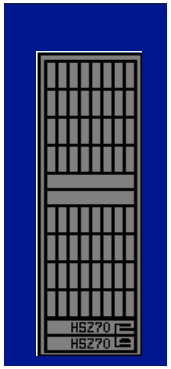

Heterogeneous SAN Configurations

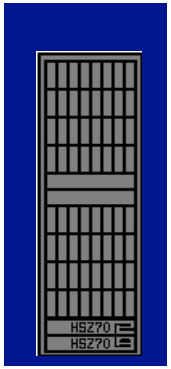

Topics

- Overview of New SAN Configurations
  - Heterogeneous SAN Platform and O/S Support
  - SAN Topologies and Fabric Rules
  - SAN Design Considerations

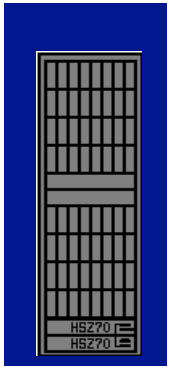

Overview of New SAN Capabilities

- Larger SAN fabrics
  - Defined and Rule based topologies
- Integrated SANs
- Longer Distances
- Additional Platform & O/S Support
- See

  http://h18006.www1.hp.com/products/storageworks/san/documentation.html

  for SAN Design Guide and Supplemental Tables.

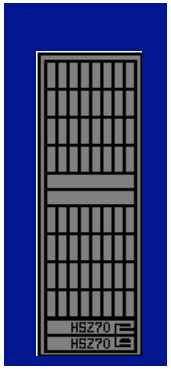

Topology Positioning and Rollout

Enterprise SAN
(~1000s Ports)

Med-Large SAN
(~400 Ports)

Small-Med SAN
(~100 Ports)

Fabric Topology Designs

"Standard Topologies" & "Rules Based Topologies"

~300/~600 Total Ports
200/400 Device Ports
- 20 Switch Backbone SAN Topology
- 20 Switch SAN Topologies

"Defined" Configurations

52/104 Usable Ports
- 4 Switch SANs

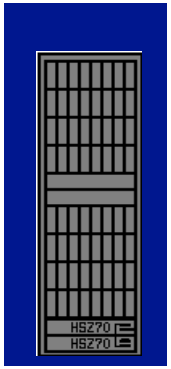

Integrated SANs, Long Distance Support

- Common DRM and Non-DRM SAN
- Common SAN with EBS
- Distances up to 100km
  - DRM, General SAN, HP OpenVMS Host-based Shadowing

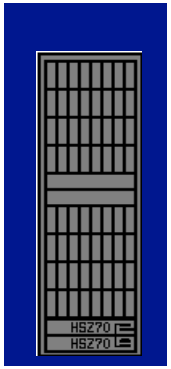

SAN Topologies & Fabric Rules

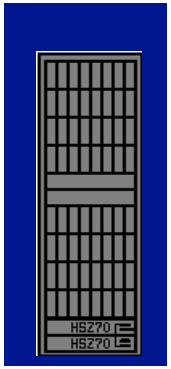

SAN Design Support Levels

- "Standard" Topologies
  - Four SAN topology designs
  - Defined inter-switch connectivity
  - Different topologies optimized for specific data locality need
- "Modified" Topologies
  - Rules allow for deviations within the standard topologies
- "Rules Based" Topologies
  - Other topologies possible based on qualified rules

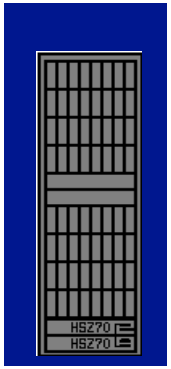

It Depends…

- Topology limits are dependent on which infrastructure switches are present
- Mixing of switch lines is NOT permitted

Interconnect product lines (June 2003)

Infrastructure alternatives that span the enterprise

markets

enterprise -class

Mid-range

entry level


core 2/64


director 2/64


director 2/140


MDS-9509


switch 2/32


edge 2/32


MDS-9216


switch 2/16


switch 2/8 EL


blade 2/8


edge 2/24

B-Series M-Series C-Series

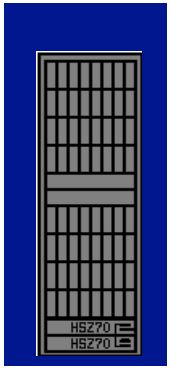

HP Standard Topologies

- SAN Fabric Topologies
  - 2 to 20 Switch "Cascaded" SAN
    - Local Access
  - 2 to 20 Switch "Meshed" SAN
    - Distributed, Local, Centralized Access
  - 2 to 14 Switch "Ring" SAN, or 20 Switch Ring SAN (10 ring switches + 10 cascaded switches)
    - Local, Centralized Access
  - 5 to 20 Switch "Skinny Tree" Backbone SAN
    - Distributed, Centralized, Local Access

Cascaded SAN Fabric

- Accommodates diverse geographic conditions
- Scales easily for additional connectivity
- Shared backup is supported
- Shared management is supported
- Optimal local access is inherent in the design

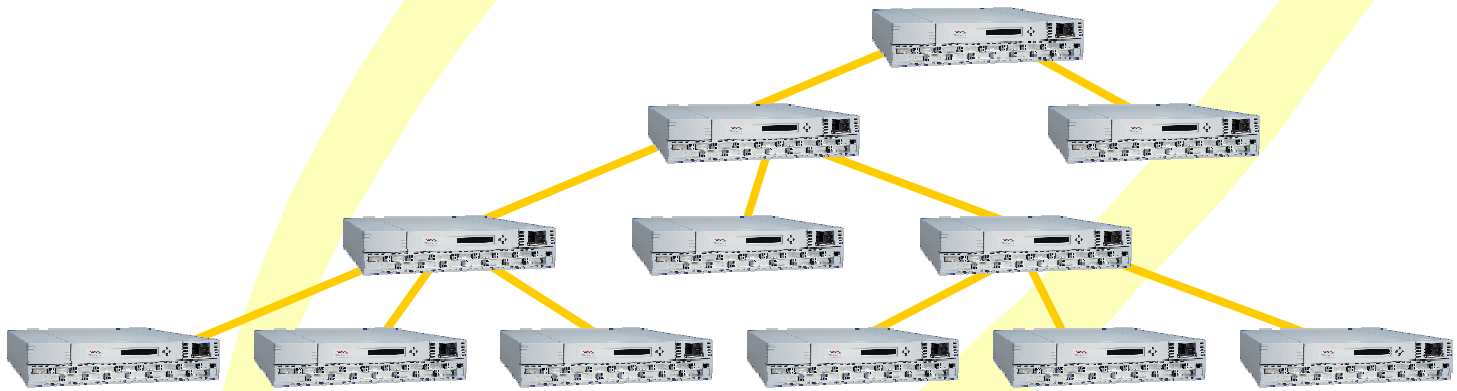

12 Switch Cascaded Fabric Example

Meshed SAN Fabric

- Can be configured for many to many or local access, or a mix
- Provides protection against link and switch port failures
- Scales easily for additional connectivity
- Shared backup is supported
- Centralized management is supported
- Optimal distributed access is inherent in the design

12 Switch Meshed Fabric Example

Ring SAN Fabric

- Easy to build
- Modular design
- Scaling is simple and non-disruptive
- Shared backup is supported
- Centralized management is supported
- Optimal local access is inherent in the design

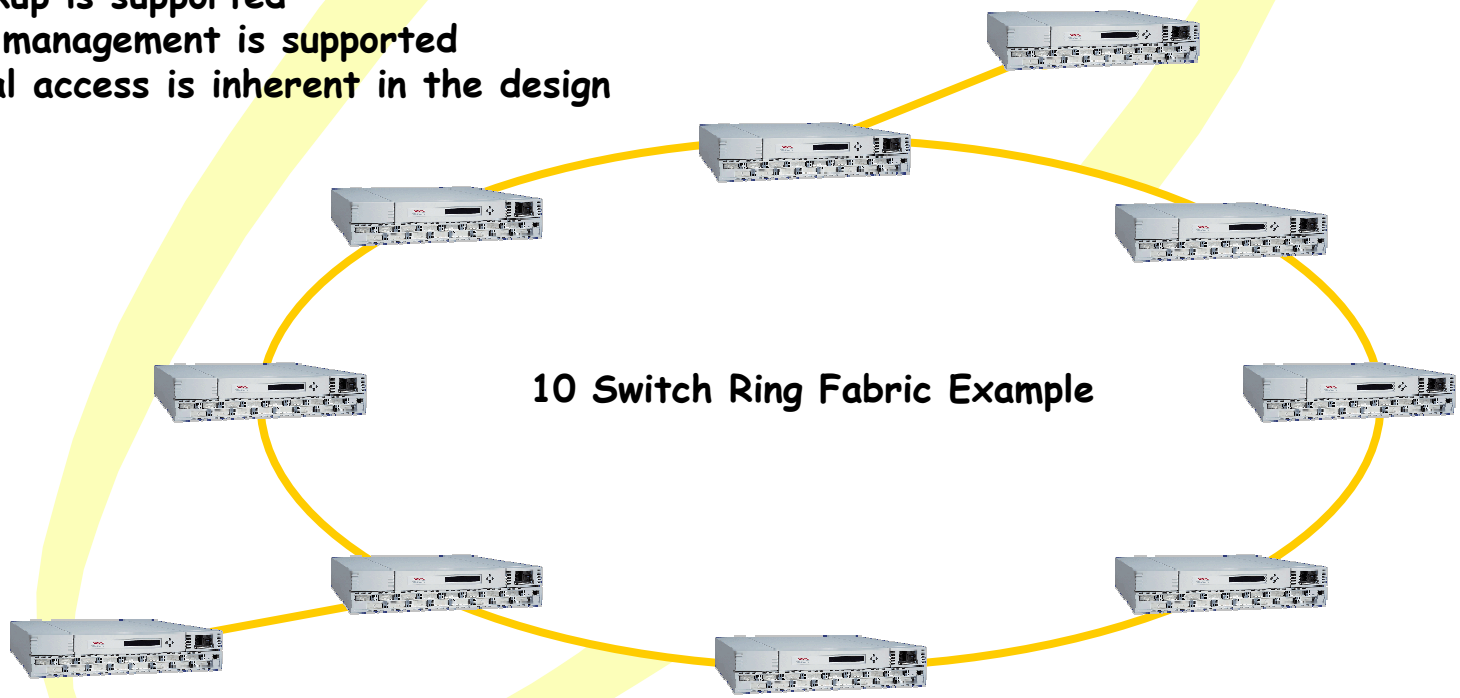

10 Switch Ring Fabric Example

Skinny Tree Backbone SAN Fabric

- Efficient port expansion
- All edge switches are only two hops apart
- Shared backup is supported
- Centralized management is supported
- With two backbone switches, provides switch redundancy
- Maximum flexibility for mixed data access types
- Best-suited to take full advantage of expected future technological developments such as storage virtualization

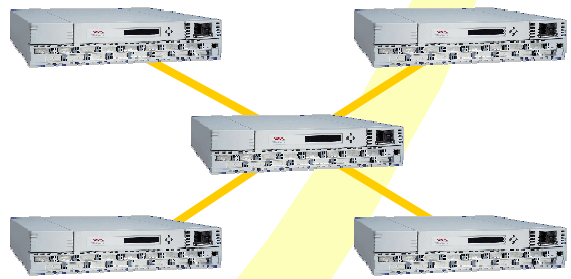

5 Switch Tree Backbone Fabric Example

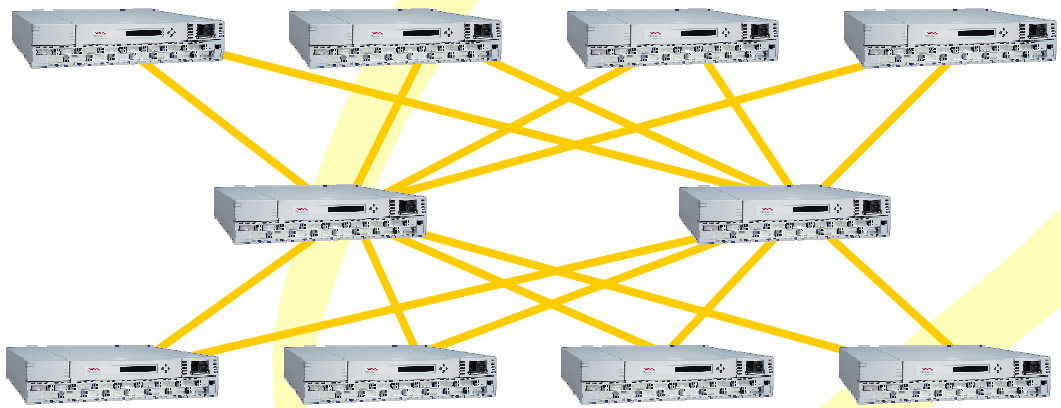

10 Switch Tree Backbone Fabric Example

Skinny Tree Backbone SAN Fabric

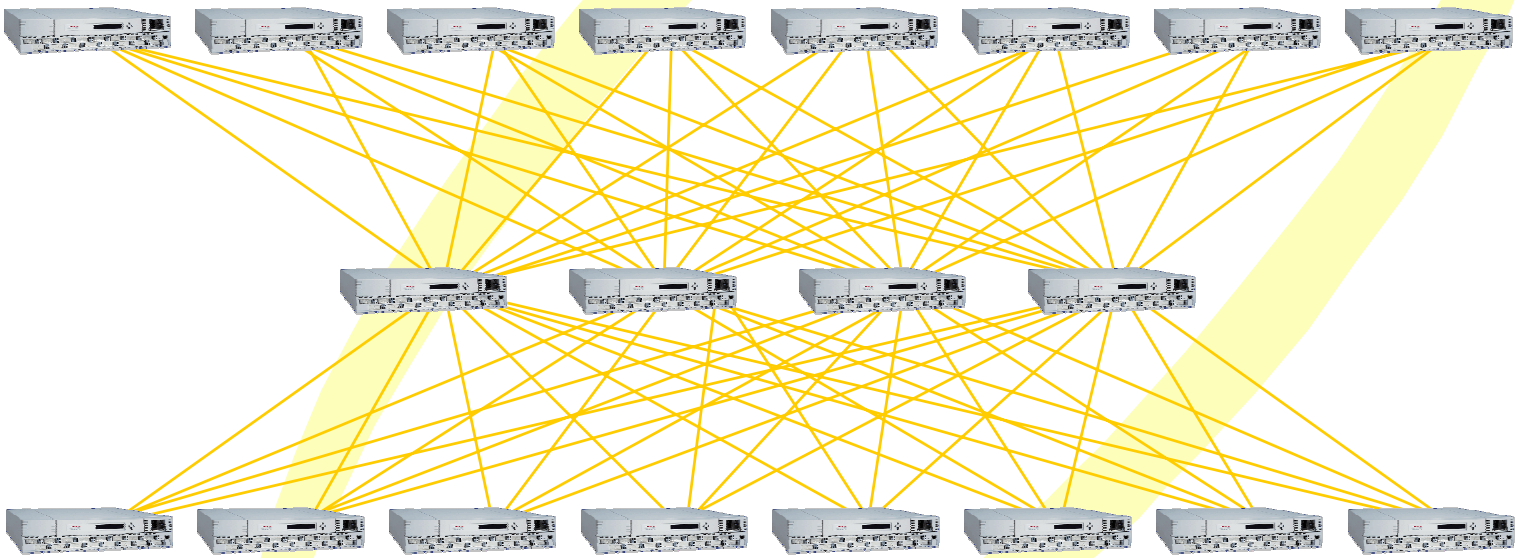

20 Switch Tree Backbone Fabric Example

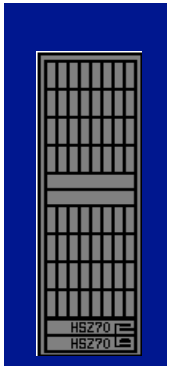

Topology Port Maximums

SAN Topology Min/Max Number of Switches	Total Number Of Ports[1]		Maximum Number Of Device Ports[2]	
	Single Fabric	Two Fabrics	Single Fabric	Two Fabrics
Cascaded, Meshed, Ring[3] Single Fabric: 2 to 20 Dual Fabric: 4 to 40	**320**	**640**	**~200**	**~400**
Tree Backbone Single Fabric: 5 to 20 Dual Fabric: 10 to 40			**192**	**384**

[1] Assumes 16 port switches.
[2] Assumes 16 port switches. Indicates the number of ports available for server and storage connectivity. For Cascaded, Meshed, and Ring fabrics the number of device ports is approximate since this is dependent on the specific switch arrangement and the number of ISLs utilized.
[3] The maximum number of switches within the ring of a Ring fabric is 14, other switches can be added to the fabric if the ring contains less than 14 switches provided there are no more then 7 hops between any two devices in the SAN.

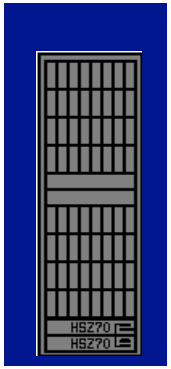

SAN Design Considerations

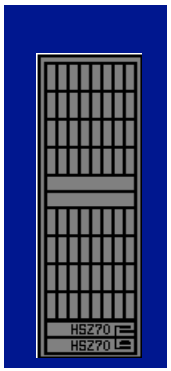

Design Requirements for a SAN

Backup
Restore

Distance

Connectivity
& Capacity

Data Locality

Scalability

Heterogeneous
Platforms

Availability

Performance

Disaster
Tolerance

Management
& Security

Cost

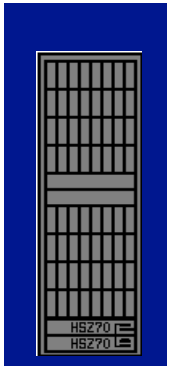

SAN Topology Metrics

- Data Locality or "Locality of Allocation"
  - The location of storage LUNs relative to servers
- Connectivity & Capacity Efficiency
  - The percentage of ports available for devices
  - The amount of storage available
- Fabric Performance
  - The cross-sectional bandwidth
- Growth Efficiency
  - Scalability
- Availability Levels
  - Fabric and server/storage availability

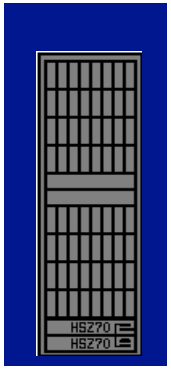

Locality of Allocation - 2 Factors

- Specific application-driven access requirements
- Where the storage is deployed or located relative to the servers that require data access
  - Local or "One to One"
    - Servers and storage typically on the same switch
  - Centralized or "Many to One"
    - All servers distributed across multiple switches throughout the fabric accessing a centralized storage pool
  - Distributed or "Many to Many"
    - Servers and storage pools are distributed across multiple switches throughout the fabric

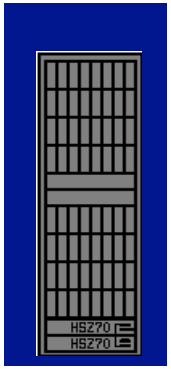

Topology Usage Ratings

SAN Topology	Data Locality		
	Local "One to One"	**Centralized "Many to One"**	**Distributed "Many to Many"**
Cascaded	Highest	Not Recommended	Not Recommended
Meshed	High	Medium	High
Ring	Highest	Medium	Not Recommended
Tree Backbone	Medium	High	Highest

- Other Topologies Possible
  - HP "Modified" Topologies
  - Custom "Rules Based" Topologies

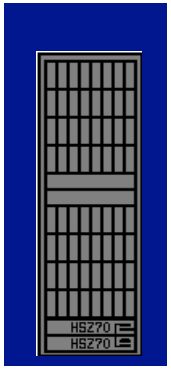

Availability Levels

- Level 1: Maximum Connectivity
  - A single fabric with single server and storage paths (equivalent to single path parallel SCSI)
- Level 2: Fabric Resiliency
  - A single fabric with multiple fabric paths and single server and storage paths
- Level 3: High Availability/Multi-Path
  - A single fabric with multiple fabric paths, and multiple server and multiple storage paths
- Level 4: High Availability/Multi-Path/Fault Tolerant/No Single Point of Failure
  - Two (or more) redundant fabrics

Availability Level Examples

- Level 1: Maximum Connectivity

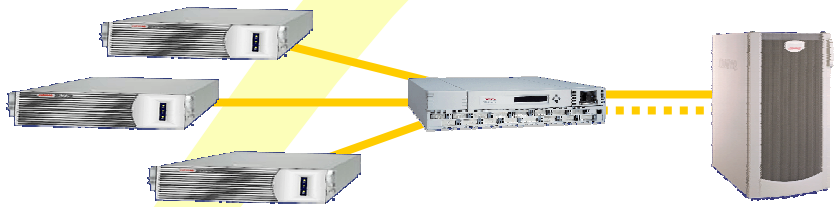

- Level 2: Fabric Resiliency

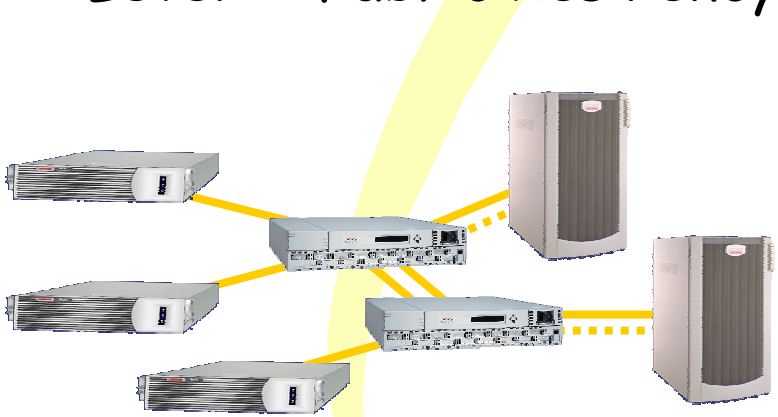

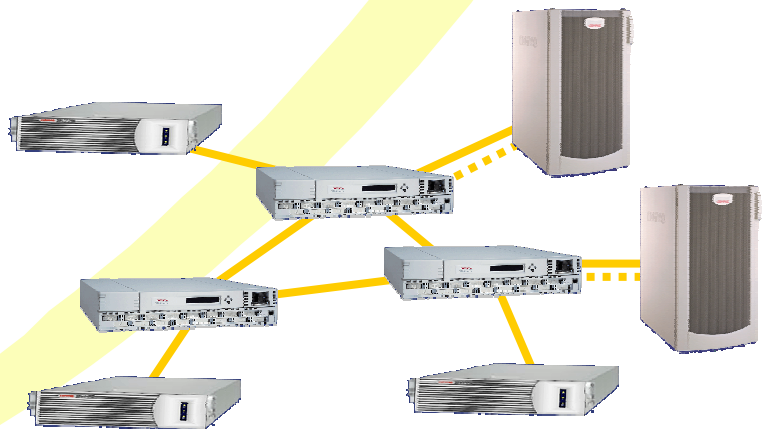

Availability Level Examples

- Level 2+: Better Fabric Resiliency - Backbone Switch Redundancy

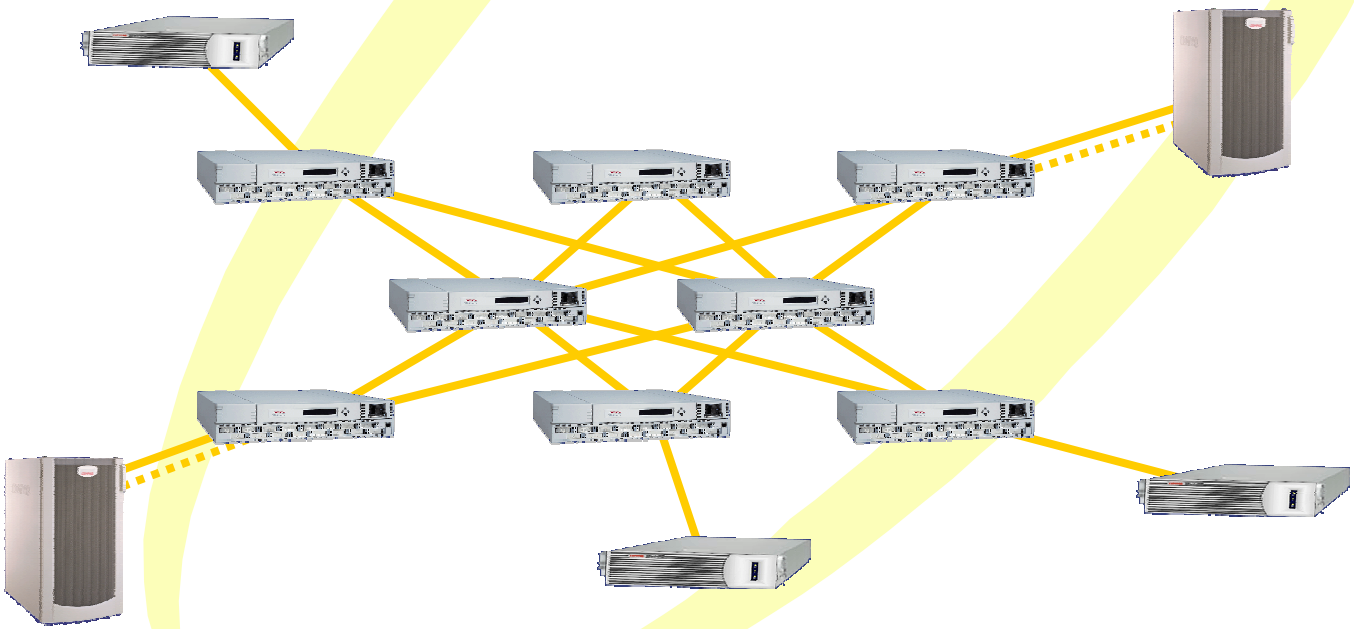

Availability Level Examples

- Level 3: Single Fabric, High Availability/Multi-Path

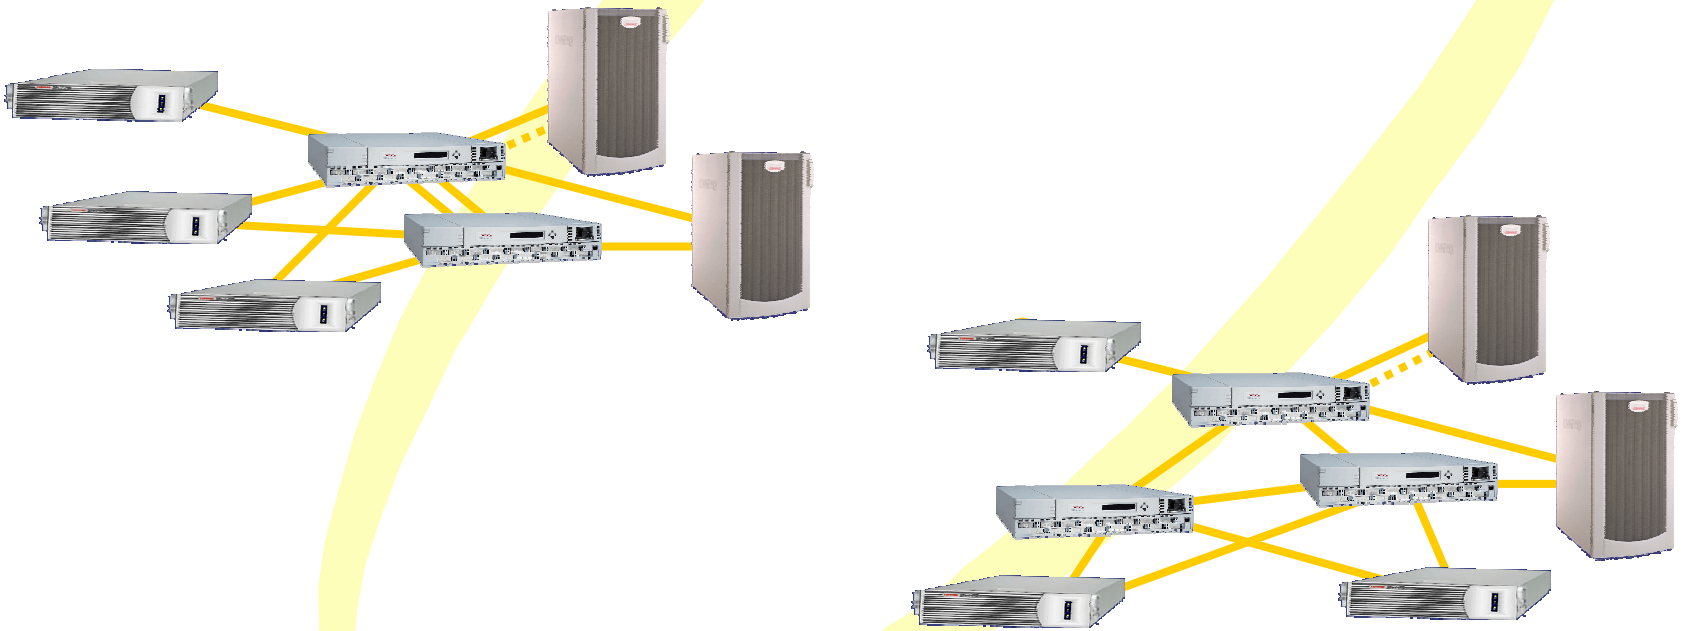

Availability Level Examples

- Level 4: Dual Fabric High Availability/Multi-Path Fault Tolerant - "No Single Point of Failure"

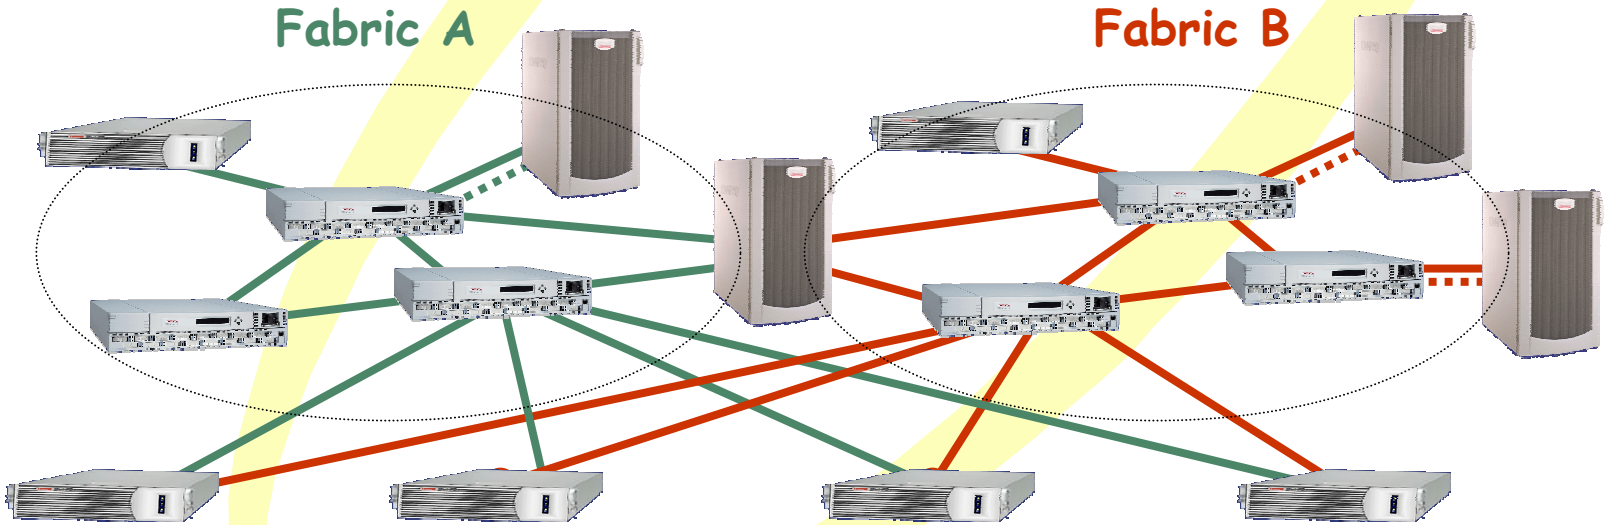

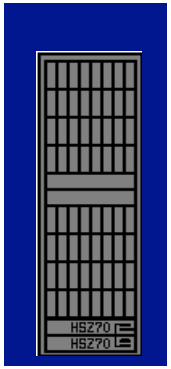

Design Rules for HP Standard and Rules-Based SAN Topologies

- Must adhere to all "Fabric Rules"
  - Maximum number of switches-hops-ISL's (20-7-16/8)
  - Fabric topology-specific rules (14 switch ring)
  - Switch model usage rules
    - Original Fibre Channel Switches: 4 switches maximum in a fabric, compatibility mode when mixing
    - #of E-Ports: SAN Switch 8-EL,1 Std, additional optional
    - QuickLoop: SAN Switch 8/16 and SAN Switch 16-EL only
  - SAN appliance usage rules/requirements
    - 1 Appliance required per fabric
- Performance Recommendations - #ISL's
  - Minimize "Oversubscription"

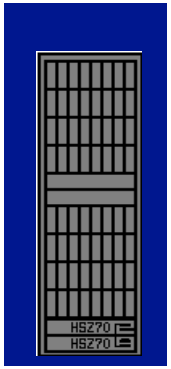

Design Rules for HP Standard and Rules-Based SAN Topologies

- Must adhere to all "Interconnect Rules"
  - F-Port, FL-Port (public & private), E-Port usage
  - Media/Interconnects/Interface Components
    - Fiber optic cable types: 50, 62.5, & 9 micron
    - Short-wave GBIC & Long-wave GBIC (switch to switch)
    - Very long distance GBIC (switch to switch)
  - Distances/GBIC/Loss Budgets
    - 50u/500 meters/SW: 4db
    - 62.5u/200 meters/SW: 3 db
    - 9u/10km/LW: 7.8 db
    - 9u/100km/VLD: 23 db

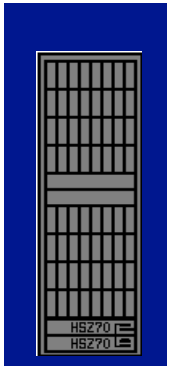

SAN Scaling and Migration

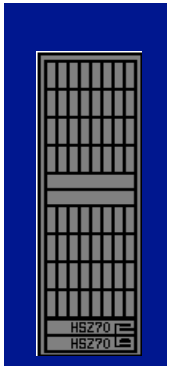

SAN Scaling

SAN Topology	Scalability (For All Topologies)	Migration
Cascaded	• Increase the number of switches • Use higher port count switches • Deploy multiple fabrics • Transition to a different topology	Convert to Meshed, Ring or Tree
Meshed		Convert to Ring or Tree
Ring		Convert to Meshed or Tree
Tree		Add additional backbone switches

- Scaling
  - Refer to OpenSAN Design Guide, Best Practices
  - Provides an overview of the expansion process for each topology

SAN Scaling

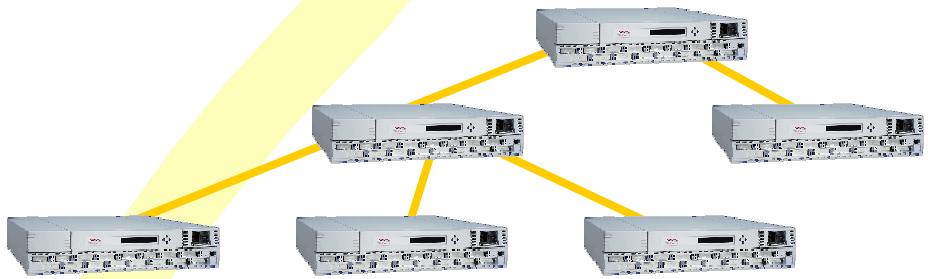

Cascaded Fabric Example

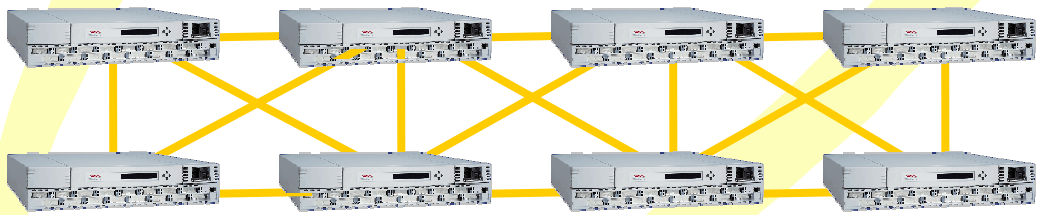

Meshed Fabric Example

SAN Scaling

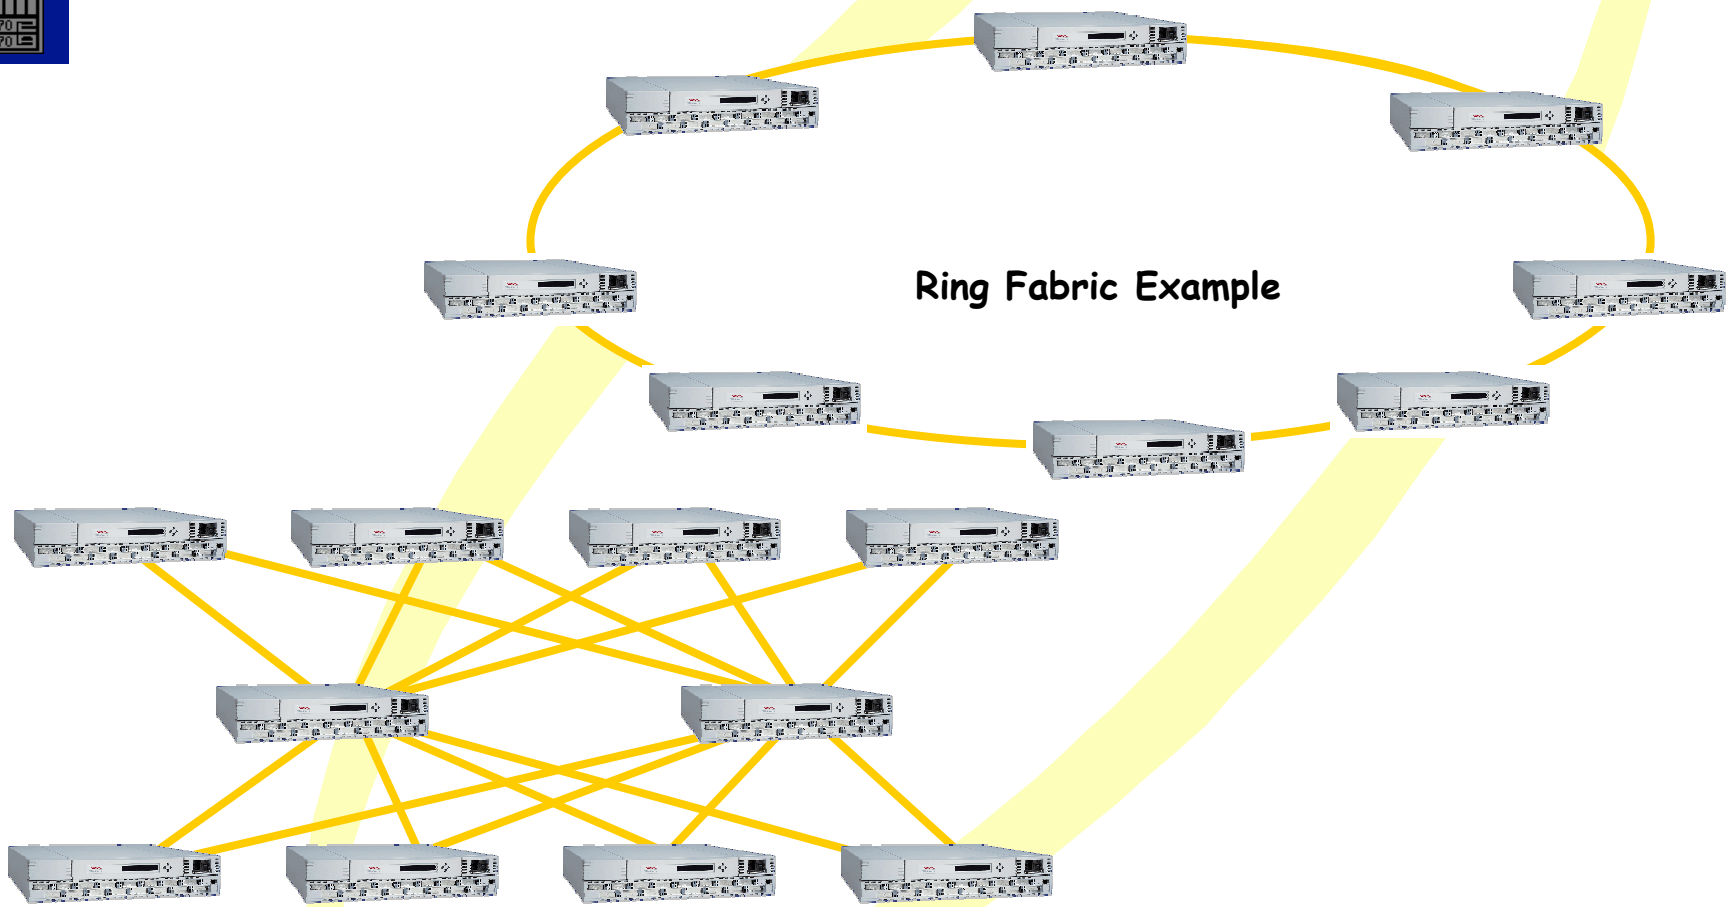

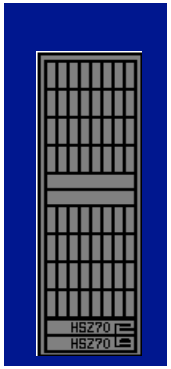

SAN Migration

SAN Topology	Scalability (For All Topologies)	Migration
Cascaded	• Increase the number of switches • Use higher port count switches • Deploy multiple fabrics • Transition to a different topology	Convert to Meshed, Ring or Tree
Meshed		Convert to Ring or Tree
Ring[1]		Convert to Meshed or Tree
Tree		Add additional backbone switches

1 The Ring fabric can be scaled by adding switches to the ring (up to 14) or by adding switches outside the ring (maximum of 20 in the fabric).

- Migration
  - Refer to OpenSAN Design Guide, Best Practices
  - Provides an overview of the migration process available for each topology
  - Some migrations are easy…some are very difficult

SAN Migration

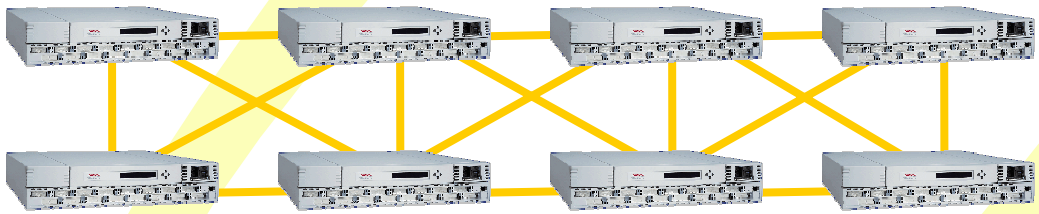

Meshed Fabric Example

SAN Migration

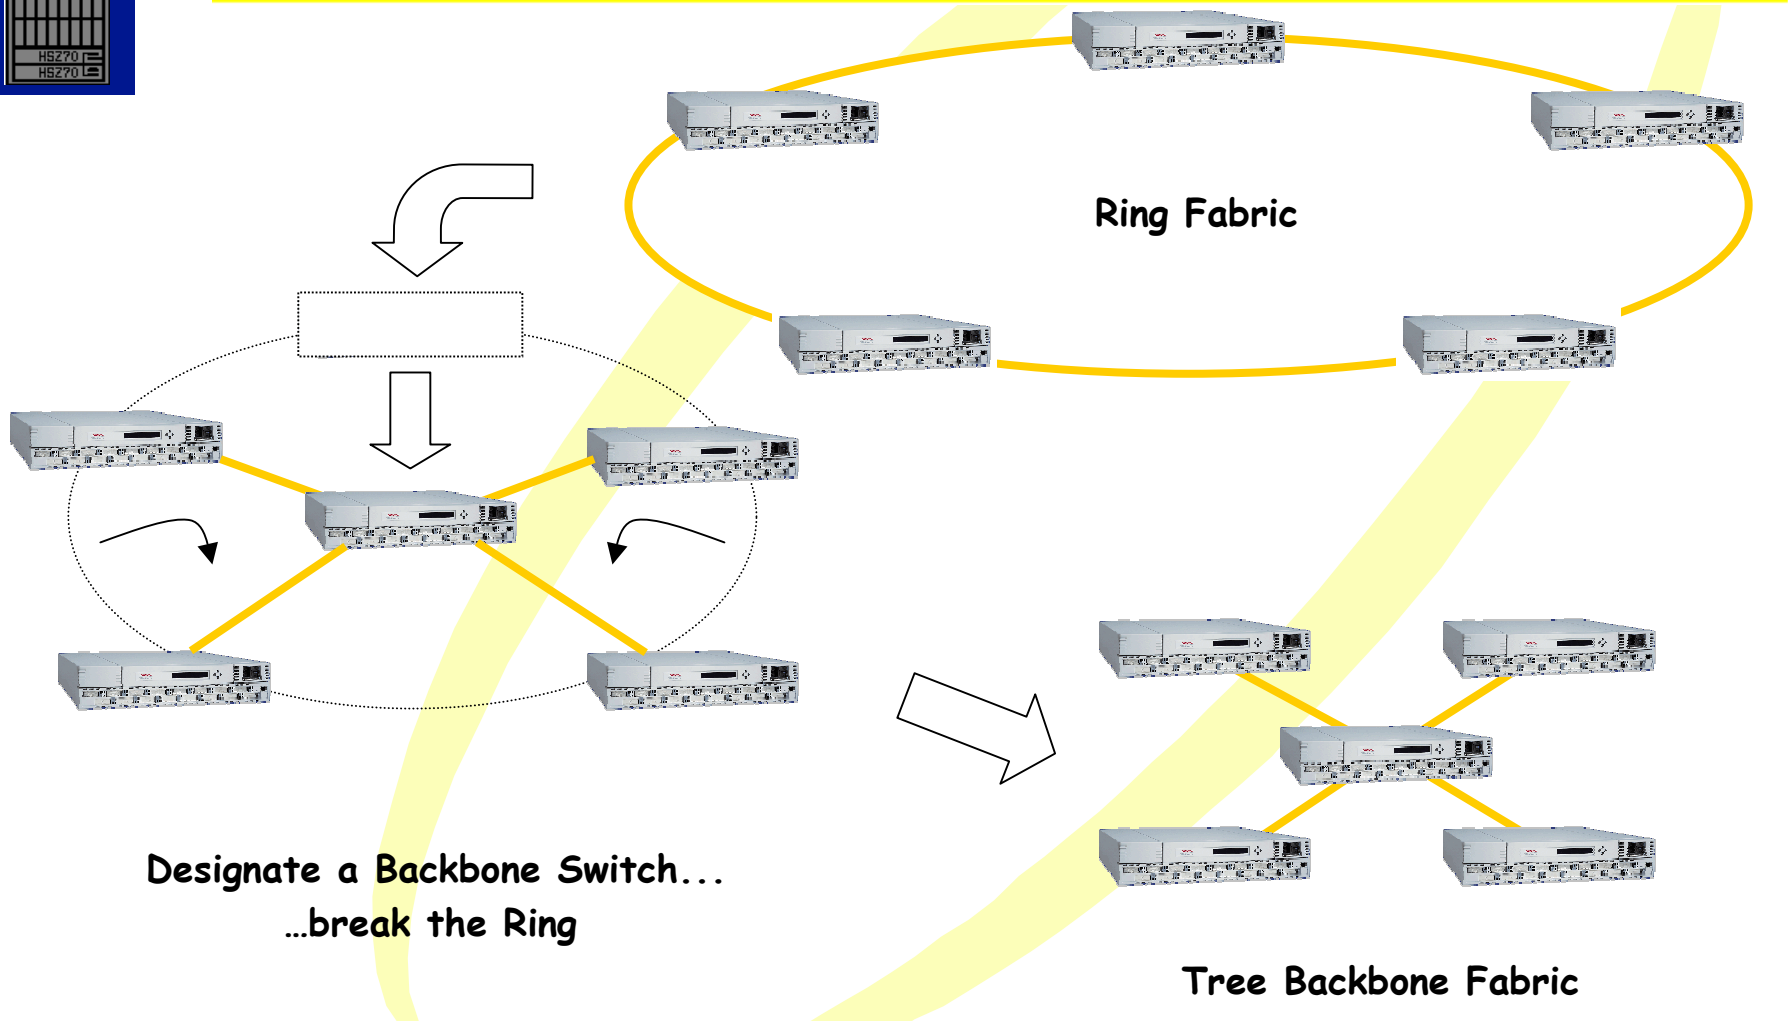

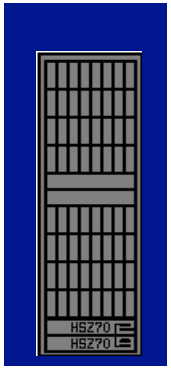

SAN Integration - DRM/EBS

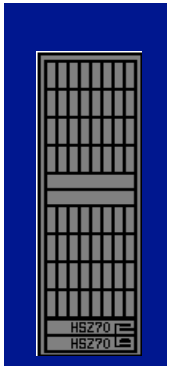

SAN Fabric Integration

- Integrated SAN/DRM/CA/HBS
  - Usage: Common SAN, Disaster Tolerance
    - Mixed DRM controllers and non-DRM controllers supported in the same physical SAN
      - Each DRM controller pair (local and remote) and the servers accessing them must be configured in two separate DRM zones (one for each path) from the rest of the SAN
      - DRM controllers are supported for homogeneous server access from servers configured for DRM only
      - DRM on NT/Windows 2000 supported with 4.41a7 driver only
    - Mixed DRM/SAN/HBS traffic supported on common 100Km fibre channel link provided SAN and HBS Luns are accessing only non-DRM controllers (v8.5F, v8.5S). Note: Must monitor link utilization to avoid adding too much traffic such that the long distance link becomes the bottleneck. Poor performance and application timeouts may be encountered if the link is overloaded.

Integrated DRM/SAN Over Long Distance

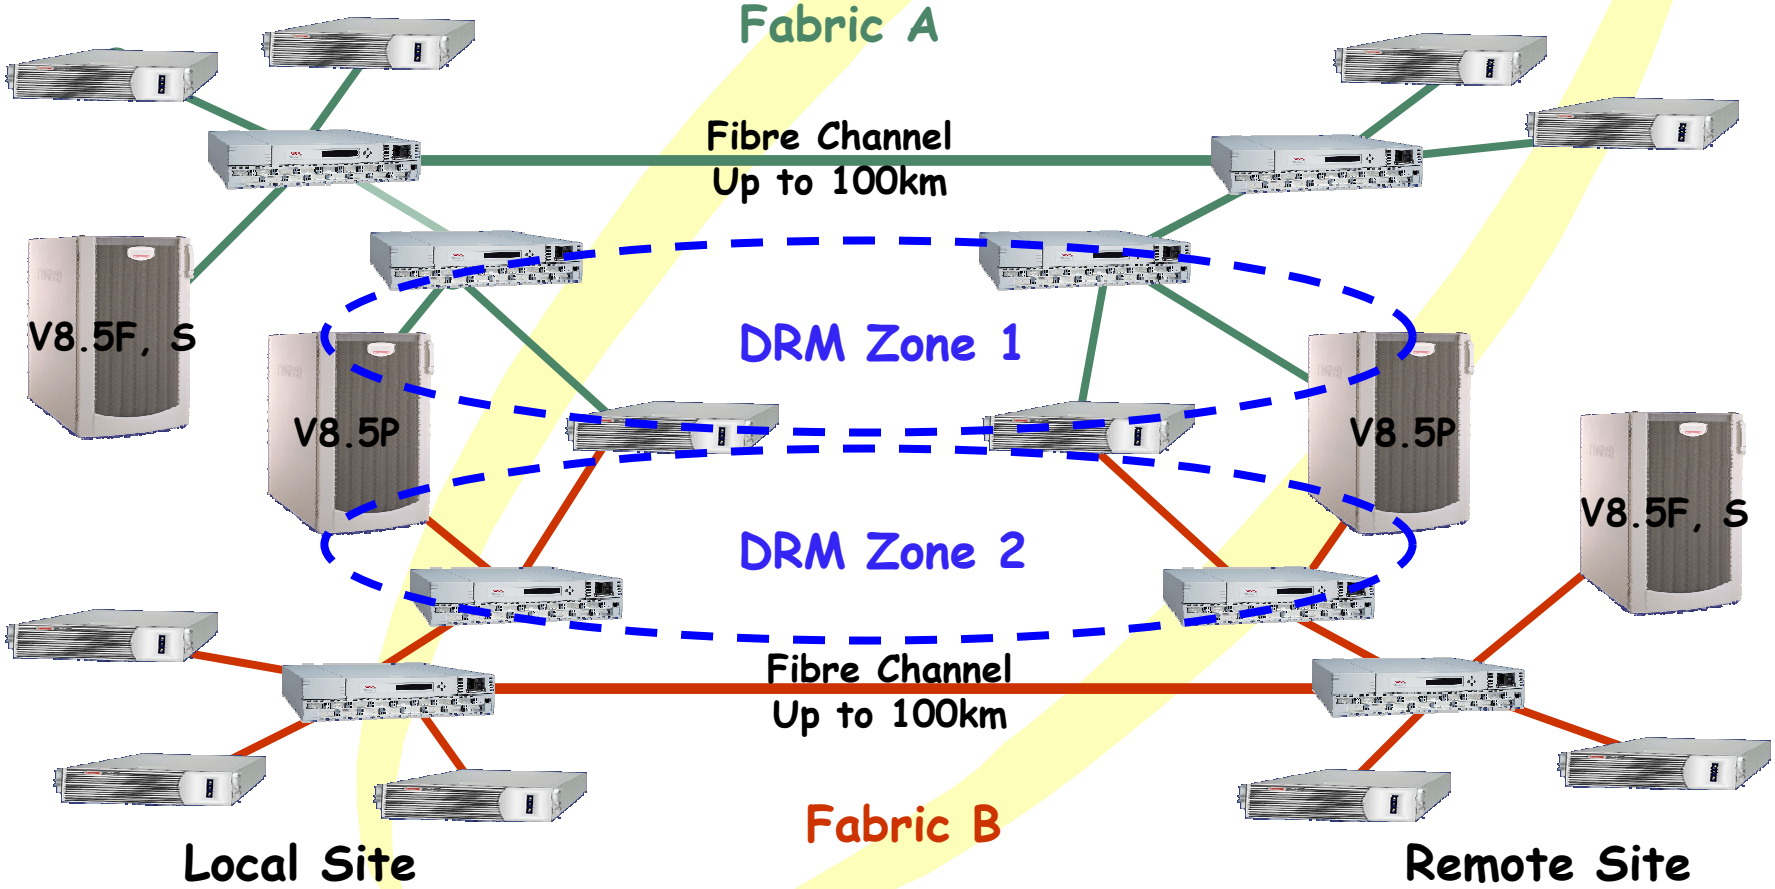

Integrated EBS/Tree Backbone SAN

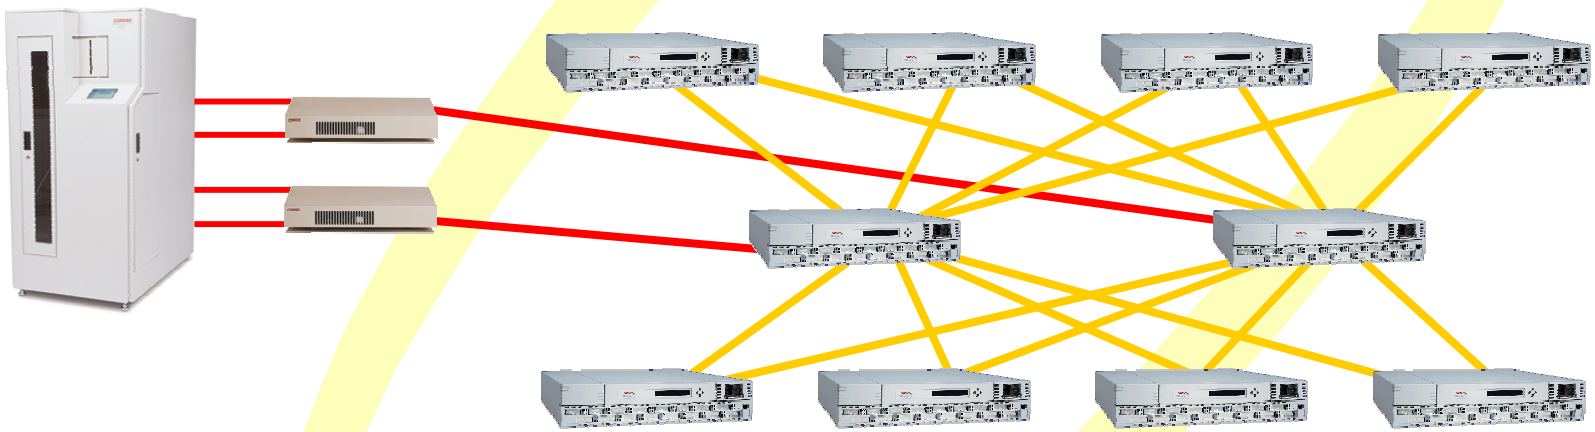

Tree Backbone Fabric/EBS Example

- Integrated EBS
  - Usage: Common SAN, Centralized/Consolidated Backup

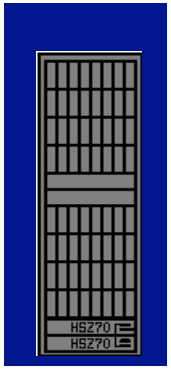

Heterogeneous SAN Platform and O/S Integration

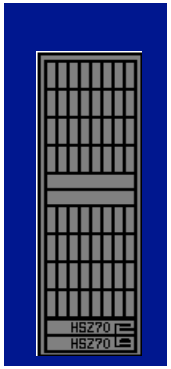

Heterogeneous SAN Platform and O/S Support

- Supported Platform and O/S Versions
  - hp OpenVMS/Clusters
    - 7.2-1H1
    - 7.2-2
    - 7.3 and higher
  - hpTru64 UNIX/TruCluster
    - 4.0F/G (ACS only),
    - 5.1
  - HP-UX
    - MC/ServiceGuard
    - 11.0, 11i
    - 10.20 (FC-AL, 8.6 only)
  - IBM AIX 4.3.3, 5.1/HACMP 4.4.1

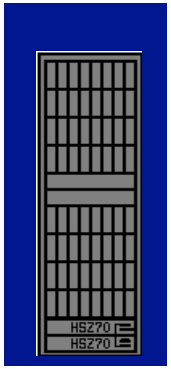

Heterogeneous SAN Platform and O/S Support

- Supported Platform and O/S Versions
  - Caldera Linux 3.1 (Intel, ACS only, single path)
  - Redhat Linux
    - 7.2, Advanced Server
  - SUSE Linux
    - 7.2, SLES 7
  - Novell Netware
    - 4.2, 5.1, 6.0
    - 5.1 Clusters v1.01/6.0 Clusters 1.06
  - Microsoft Windows NT 4.0 SP6a/MSCS
  - Microsoft Windows 2000 SP2, SP3/MSCS

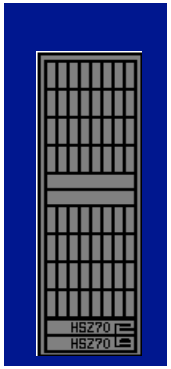

Heterogeneous SAN Platform and O/S Support

- Supported Platform and O/S Versions
  - SGI IRIX (ACS 8.6 only, single path)
    - 6.5.11, 6.5.12
  - Sun 2.6, 7, 8 /Sun Clusters 2.6, 7, 8/Veritas Clusters 1.3

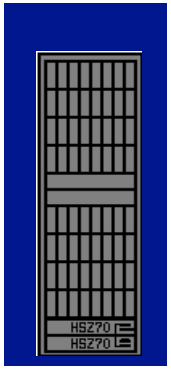

Design Rules for HP Standard and Rules-Based SAN Topologies

- Must adhere to all "Platform & O/S Specific Rules"
- Must adhere to "Heterogeneous Rules"
  - Reside in the same physical SAN?
    - Zoning may be required: HBA interactions, connection table limits, or multi-pathing in a single fabric
  - Share the same physical storage system?
    - Shared storage system access allowed if a match of failover mode & SCSI-mode
- Complete SAN Product Support List
  - Refer to SAN Design Reference Guide and the SAN support WEB page

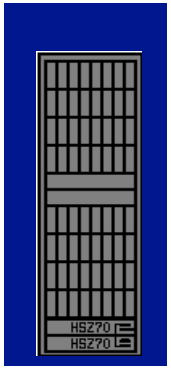

Heterogeneous SAN Platform and O/S Support

- Zoning Is Required When:
  - Known O/S interaction problems
    - Multiple Tru64 TruClusters, Tru64 with Windows NT or Windows 2000
    - HP-UX or Linux with all other O/S's
  - Unknown whether or not there are interaction problems
    - When a platform or O/S is supported in a homogeneous SAN, but has not yet been tested in a heterogeneous SAN (IBM AIX)
  - HSG60/80 storage controller connection table limits
    - ACS 8.5 maximum of 64 connection table entries
    - ACS 8.6 and 8.7 maximum of 96 connection table entries
  - Multi-Pathing in a single meshed fabric
    - Two paths must be in separate fabric zones

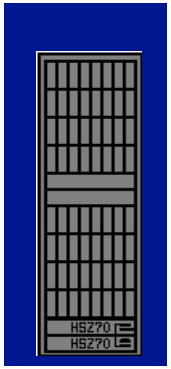

Heterogeneous SAN Platform and O/S Support - Maximums Table

Check latest version of SAN Design Guide for changes!!!

Platform or Operating System	Host Bus Adapters per Server[1]	Active Controller Ports per HBA	LUNs per HBA Target[2]	Servers per Active Controller Port[3]	Servers per Storage System[4]
Compaq OpenVMS 7.2-1, 7.2-1H1	4	16	128	16	32
Compaq Tru64 UNIX 4.0F	32	4	8	4	8
Compaq Tru64 UNIX 5.0A, 5.1	64	128	128	32[5]	64[5]
HP-UX 10.20, 11.0	16	4	8	8	16
IBM AIX 4.2.1, 4.3.0, 4.3.1, 4.3.2, 4.3.3	4	4	16	4	8
Microsoft Windows NT 4.0 SP5, SP6a Windows 2000 Advanced Server, SP1	4	4	8	8	16
Novell Netware 4.2, 5.x,	4	4	32	8	16
Redhat Linux (Alpha/Intel) 6.1 & 6.2	2	4	64	2	4
SUSE Linux (Alpha) 6.3, (Intel) 6.3	2	4	64	2	4
SGI IRIX 6.5.7, 6.5.8	4	4	64	2	4
SUN Solaris 2.6, 7 & 8 (32/64 bit)	16	4	64	8	16
Heterogeneous SAN				See Note[6]	

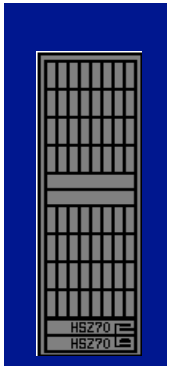

Heterogeneous SAN Platform and O/S Support - Maximums Table Notes

1. The maximum number of host bus adapters supported per server is dependent on the specific server model.
2. The number shown in this column is reduced by 1 if the command console lun is enabled.
3. The maximum number of host bus adapters that can be configured for access to an active controller port. Assumes 1 host bus adapter per server for single path using controller transparent failover or 2 host bus adapters per server for multi-path using controller multiple-bus failover. For transparent failover, the limit is specified by controller port pair – one active and one standby controller port. For multiple-bus failover, the limit is specified per single active port.
4. For example, under Windows NT up to 16 servers can be configured for access to a controller pair whether configured for transparent or multiple-bus failover mode. In transparent failover mode up to 8 servers with one HBA in each can be configured for access to the port 1 active/standby port pair, and up to another 8 servers with one HBA in each can be configured for access to the port 2 active/standby port pair.
5. In multiple-bus failover mode up to 16 servers with two HBAs in each can be configured for access. Up to 8 servers are configured for access with 8 HBAs on port 1 of the first controller and 8 HBAs on port 1 of the second controller. Up to another 8 servers are configured with 8 HBAs on port 2 of the first controller and 8 HBAs on port 2 of the second controller. In this example, for both failover modes the limit of 8 HBAs/active controller port is never exceeded, and both support up to 16 servers per controller pair.
6. Assumes 1 host bus adapter per server for single path using controller transparent failover or 2 host bus adapters per server for multi-path using controller multiple-bus failover.
7. Requires the use of zoning to limit the number of connections to 32 per port pair.
8. In a heterogeneous SAN, the maximum number of servers per controller pair, per controller port-pair (transparent failover) or per controller port (multiple-bus failover) is equal to the lowest maximum listed in these columns for the operating systems that are sharing the storage system.

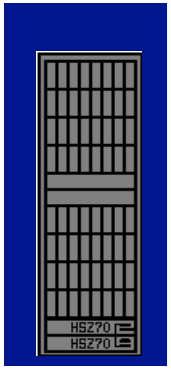

Best Practices and Troubleshooting

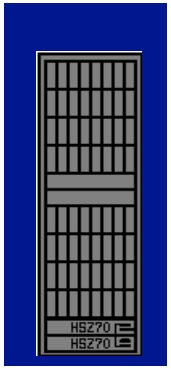

Best Practices

- Planning a SAN
  - Topology Map, Configuration Layout, Zoning Map, Storage Map
- Configuring a SAN
  - Recording, Cabling, Labeling, Naming
- Upgrading a SAN
  - Adding, replacing a switch
  - Updating switch firmware
  - Merging Fabrics

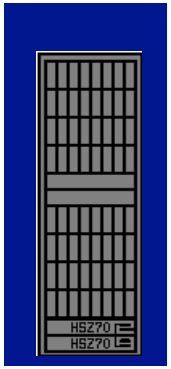

Troubleshooting

- Isolating Problems
  - Servers
  - Switches
  - Storage
- Fabric Segmentation
  - Mismatch of Zoning information
    - Configuration mismatch
    - Type mismatch
    - Content mismatch
  - Mismatch of certain switch configuration settings
    - bb_credit, r_a_tov, e_d_tov, data field size, device probing, vc encoded address mode, translative mode, per-frame route priority