Hr</i>
Home Shopping Channels	<i>\$113,750/Hr</i>
Airline Reservation Cente	<i>\$89,500/Hr</i>
Package Shipping Service	<i>\$28,250/Hr</i>

Loss of Capability to Perform Basic Business Functions



No longer just a computer disaster after 1 or 2 days

The Effects of Data Loss

- When Companies lose data in a disaster, 50% never reopen and 90% are out of business in two years
 - Source: University of Texas
- 93% of companies that suffer a significant data loss are out of business within five years
 - Source: U.S. Bureau of Labor
- 43% of U.S. businesses never reopen after a disaster and 29% close within two years
 - Source: University of Wisconsin
- 30% of computer users say they spend the equivalent of one week per year reconstructing lost data.
 - Source: 3M Corporation

HP SAN's Virtualization Concepts (Differences and Considerations)

Virtualization Technology

Virtualization: Storage

- **Enables vast amount of physical capacity scattered across the enterprise to appear as a single large pool of data for presentation to application servers.**
 - **Storage pool is a single manageable entity**
 - **Virtualization can occur at any of three levels.**
 - **All levels can be used together, or independently, to maximize the benefits to customers:**
 - ***Server level:* Ideal for small entry level needs**
 - ***SAN Fabric level:* SAN-wide virtualization**
 - **increase efficiency in SAN development, management, and service**
 - ***Storage System level:* Ideal for large volumes without sacrificing performance or reliability**

Why is Virtualization Important Anyway ?

- Virtualization will:
 - Provide investment protection for non-virtualized arrays
 - Provide a Common Storage Pool that will allow for greater efficiency and not a load of stranded capacity in servers or on SAN RAID volumes/arrays
 - Data Migration that is transparent to the OS for upgrades, changes, etc to dataset for applications
 - Reduction of ownership costs

EVA Virtualization Primer

- Principles of EVA virtualization
 - “Disk blocks are disk blocks”
 - All disk blocks are equivalent*
 - Foundation object is the “Virtual Disk”
 - A SCSI LUN
 - Attributes
 - Data
 - It doesn’t matter where or how data actually stored*

EVA Virtualization Primer (con't)

– Data location independence

- Makes new benefits possible
- BUT - to realize those benefits, we must change some old habits
 - Backend disk access pattern no longer correlated with host access pattern
 - Stop using configuration as “organizing” tool
 - Use EVM capabilities instead
 - Resist temptation to micromanage
 - Avoid temptation to utilize every last byte

Significantly Higher “Utilization” of Purchased Capacity

- Up to twice the typical 40-50% Open Systems utilization . . . Based upon dynamic pool/LUN expansion, etc.

- Importance to Business/IT:

- Customer minimizes purchase of unusable capacity
- Just-In-Time capacity increments for application growth
 - Even to the point of adding one disk-at-a-time
 - Dynamic Pool/LUN Expansion (w/Server support)
- No “droop” effect in performance
- No intensive storage administration “gyrations”

- BOTTOM LINE:

Much lower “effective” price/MB . . . Easily “justifiable”

StorageWorks Enterprise Virtual Array

● HSV110 Array Controller Virtualization:

- All raw storage is pooled
- Virtual Disks are drawn from a pool
- Virtual Disks managed by customer to these constraints:
 - Size range 1GB - 2TB, in 1GB increments
 - Up to 256 Virtual Disks - selectively presented to hosts
 - Each Virtual Disk can have no, medium or high redundancy
 - RAID V0, V5, V1
 - A Virtual Disk's size can be dynamically expanded, but not shrunk

What is StorageWorks Virtualization?

- Virtualization will bring new terms with it
 - Aggregation
 - Fine Grain Mapping
 - Distributed RAID
 - Capacity Free Snapshots
 - Storage Pools (Local & Global)

Strategic Storage Virtualization Views

Server Level

Local Storage Pools; Direct Access to Virtual Disks; Network Mapped Disks Served to other Clients; Local Snapshot

Fabric Level


Virtualization Within Switch

SAN Level

Global Storage Pools; Direct Access to Virtual Disks; Attribute Based Storage; Multiple Models for data delivery: Asymmetric & Symmetric; Builds upon RAID Volumes on SAN ;

Storage Level

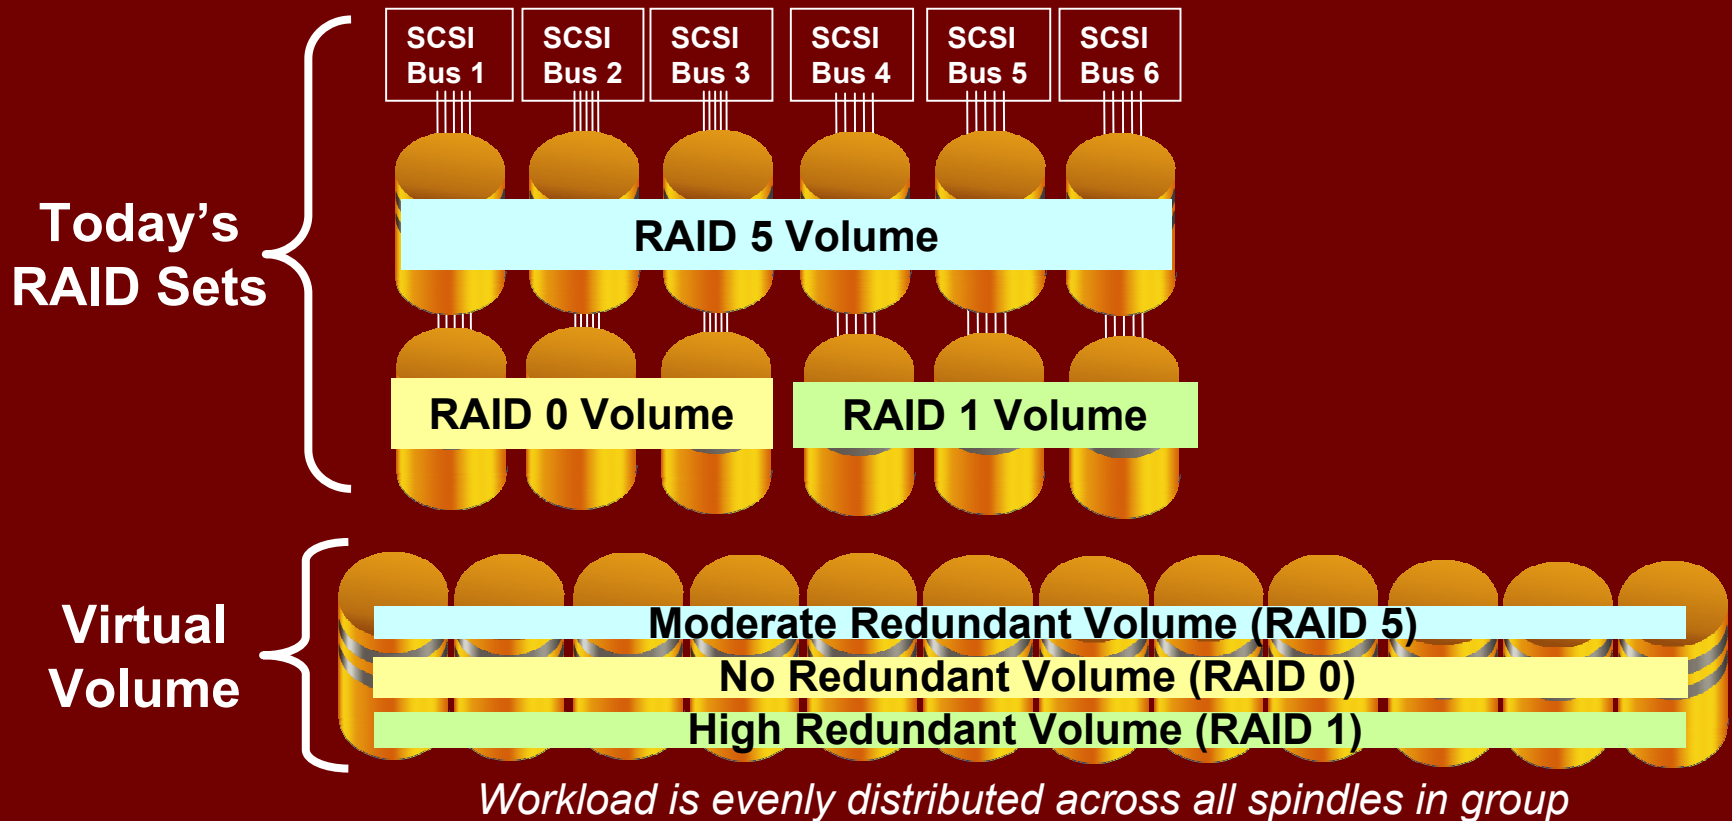
Large Local Pools; Distributed RAID; Ultra High Performance



**SAN
Virtualization
Appliance**

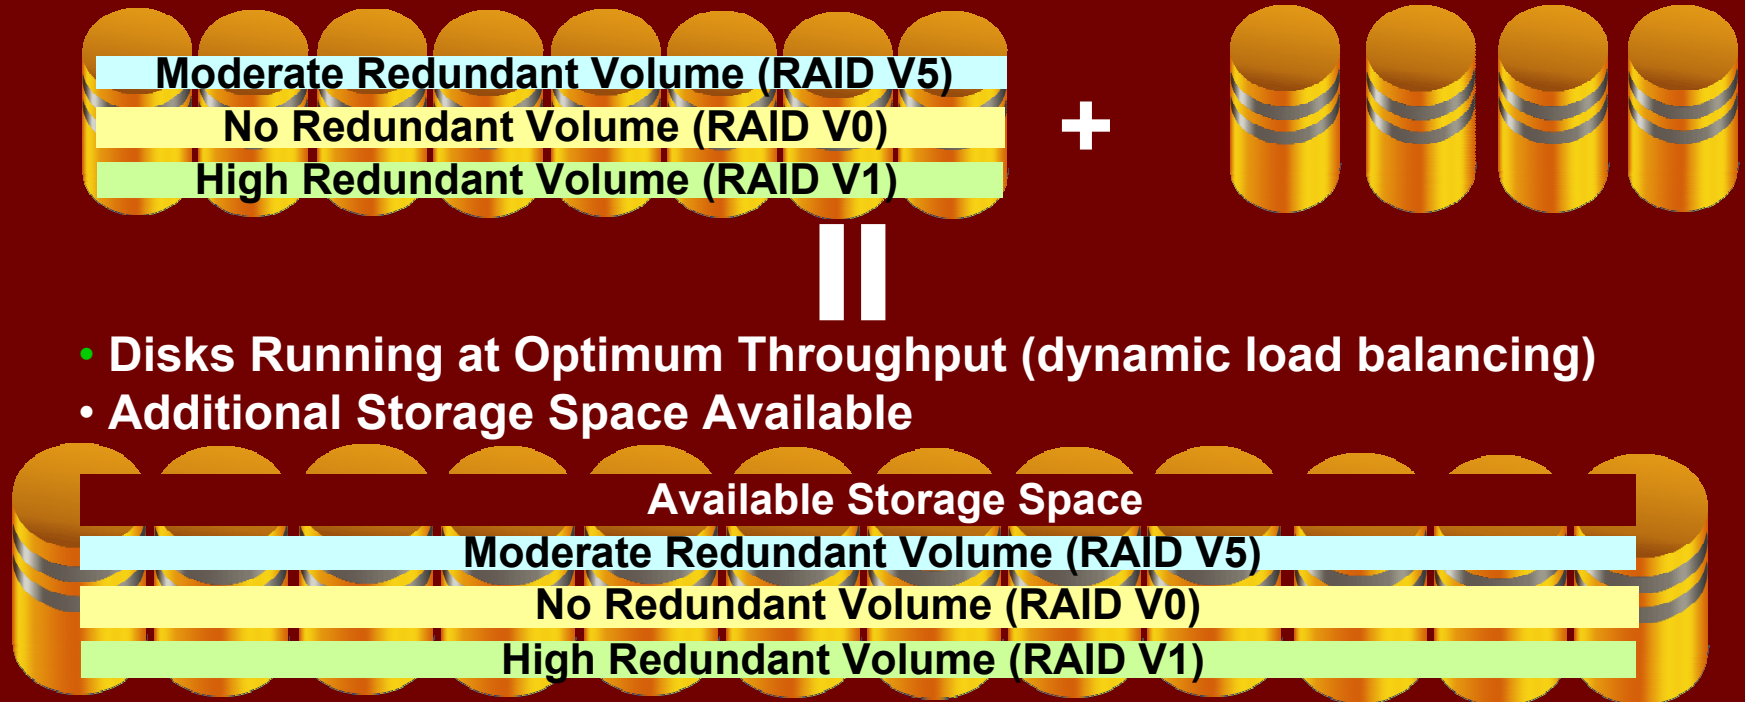
Virtualization

- Eliminate throughput bottlenecks
- Eliminate load balancing procedures for application and Data Base



Virtual Storage Pools



- Dynamic pool capacity changes
 - Pool capacity can be expanded by spindle(s)
 - Virtual Disk blocks are automatically relocated to level spindle use
-
- **Disk Spindles Becoming a Throughput Bottleneck**
 - **No Additional Storage Space Available**



- **Disks Running at Optimum Throughput (dynamic load balancing)**
- **Additional Storage Space Available**

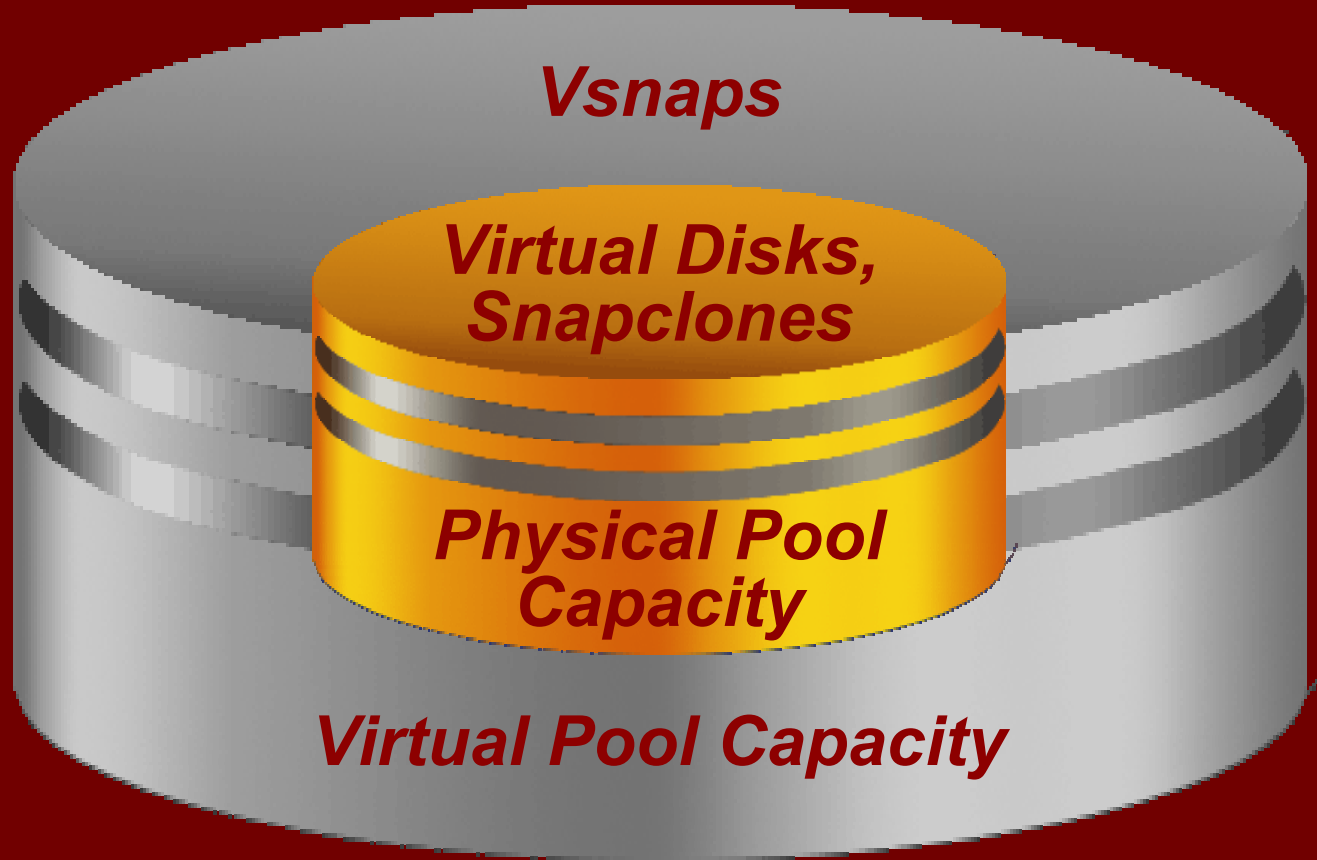
Vsnaps and Snapclones

- Double effective capacity
- Stretch virtual pool capacity for free
- Simplify data mining operations, reduce application development backlog
- Eliminate backup windows, reduce recovery time to raise service levels
- Make ad hoc point-in-time copies with minimal effort to maximize IT flexibility
- Provide administrators with powerful tools to multiply management efficiency




Stretching

Capacity with Capacity-Free Vsnaps



Virtually
Boundless
Storage


Heterogeneous Connections – The Problem



■ The Problem


- “Mode Set” determines port behavior.
- Different mode sets for Solaris, AIX, NT, etc.
- Port can be set in only one mode
- Sharing same port between different host platform types creates problems.

Heterogeneous Solutions Today



- The Solution
 - Must configure separate physical ports for each host platform type
- Consequences
 - Expense of extra ports
 - Max number of channel features, especially with FC, ESCON, FICON mix limits heterogeneous consolidation.

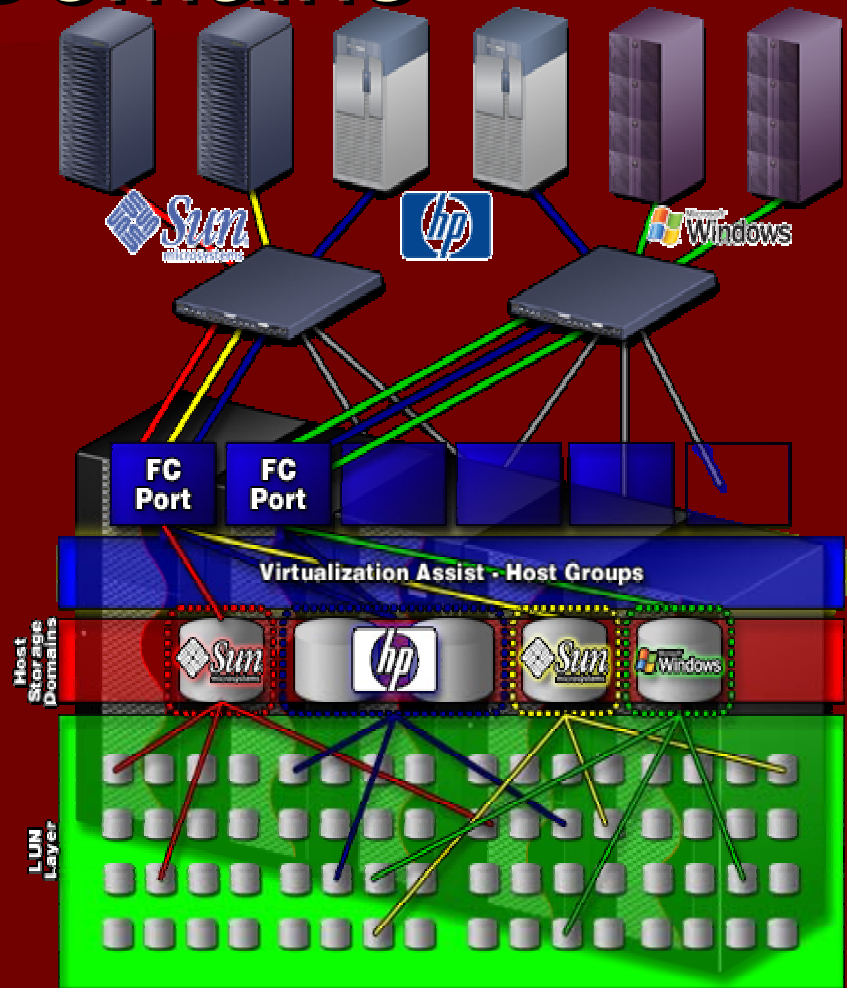
The HP XP Series Heterogeneous Solution



- **Host Storage Domains (HSD)**
 - Each HSD has its own logical (virtual) FC port.
 - FC port “mode set” applies to logical FC port in each HSD.
 - Each HSD has its own set of LUNs.
 - Hosts matched to their assigned HSD based upon WWN
 - Hosts can only “see” their own HSD

Virtualization Assist: Host Storage Domains

- Multiple Host Storage Domains can share same physical port.
- Each Host Storage Domain has its own logical FC port and its own independent set of LUNs.
 - Multiple LUN 0's
- Host connections routed to HSD based upon WWN.
- Fewer physical ports needed
 - Reduces complexity & cost
- More overall connections
- Enables consolidation



HP XP Virtualization Assist Summary

- Many of virtualization benefits immediately
 - Single span of control large amount of data
 - Optimum capacity utilization
 - 35% => 70%+
 - Storage pooling
 - By application, by host, by host groups
 - QS to manage workload priorities
 - Significant reduction in physical complexity
 - Fewer storage units
 - Fewer physical connections (storage units and switches)
- Compatible with evolving virtualization architectures

Virtualization Assist

- Host Storage Domains storage pools
 - Better granularity than port level
 - Enables fine granularity pooling
 - Security at WWN name level
 - Performance management/control WWN level
- Works with/without higher levels virtualization
 - Complements other virtualization methods
 - Delegate to lowest common denominator
 - Protects investment as higher levels evolve
 - Many of benefits now, flexibility in future

Storage-Centric Project Life Cycle

➤ SAN Project Lifecycle Phases

- Unique SAN specific project lifecycle methodology
- Details a series of step-by-step procedures or phases

Lifecycle Phases

- **Business & Data Requirements Collection & Analysis**
- **SAN Design Verification**
- **SAN Component Selection**
 - ✓ **Host Bus Adapter**
 - ✓ **Fabric Vendor**
 - ✓ **Disk Array**
 - ✓ **Integrated NAS Head (Optional)**

Lifecycle Phases (con't.)

- ✓ **SAN Layered Products**
- ✓ **SAN Component Analysis Matrix**
- **SAN Procurement & Implementation**
- **SAN Project Plan**
- **SAN Functional Testing**
 - ❖ **SAN component testing**
 - ❖ **SAN integration testing**
 - ❖ **SAN communication testing**
 - ❖ **SAN fault tolerance testing**

Lifecycle Phases (con't.)

- ✓ **SAN Functional Testing (con't.)**
 - ❖ **Network component integration and testing**
 - ❖ **Data load testing and analysis**
 - ❖ **SAN Disaster Recovery testing**

Lifecycle Phases (con't.)

- Production Preparation & System Cutover
- Performance Measurement & Requirements Verification
- Continuing Analysis Services

SAN management software checklist

Feature	Key benefit
Plug-in, or modular, architecture	Timely support of new devices
Support for new discovery protocols	Timely support of new protocols
Layered architecture	Can handle proprietary management systems
Single management console	Ability to monitor/manage the entire SAN with a single interface
Automatic launch of native software tools	Unique device vendor error reporting and correction
Association of devices with the desired	Flexibility in application launching subsystem software
Pass-through of SNMP traps (events)	Interoperability with other applications
Forwarding of events to multiple listeners	Remote monitoring
Topology import/export	Ability to quickly restore a SAN after a move
Bill of materials export	Asset management
Multiple OS support	Portability to a variety of OSs
Hierarchical user interface	Scalability
Single vendor owns the management software	Optimizes scalability, configuration, and the user interface presentation, and support
Software quality engineering program	Ensures reliability
Intuitive management interface	Easier to visualize components of SAN
Restricted ability to change those items that can negatively affect the network	Protect the network from administration errors
Role-based security	User rights reflect authority and capability
Automatic device discovery	Enables SAN monitoring
Topology mapping	Illustrates all SAN connections and devices
Persistence of topology and properties	Enables software to notify administrator of topology changes
User-specified properties	Enables user to track properties not reported by the devices
User-specified polling interval	Allows optimization based on user needs and network performance
User-specified polling time-out	Allows optimization based on user needs and network performance
Status monitoring—polling	Allows application to identify catastrophic component events
Status monitoring—event listener	Allows application to identify non-catastrophic component events
Event logging	Provides a historical record of events
Event notification (e-mail and pager support)	Eliminates need to constantly monitor the SAN
Mini-map navigational aid	Aids navigation in large topologies
Data-path zoning	Ability to change data-path zones
Discovery and mapping of subnets	Support subnets anywhere in world
Remote monitoring services	Enables remote monitoring of SAN

Thank You For Attending Questions ?

For Additional Information

Please Call Jim Gursha at HPSS

Email: JimGursha@sandisks.com

Telephone: 212-831-0291 (office)

917-359-2087 (cell)