# Technical Overview of the Enterprise Virtual Array

# Objectives

- Discuss storage virtualization
- Describe the hardware features and functions of the Enterprise Virtual Array
- Describe the Software features of the Virtual Controller Software (VCS).
- Identify examples of SAN-based solutions that incorporate the Enterprise Virtual Array

# HP StorageWorks
# The industry's most complete array portfolio

MSA1500

**Low Cost SAN**
**24 T Bytes**

- WEB, Exchange, SQL
- Simple DAS-to-SAN (ProLiant)
- Tiered Storage
- Windows, Linux, Netware + more

eva3000

**Midrange Modular**
**≤14 T Bytes**

- Mid—sized DB/Remote office applications
- Simplification through virtualization
- Windows, HP-UX, Linux, + more

eva5000

**Enterprise Modular**
**≤35 T Bytes**

- Storage consolidation + disaster recovery
- Simplification through virtualization
- Windows, HP-UX, Linux, + more

XP128 XP1024

**Enterprise Frame**
**≤149 T Bytes**

- Data center consolidation + disaster recovery
- Large scale Oracle/SAP applications
- HP-UX, Windows, + 20 more including mainframe

# Positioning HP StorageWorks

NT, W2K, NetWare, Linux

Broad Multi-vendor OS

| Direct Attached | Entry Level Small Clusters | Mid-Range Homogenous SAN | Enterprise Heterogeneous SAN |
|---|---|---|---|

**Data Center**
- Multi platform & OS support (broad base)
- Zero downtime
- High capacity
- High performance

**Dept/Branch Office/Medium Business**
- Multi platform & OS support (core base)
- Cluster solutions
- Limited technical resources on-site
- Moderate price sensitivity

**Workgroups/Small Business**
- Lower entry price point
- Simple deployment
- Partners are technical resources

XP128/XP1024 (149 TB max)

EVA5000 (35 TB max)

EVA3000 (8 TB max)

MSA1000/1500 (24 TB max)

MSA500 G2

MSA30

Connectivity

Scalability

# The Value of Virtualization

hp StorageWorks

eva3000

eva5000

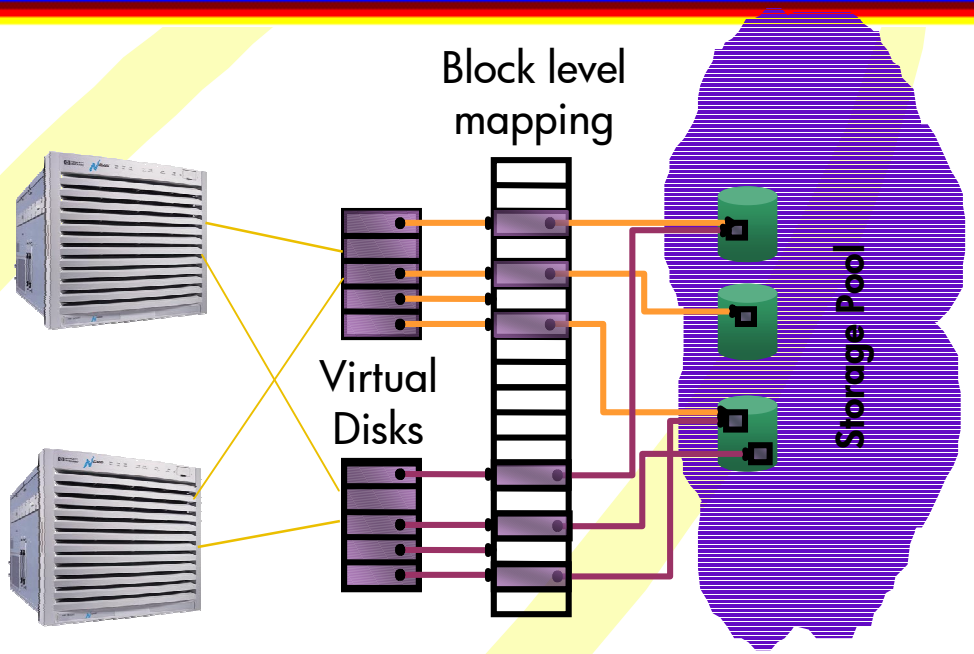


**the technology . . .**

- **capacity is treated as a pool of storage blocks, not discrete disks**
- **powerful mapping techniques present a logical view of storage to host servers**

EVA family values

# What is Storage Virtualization?

- Combination of hardware, software, and networking gear that turns a SAN into a "coherent, intelligent storage subsystem capable of automating information management processes." (Illuminata, April 2001)
- Allows heterogeneous storage devices to be combined into a SAN-attached storage pool that is managed as a single storage resource. (IDC, November 2000)
- The act of integrating one or more (back end) services or functions with additional (front end) functionality for the purpose of providing useful abstractions. (SNIA Dictionary of storage terminology, 2004)

# Virtualization on EVA

- All raw disk storage is pooled
- Virtual Disks (VDs) (or LUNs) are drawn from a pool
- Less overhead for:
  - Controller state
  - Mapping metadata
  - Audit / error logs
- Virtual Disks managed by customer to these constraints:
  - Size range 1GB – 2TB, in 1GB increments
  - Up to 512 Virtually Disks — selectively presented to hosts
  - Each Virtual Disk can be zero (none), medium or high redundancy
    - Vraid0, Vraid5, or Vraid1
  - Virtual Disk size can be dynamically expanded, but not shrunk

# Virtualization

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

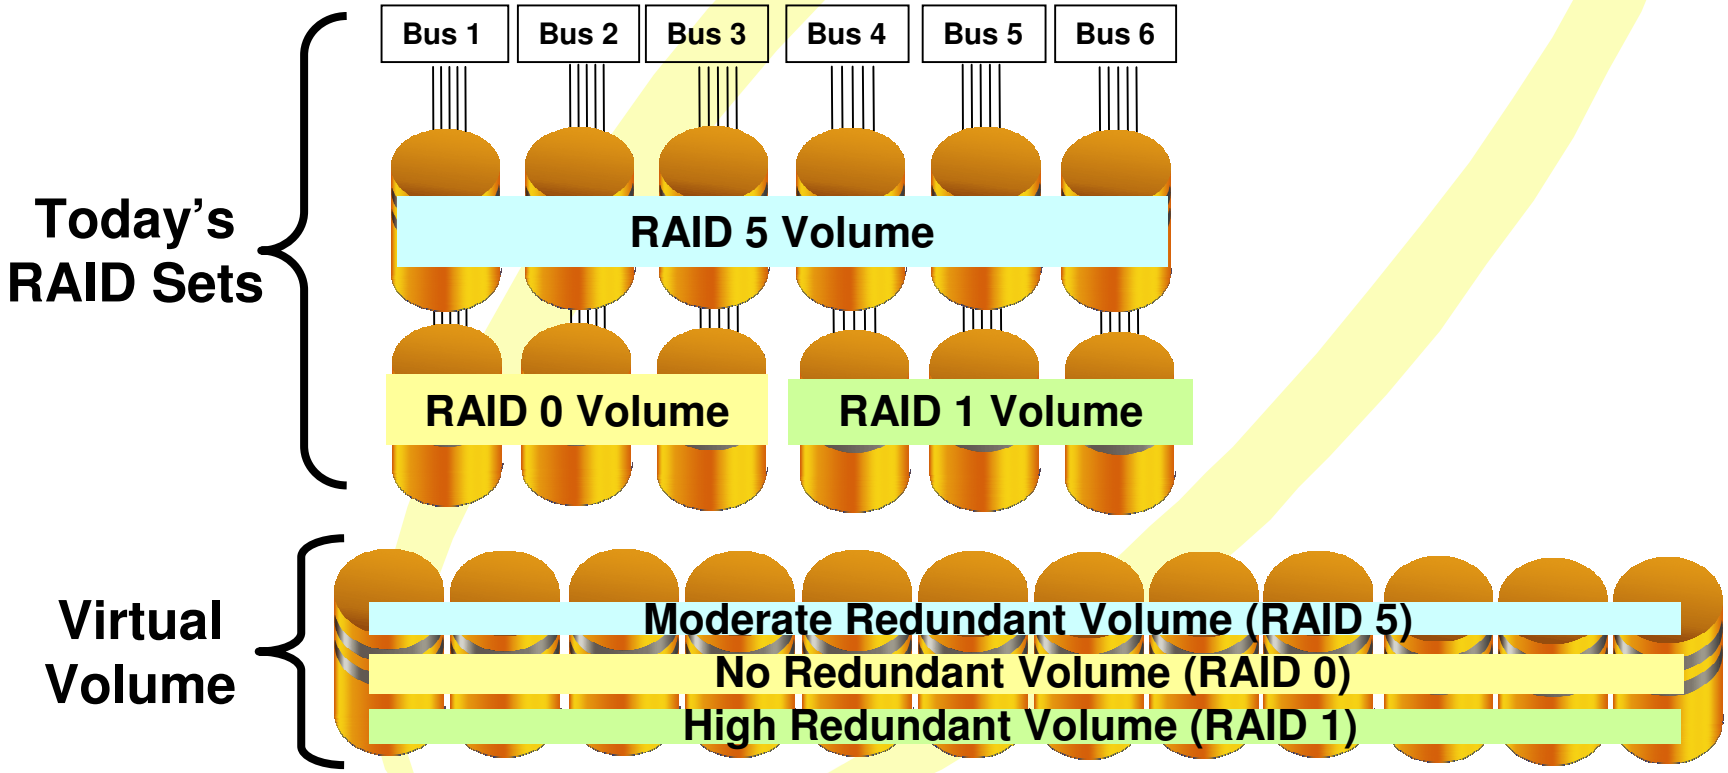


*Workload is evenly distributed across all spindles in group*

# 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**

**Add More Disks**

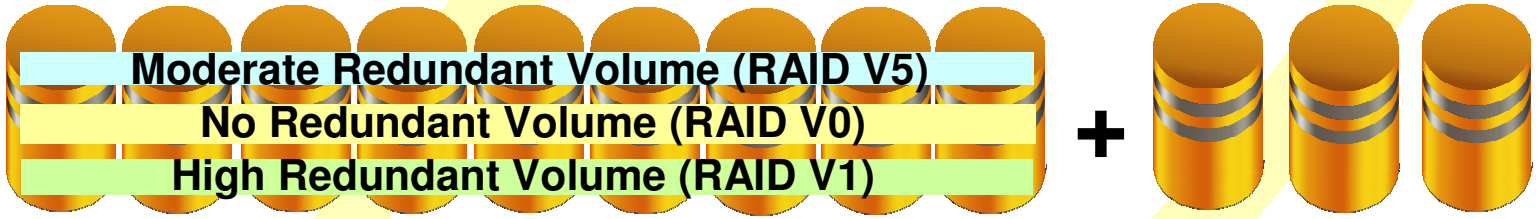


• **Disks Running at Optimum Throughput (dynamic load balancing)**
• **Additional Storage Space Available**

# Summary of Virtualization Key Points

- ✓ The EVA virtualization technology is state-of-the-art but does not require customers to "cross the chasm" to a new storage paradigm
- ✓ It is not risky or untried
- ✓ It **automates** array management processes – the ones that are repetitive, time consuming, or have the potential for human error
- ✓ It **aggregates** "things" like LUNs or RAID groups so in the end, customers have fewer things to manage
- ✓ It frees up valuable system administrators – SAVE MONEY!
- ✓ It reduces complexity in data center – WORK SMARTER!

# Enterprise Virtual Array

- Modular, scalable and highly available design
  - Redundant Power
  - Redundant Cooling
  - Distributed hot spare disk drives
  - Mirrored Cache with battery backup
- virtual RAID architecture
  - vRAID 0
  - vRAID 1
  - vRAID 5
- Full 2Gb fibre channel front to back
  - 1Gb fibre channel SAN compatible
- Centralized, unobtrusive "LiteTouch" manageability

# vRAID Levels

- Virtual Disk redundancy:
  - None (VRAID0): Data is striped across all physical disks in the Disk Group.
  - Moderate (VRAID5): Data is striped with parity across all physical disks in the Disk Group.
  - High (VRAID1): Data is striped mirrored across all physical disks in the Disk Group. Established pairs of physical disks mirror each other.

# EVA Family

- EVA 3000
- EVA 5000

# EVA Family Characteristics

- 2Gb FC Array Controller
- Automatic load balancing
- Virtualized data distribution
- Snapshot/SnapClone
- Dynamically expandable disk pools
- Load leveling & Auto Sparing
- Mirrored write-back cache
- NSPOF design
- New Performance Architecture
- Supports up to 240 drives
- Dual ported 2 Gb FC/FATA drives

# Enterprise Virtual Array (EVA) Platform EVA3000 & EVA5000

EVA3000 EVA5000

## Platform characteristics

- Midrange to enterprise customers
- Business-critical applications
- No single-point-of-failure
- High scalability – 8/35TB
- High performance
- State-of-the-art virtualization
  - vRAID, vSNAPS, snapclones
  - Auto pool expansion & tuning
- Compatibility across major operating systems
- Remote replication through HP's "Continuous Access"

## Benefits

- Virtualization technology automatically manages and optimizes the storage array, saving time and lowering your customer's overall TCO
- Saves disk space and money with StorageWorks Business Copy EVA
- Scales up to 35TB, offering consolidation for midrange customers
- Simplifies storage management



# EVA3000 & EVA5000 Review

| | EVA 3000  | EVA 5000  |
|---|---|---|
| Fiber | 2 Gb | 2 Gb |
| Capacity | 56 drives<br>8 TB | 240 drives<br>35 TB |
| Loop Switches | Not Required | Required over 56 drives |
| Controllers | Dual HSV100 | Dual HSV110 |
| Cost | 20% Less Expensive for initial enclosure | 20% More Expensive for initial enclosure |

# Upgrading EVA3000 to EVA5000

- eva3000 can be upgraded to eva5000 by swapping out the controllers and adding loop switches for > 4 drive enclosures
- Disk drives, drive enclosures, rack, and storage management appliance are reusable
- No retraining on software or device management
- The result: Outstanding Investment Protection

# FATA - New Fibre Hybrid Drive

- FATA – Fibre Attached Technology Adapted
- Innovative program from HP, Seagate and Hitachi
- Low-cost fibre channel storage for infrequently accessed data
- FATA Features:
  - ATA drive with dual 2Gb/sec Fibre interface
  - Hot-pluggable directly into FC enclosures
  - Some Enterprise-class functionality
  - Low I/O duty-cycle drive –20-30%
- Mixable with FC drives in the same shelf
- All EVA capabilities are supported
  - Business Copy, Continuous Access, …
- Competitively priced with today's ATA solutions
- Requires VCS V3.020

# Each drive type for a reason

Low-duty cycle drives

Enterprise-class drives

Performance

Price

**FATA**
**=> RPM than SATA**
**20-30% I/O duty cycle**

**SATA**
**7200 RPM**
**20-30% I/O duty cycle**

**PATA**
**5400 RPM**
**20-30% I/O duty cycle**

**SCSI**
**15K RPM**
**80-100% I/O**
**duty cycle**

**FC**
**15K RPM**
**80-100% I/O**
**duty cycle**

**SCSI**
**10K RPM**
**80-100% I/O**
**duty cycle**

**FC**
**10K RPM**
**80-100% I/O**
**duty cycle**

# Online/near-online storage Example: Business Copy

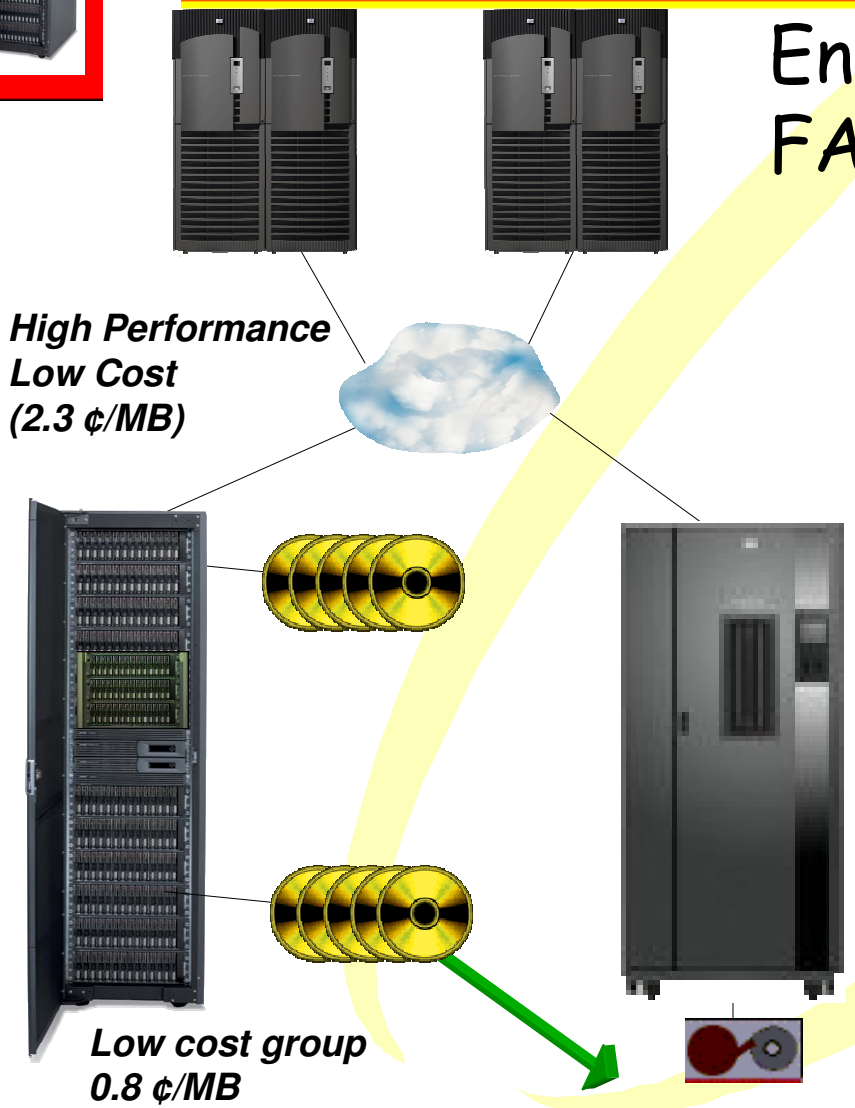


## Enterprise Virtual Array using FATA disk drives

- Business Copy
  - Create vsnap/snapclone
- Mount, snap and backup
  - HP Data Protector

# HP StorageWorks EVA3000
# Enterprise Functionality to the Mid-range

Highlights

- End-to-end 2Gb/s FC architecture
- Redundant HSV100 controllers
- Up to 4 drive enclosures per controller-pair
- Up to 56 drives per controller-pair
- Up to 335MB/s per controller-pair
- Great for moderate storage growth environments
- Scalable up to 8TB per controller pair
- Optional: business-data protection, business continuance and storage management software

# HP StorageWorks EVA5000
# The New Enterprise Storage Benchmark

Highlights

- End-to-end 2Gb/s FC architecture
- Redundant HSV110 controllers
- Up to 18 drive enclosures per controller-pair
- Up to 240 drives per controller-pair
- Up to 560MB/s w/one controller-pair
- Scalable up to 24TB/cabinet
- Scalable up to 35TB per controller-pair
- Optional: business-data protection, business continuance and storage management software

# Comparison Between EVA3000 and EVA5000

## Enterprise Virtual Array Platform EVA3000 and EVA5000

EVA3000

EVA5000

### Similarities:

- **same functionality**
- **same performance** (at same size)
- **same compatibility** (across ALL major operating systems: NT/W2K/Win2003, HP-UX, Tru64, Open VMS, Solaris, AIX, Linux, & NetWare)

### Differences:

**eva3000**

- **scales to 56 drives (8TB max)**
- **lower entry cost**

**eva5000**

- **scales to 240 drives (35TB max)**
- **back-end loop switches** (required when installing over 4 drive enclosures)

# Array Comparisons

| Array | eva3000 | eva5000 |
|---|---|---|
| # controllers | 2 | 2 |
| host interface | FC | FC |
| host ports | 4 | 4 |
| drive interface | FC | FC |
| drive enclosure density | 14 | 14 |
| maximum capacity (raw) | 8TB | 35TB |
| disk sizes / speeds (rpm) | 146GB/10k<br>73GB/15k,10k<br>36GB/15k,10k | 146GB/10k<br>73GB/15k,10k<br>36GB/15k,10k |
| maximum number of drives supported | 56 | 240 |

# Product Description

- Inside the cabinet, things are different
  - HSV Controllers
  - Fibre Channel Drive Enclosures
  - Fibre Channel Drives
  - FC-AL Cabling
  - FC Loop Switch

# EVA5000 2C2D Configuration

- 2C2D configured to be easily expanded
  - CTO in factory up to 2C5D with or without FC loop switches
  - Field upgradeable all the way to 2C12D (with FC loop switches)

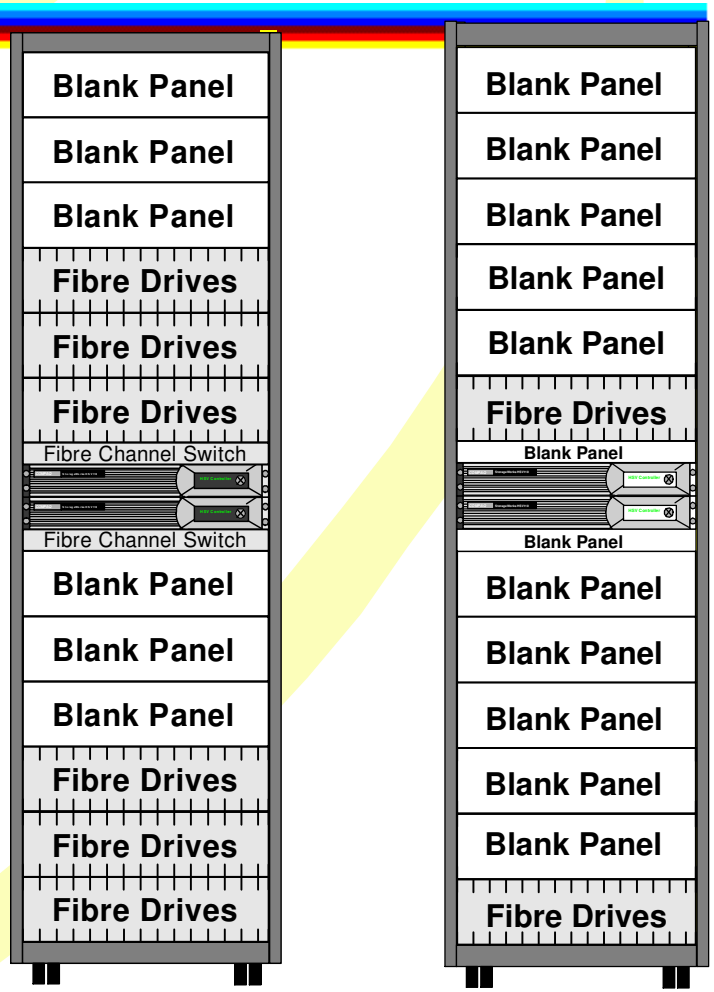


2C6D-B
with FC Loop Switches

2C2D-B
without FC loop switches

# EVA5000 2C6D Configuration

**Model 2C6D**

(1) 42U Modular storage cabinet
(1) M3220 Enclosure pair w
(2) HSV110 Controllers
(2) Cache batteries per enclosure
(6) 14-bay FC enclosures
(17) Int FC Cables
(7) 2-port EMU boxes
(8) AC strips
(2) 0u PDUs

**Disks ordered separately**

3 TBytes (36GB)
6 Tbytes (72GB)
12 TBytes (146GB)

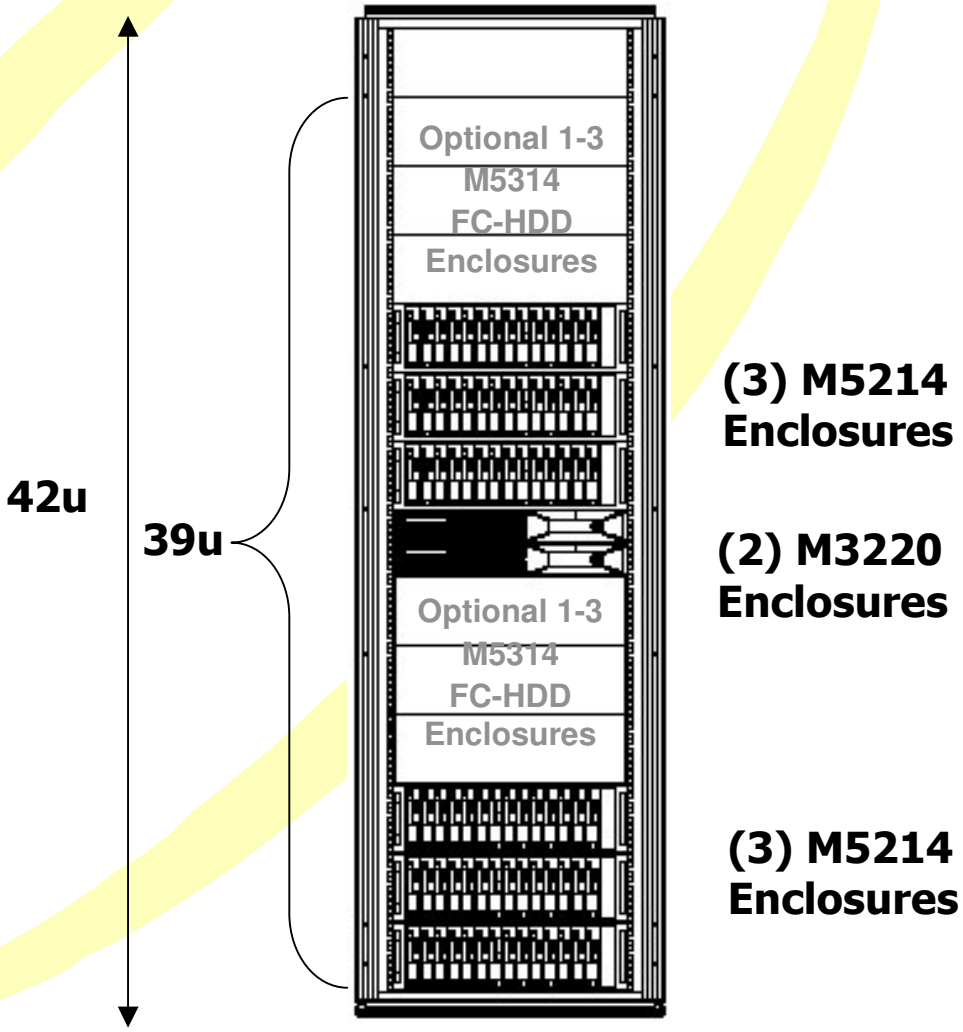

# EVA5000 8C8D
## Write Performance Model

**Model 8C8D**
(1) 42U Modular Storage Cabinet
(4) M3220 Enclosure pair w
(8) HSV110 Controllers
(8) Cache Batteries
(8) 14-bay FC Enclosures
(36) Int FCCables
(7) 2-port EMU boxes
(8) AC strips
(2) 0u PDUs

**Disks ordered separately**

16.0 TB (112x 146GB) - Max Drive Configuration

**This model is for "memory dump" (high speed, large block, consecutive 100% write) operations typically for Scientific Applications*

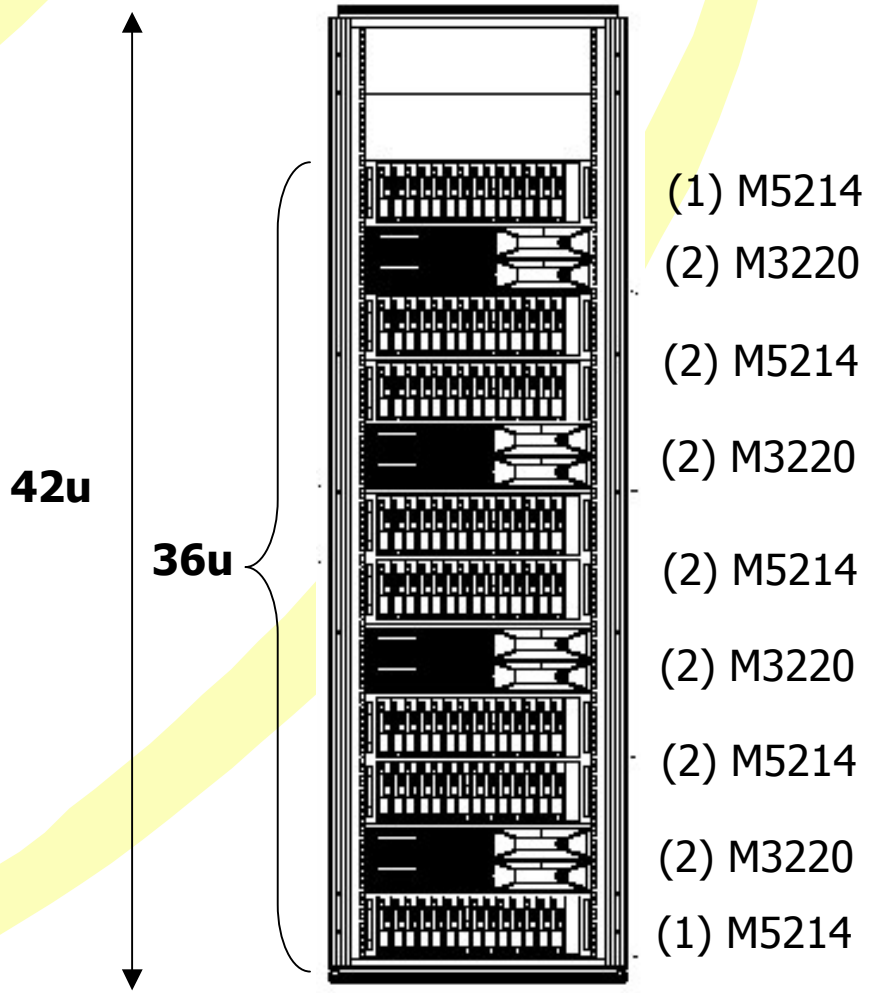

# EVA5000 - 2C12D Configurations

**Model 2C12D**

(1) 42U Modular storage cabinet
(2) M3220 Enclosure pair w
(2) HSV110 Controllers
(2) Cache batteries per enclosure
(12) 14-bay FC enclosures
(37) Int FC cables
(7) 2-port EMU boxes
(8) AC strips
(2) 0u PDUs

**Disks ordered separately**

6 TBytes (36GB)
12 Tbytes (72GB)
24 Tbytes (146GB)

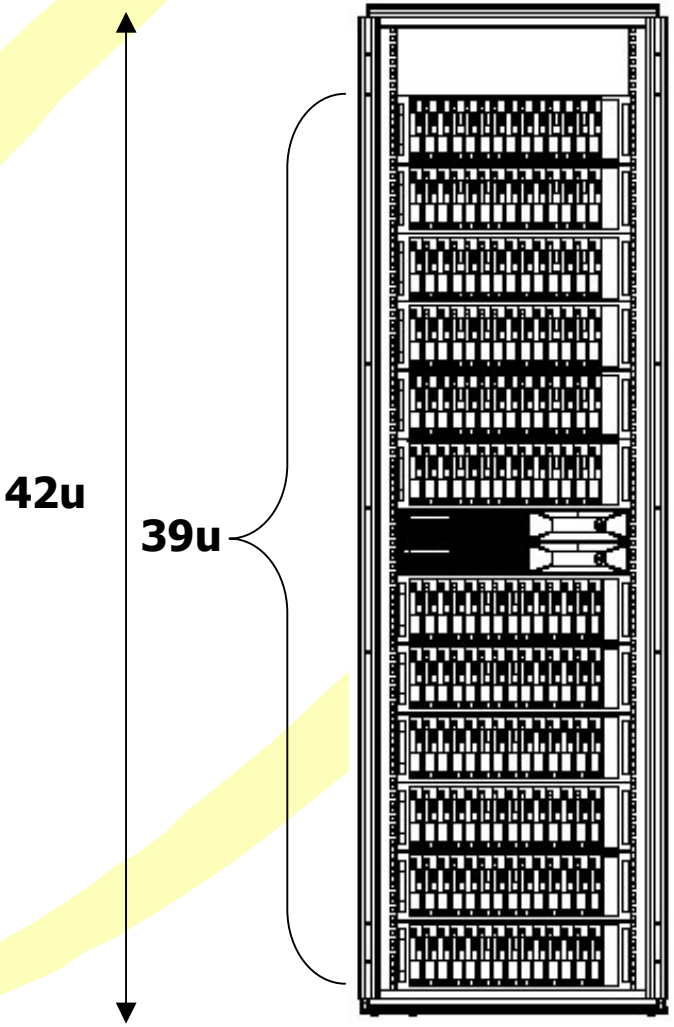

# 2C12Ds with 0C12D Expansion Cabinet

Expansion Cabinet

- 70 Terabytes
  - 480, 146GB FC disks
- 4 controllers
- Single footprint: 17.7 $ft^2$ (1.5 $m^2$)

# EVA3000

# EVA3000 Configuration

**EVA3000 Base model**

(1) 42U Modular storage cabinet
(1) M3200 Enclosure pair w
(2) HSV100 Controllers
(2) Cache batteries per enclosure
(2) 14-bay M5114 FC enclosures
(7) 2-port EMU boxes
(8) AC strips
(2) 0u PDUs

**Disks ordered separately**

2 TBytes (36GB)
4 TBytes (72GB)
8 TBytes (146GB)

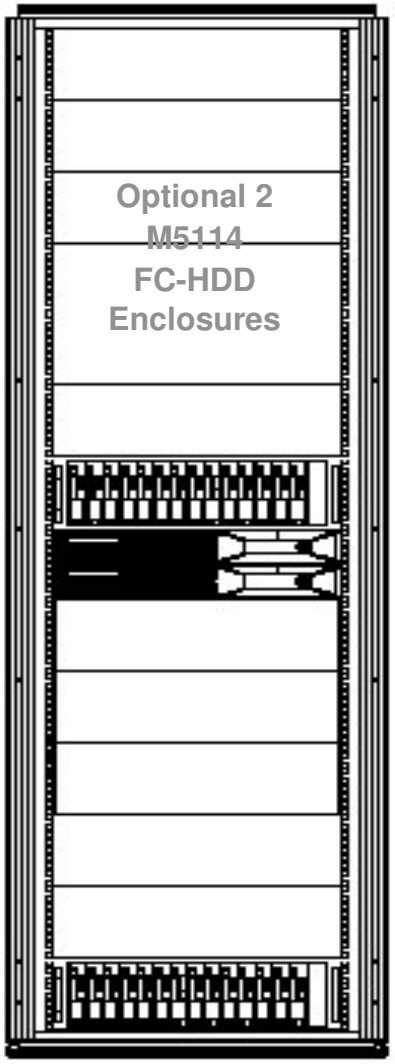


**(1) M5114 Enclosure**

**(2) M3200 Enclosures**

**(1) M5114 Enclosure**

# EVA3000 architecture

- Base eva3000 module:
  - Dual HSV100 controllers in an M3200 enclosure
    - Four optical 2Gbit front-end loops to host
    - Two copper 2Gbit back-end loops to drive enclosures
    - 2GB cache and dual power supplies
  - Two M5114 FC drive enclosures
  - One 42U 10,000 series graphite metallic rack
  - Redundant zeroU PDUs
  - Scales up to 56 drives (with the addition of two drive enclosures to the initial two enclosure subsystem)

# EVA5000 architecture

- Base eva5000 module:
  - Dual HSV110 controllers in an M3220 enclosure
    - Four optical 2Gbit front-end loops to host
    - Four optical 2Gbit back-end loops to drive enclosures
    - 2GB cache and dual power supplies
  - M5214 FC drive enclosures
  - One 42U 10,000 series graphite metallic rack
  - Redundant zeroU PDUs
  - Scales up to 240 drives (with the addition of expansion enclosure)

# M32x0—Array Controller Enclosures

- High performance HSV1x0 controllers
  - High performance Power PC microprocessor
  - Two 2Gb/s "ready" FC-Switch Fabric host ports
    (2Gb/s switches and HBAs not available at first release)
  - Two or Four 2Gb/s FC-AL device ports
    - Arranged in redundant pairs
    - Data load / performance is balanced across a pair of device ports
    - Supports up to 240 disks (120 disks per pair of device ports)
  - 1GB cache per controller, mirrored, with battery backup
  - 2Gb/s FC cache mirroring port (device ports as backups)

# Controller Enclosure Pair—Front View


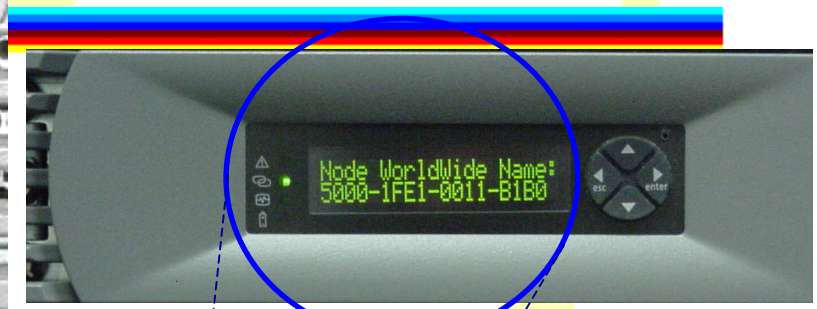


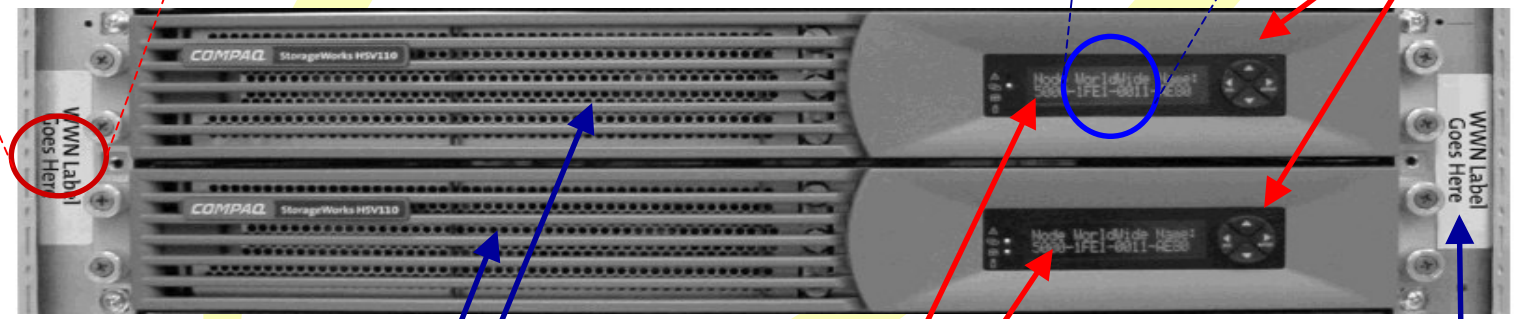


Push Buttons

Plastic Bezels

LCD Displays

World Wide Name

# Controller Enclosure - Back View

- 1.5U single controller enclosure (offered only as 3U pair)

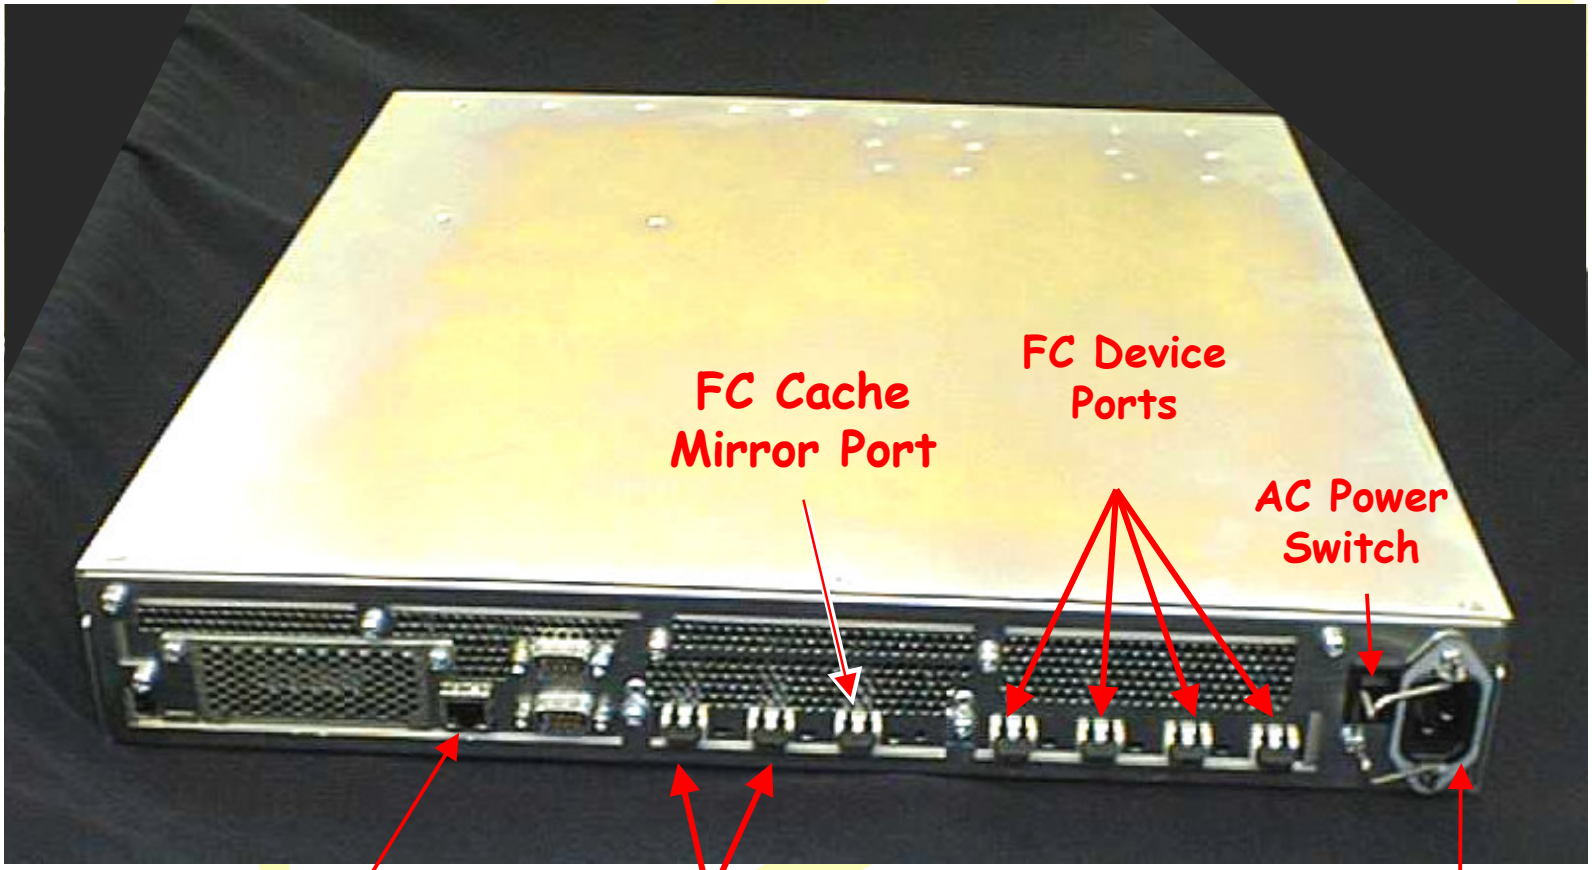

# Controller Enclosure - Front View, Fans and Cache Batteries Access Panels

# M5214—FC Drive Enclosures - Front View

- 3U disk enclosure
- Dual redundant active-active 2Gb/s FC busses
- Fourteen 1-in. FC disks per enclosure (4 drive minimum)

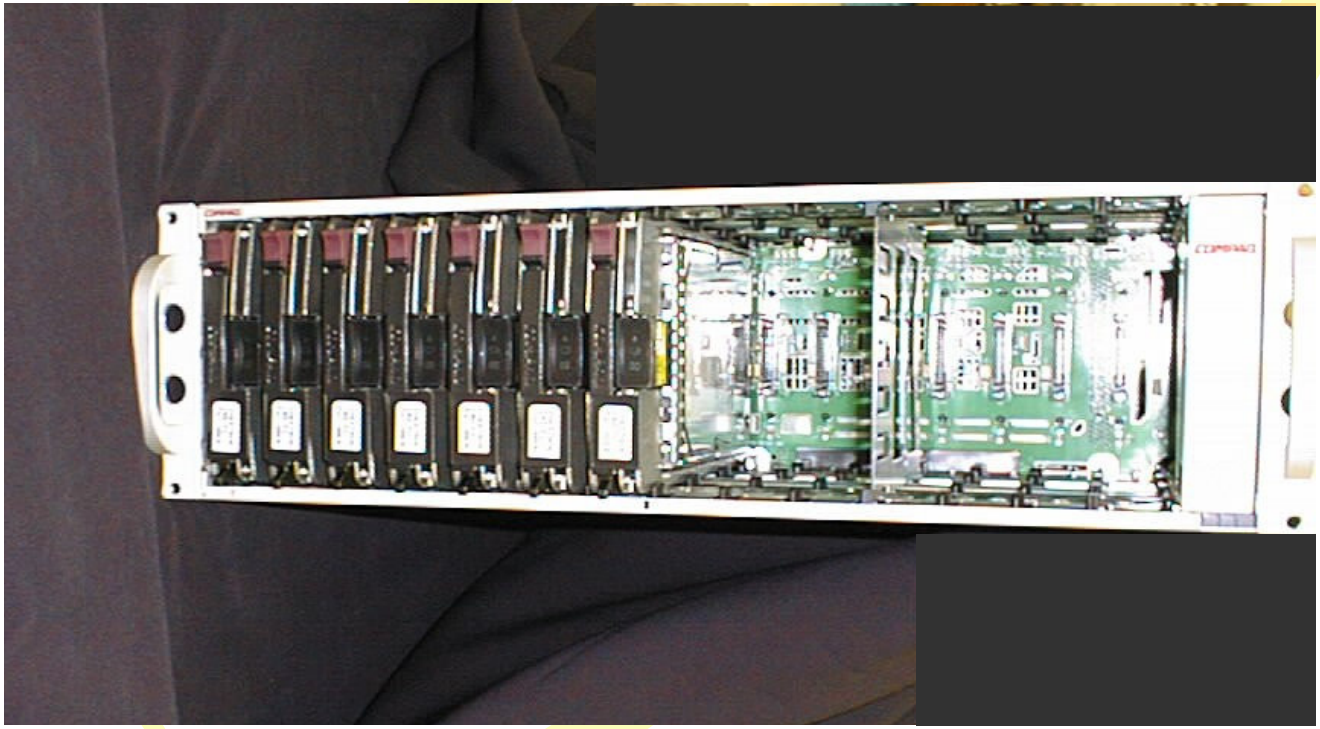

# FC Drive Enclosure - Back View

- Environmental Monitor Unit
- Dual 2Gb/s FC I/O module
  - B (left-side) A (right-side)
- Dual 500 watt redundant hot plug power supplies and fans

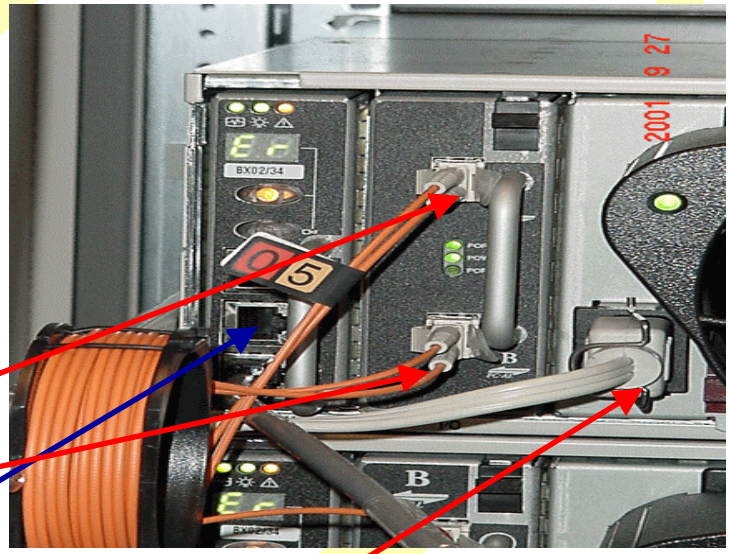

# FC Drive Enclosures - Rear Cabinet View

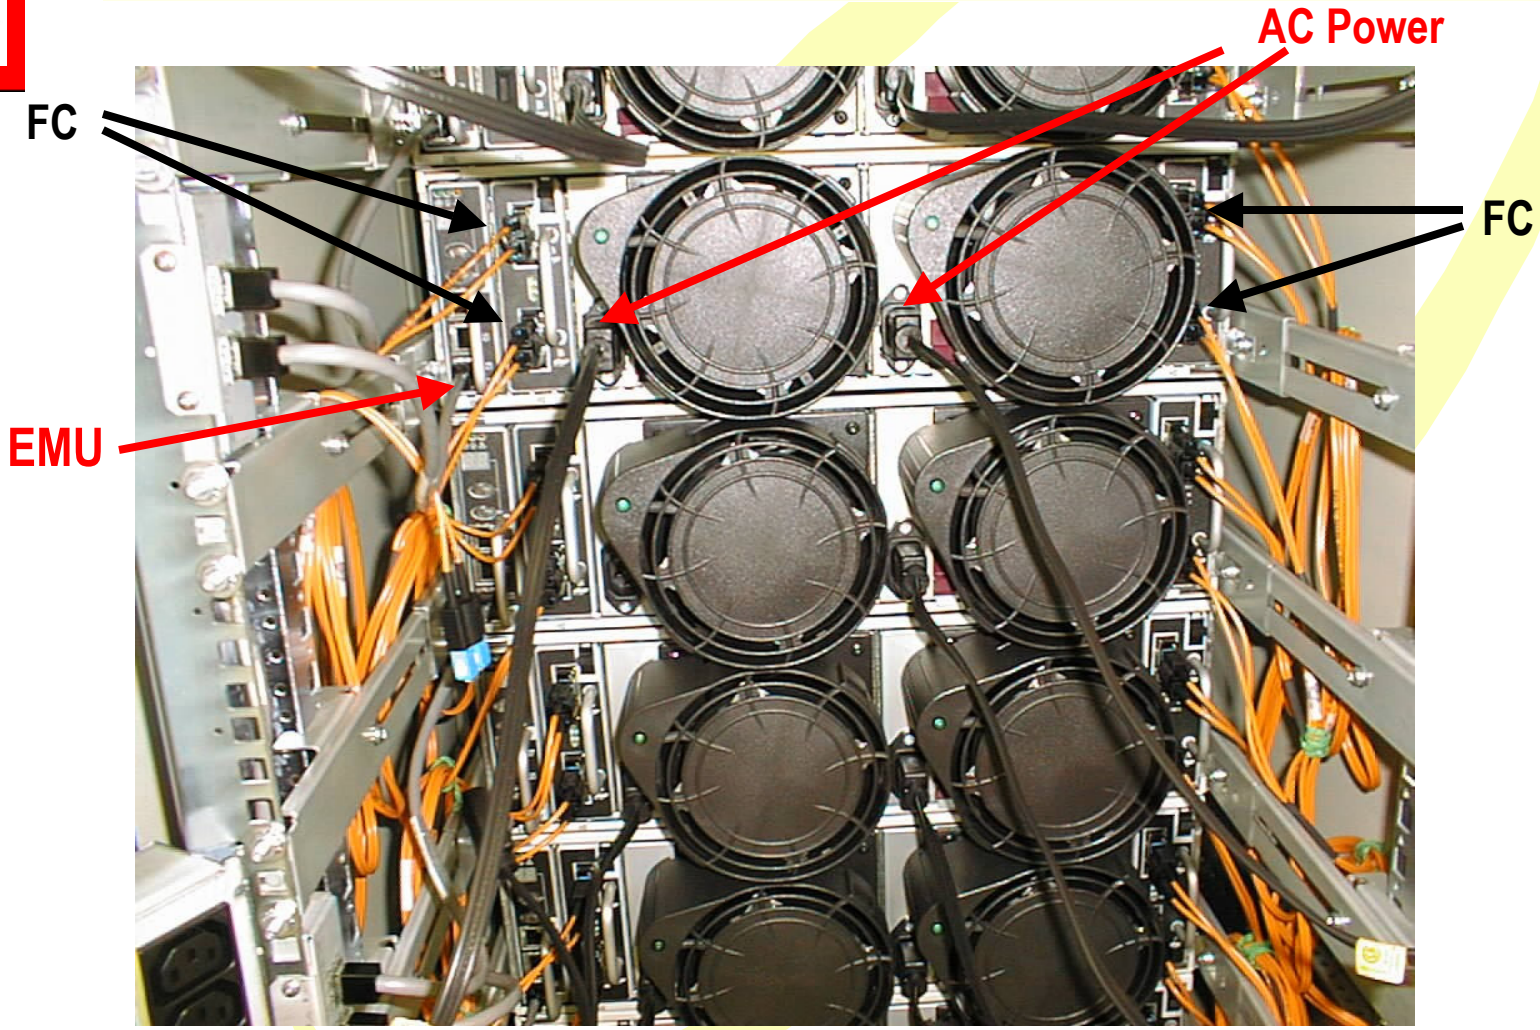

# FC Disk Drives

- Dual-ported 2Gb/sec FC-AL
- 72B and 146GB 10K rpm
- 36GB and 72GB 15k rpm
- Up to 120 drives to be supported per FC-AL pair


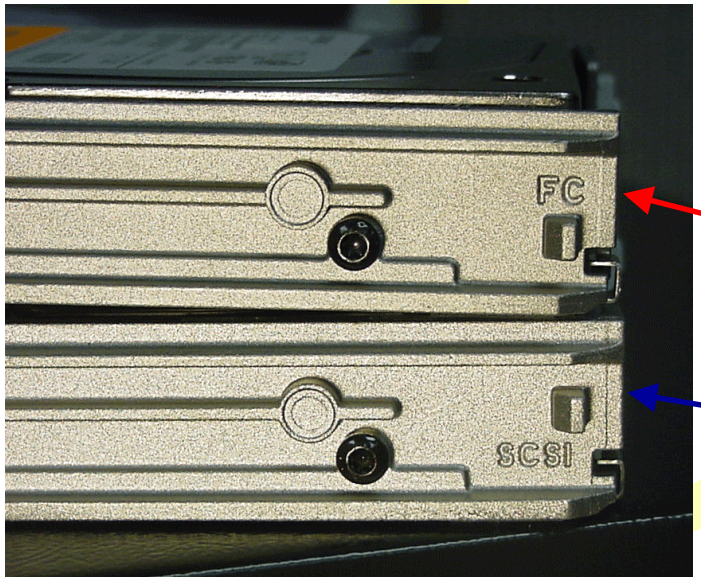


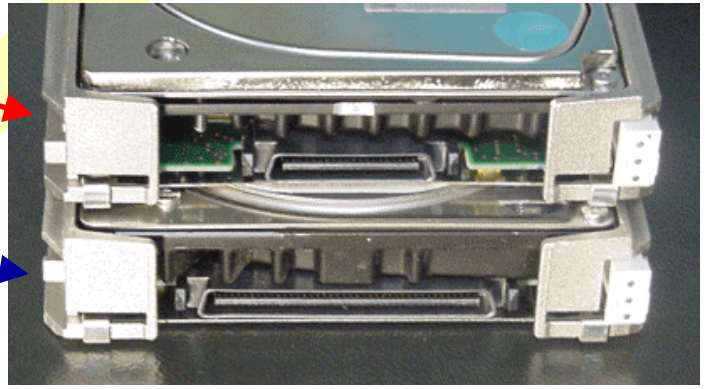

# Backend FC Loop Switches

- (4) four 1u half-rack-wide 12-port (11 utilized with expansion cab) FC loop switches
- improves diagnostic/service ability:
  - ability to add or remove shelves
  - ability to add expansion cabs

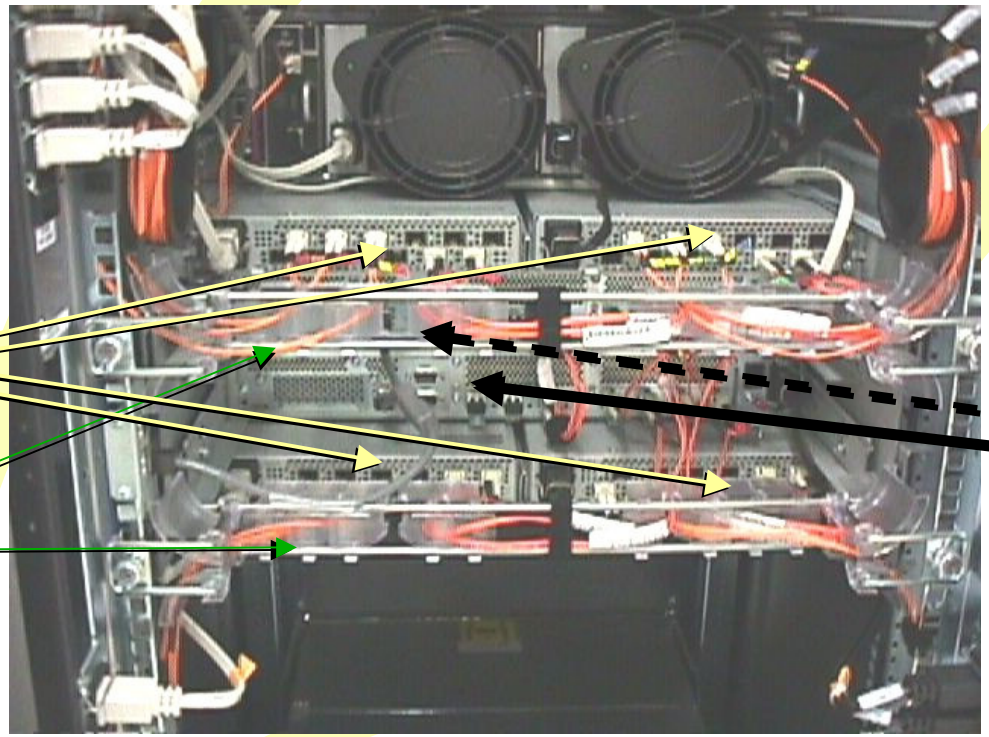

# Cable Routing – EVA-2 Cabinet

**E2 2C6D-A, -B Config**

- 2 controllers
- 6 drive shelves
- up to 84 drives
- 2  0U PDUs
- 42 U cabinet
- **4 FC loop switches**

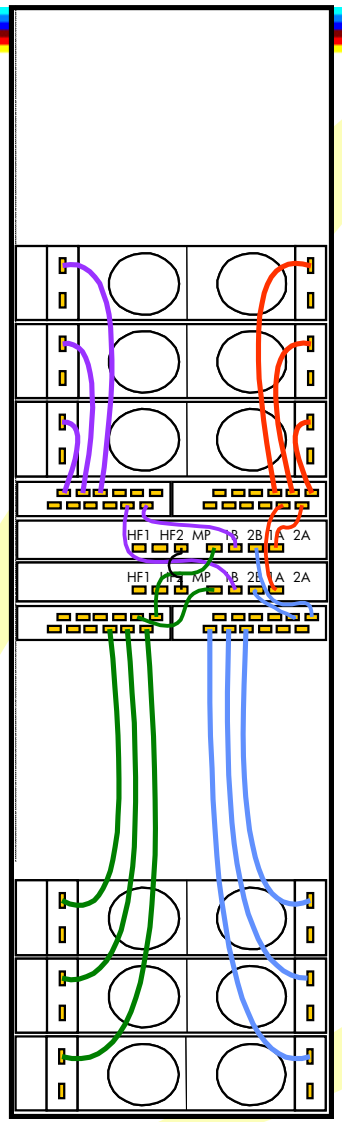


**E2 2C12D-A, -B Config**

- 2 controllers
- 12 drive shelves
- up to 168 drives
- 2  0U PDUs
- 41U cabinet
- **4 FC loop switches**

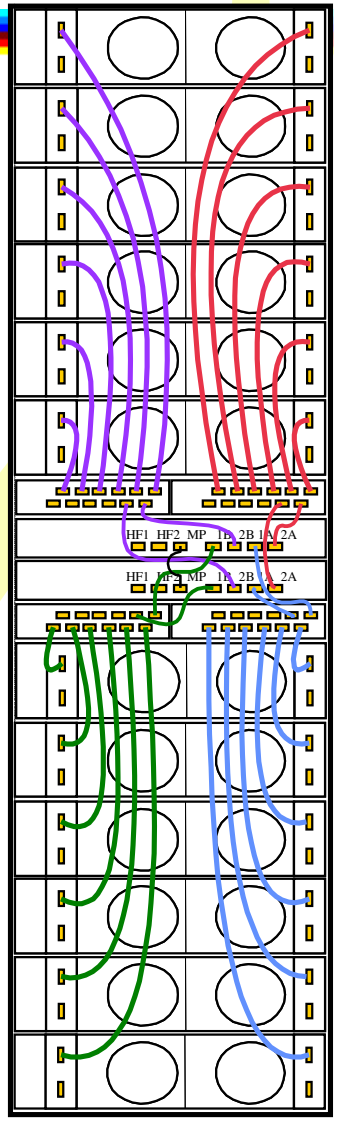

# Cabinet Cabling

- Fibre Channel, Power, and Cabinet Area Network (CAN)

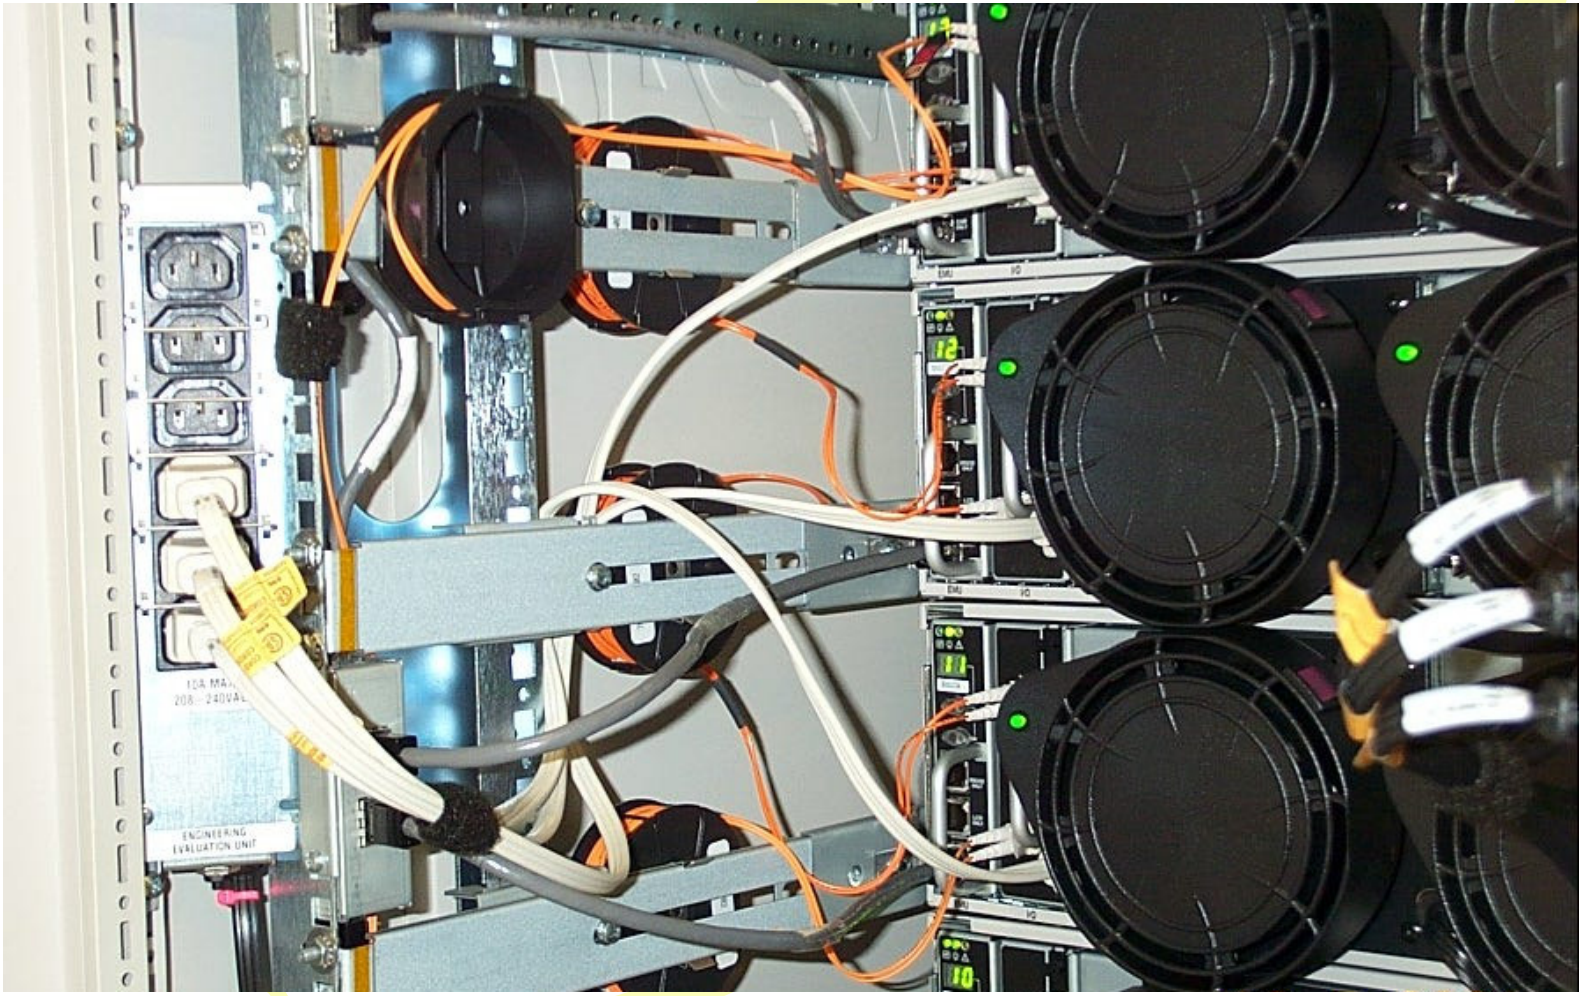

# HSV Storage System - Cabinet Bus

- Cabinet Area Network (CAN)
- Requires Terminator at top and bottom of bus
- Provides interconnection for Environmental Monitoring Units (EMU's) and Controllers
- Assigns Enclosure Number / Address to the EMU's
- Provides support for FC-AL addressing
- Provides common path for event logging

# CAN (EMU) bus cabling

CAN = Cabinet Area Network

- Designed for automobiles
- Resistance to noise

- CAN bus junction boxes
  - each 6U increments
  - whether fully populated or not
  - between enclosure mounting rails

- Each RJ-45 connector provides a different enclosure number

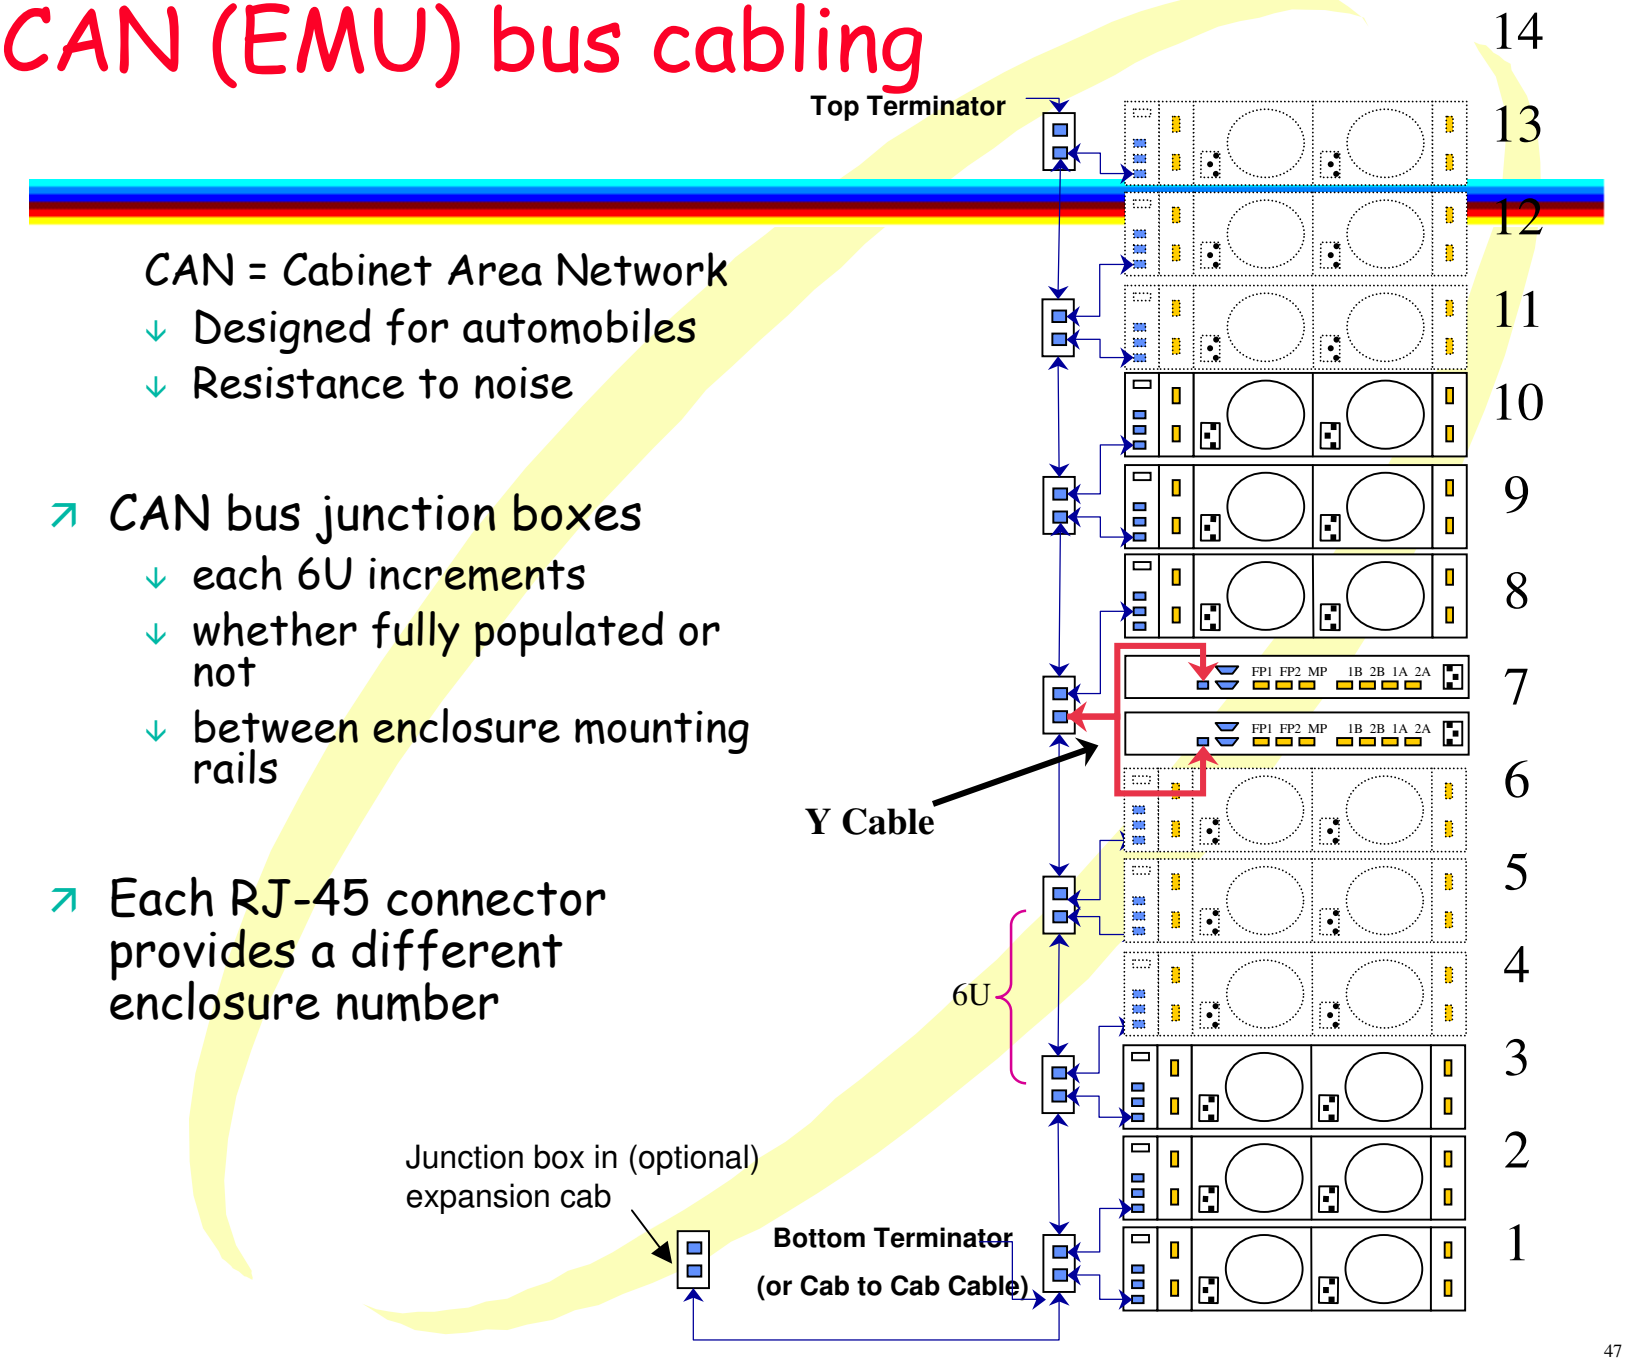

# EMU Monitors & Controls

- Disk Drives
- Power Supplies
- Fibre Channel Transceivers (GBICs)
- EMU
- Fans
- I/O Modules
- Temperatures (Alerts and Errors)
- Voltage Sensors (Alerts and Errors)
- Current Sensors (Alerts and Errors)

# What is SES?

- SCSI-3 Enclosure Services (SES)
- SCSI-3 Specification:
  - A family of documents
  - SES is a member of the family
- SCSI-3 Primary Commands (SPC) referenced by SES

# SES EMU Monitors

- Disk Drives
  - Fibre Channel link rate (1G vs. 2G)
  - Fibre Channel addresses relative to physical location
  - Presence vs. absence
  - Faults reported by drives
  - Bypass state
  - Drive WWN
- Power Supplies
  - Presence vs. absence
  - AC presence vs. absence
  - Emergency shutdown status

# SES EMU Monitors

- Fibre Channel Transceivers (GBICs)
  - Presence vs. absence
  - Faults
  - Missing cables
- EMU
  - Internal hardware diagnostics
  - Enclosure number conflicts
  - Environmental data validation
  - Enclosure WWN and serial number
  - Power shutdown
  - Error code translation to text
  - Code load

# SES EMU Monitors

- Fans
  - Presence vs. absence
  - Speed
- I/O Modules
  - Presence vs. absence
  - Enclosure link rate (1G vs. 2G)

# SES EMU Monitors

- Temperatures (Alerts and Errors)
  - Inlet temperature (at EMU)
  - Power supplies
  - Drives
  - Meltdown protection Voltage Sensors (Alerts and Errors)
  - Supply output voltages for 5V
  - Supply output voltages for 12V
- Current Sensors (Alerts and Errors)
  - Supply output currents for 5V
  - Supply output currents for 12V

# SES EMU Controls

- Disk Drives
  - Bypass state
  - Element identification (location)
  - Drive spinup
- Power Supplies
  - Emergency shutdown
  - Element identification (location)

# SES EMU Controls

- Fibre Channel Transceivers (GBICs)
  - Element identification (location)
- EMU
  - Power shutdown
  - Error code translation to text
  - Code load
  - Element identification (location)
- Fans
  - Speed
  - Element identification (location)

# SES EMU Controls

- I/O Modules
  - Element identification (location)
- Audible Alarm
  - Severity
  - Muting

# SES EMU Operator Interface

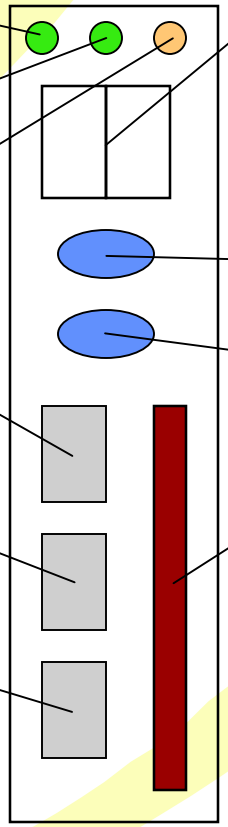

Heartbeat – Flashing is good

Power – On is good

Fault – Off is good

Two-character display – A two-digit number is good
**Er** indicates an error

Top push button with backlight

Bottom push button with backlight

RJ-45 connector – RS232 Only

RJ-45 connector – LCD Only

RJ-45 connector – Can Only

Handle

# SES EMU Operator Interface

- Push button terminology
  - Push and hold
    - Push the button and hold it depressed until the operator interface changes (within about 2 sec.)
    - Release the button
  - Push (and release)
    - Push the button
    - Release the button immediately

# SES EMU Operator Interface

- Push the bottom button to cycle among the following choices:
  - **En** – Examine the enclosure number
  - **Li** – Modify the enclosure's loop id
  - **rG** – Modify the enclosure's reporting group #
  - **Au** – Enable or disable the audible alarm
  - **Er** – Examine the currently active alarm(s)
    - **Er** is not available when there are no currently active alarms

# SES EMU Operator Interface

- **Er** – Examine the currently active alarm(s)
  - Push and hold the top button to select the next currently active alarm
    - When the top button's backlight goes out, the last active alarm has been selected
    - Push the bottom button to return to the **Er** display
  - Push the top button to cycle among the digits of the alarm's numbers
  - Represented as **#.#.##.##**

# SES EMU Alarm Codes (Example)

- **0.3.02.01 N** Fan speed alert; replace fan soon
  - **0.3.02** indicates fan #2
  - **01** indicates fan alarm #1
  - **N** indicates that the severity is Non-Critical
    - **U** – Unrecoverable (most severe)
    - **C** – Critical
    - **N** – Non-Critical
    - **I** – Informational (least severe)
  - Alarm description; remedy

# SES EMU Unusual Displays

- **Ld** – Code load in progress or incomplete
  - Removing power disables the EMU – RS232 cable required for repair!
- **-A** – EMU firmware application code is missing
  - RS232 cable required for repair!
- **--** – Both push buttons inadvertently held/pushed while inserting the EMU
  - Remove and reseat the EMU without pushing buttons
- **8.8.** – EMU is not completely seated
  - Remove and reseat the EMU

# EMU Error Code

- Current Errors
  - Audible—beeping
  - Code available on EMU display
  - 3 part error code xx-xx-xx
  - May have multiple errors available
- Error log
  - Available through HSV Element Manager event page
  - 62 entries per EMU (enclosure)

# EMU Error Code

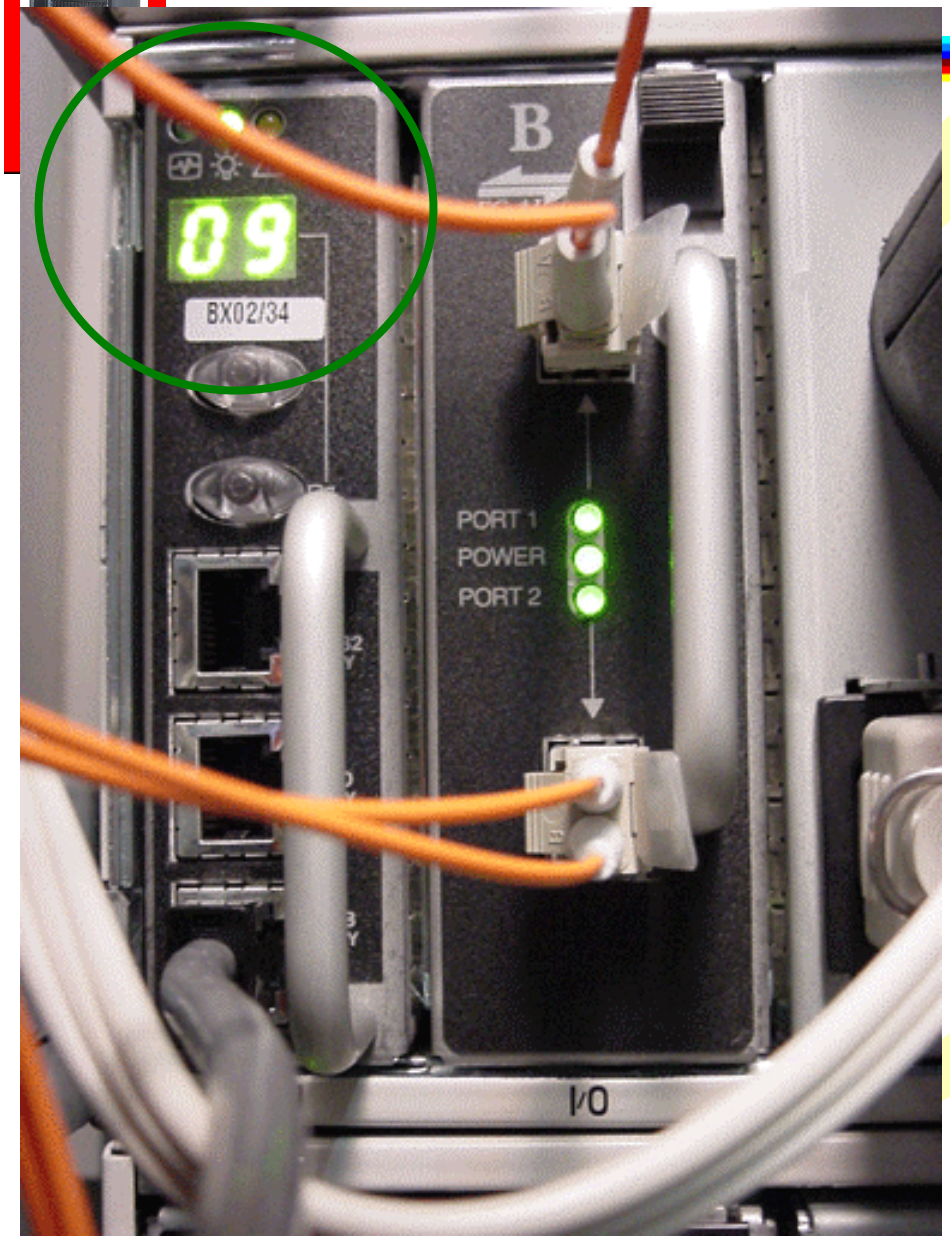


# Power Distribution Modules (PDMs)

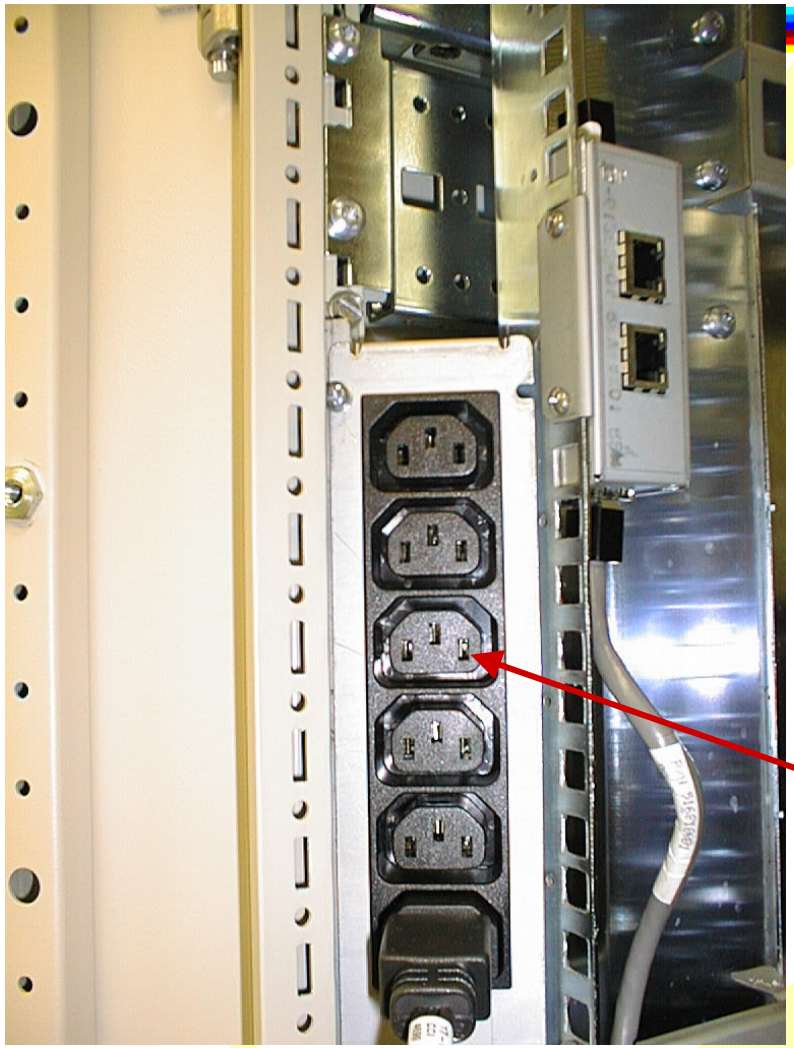

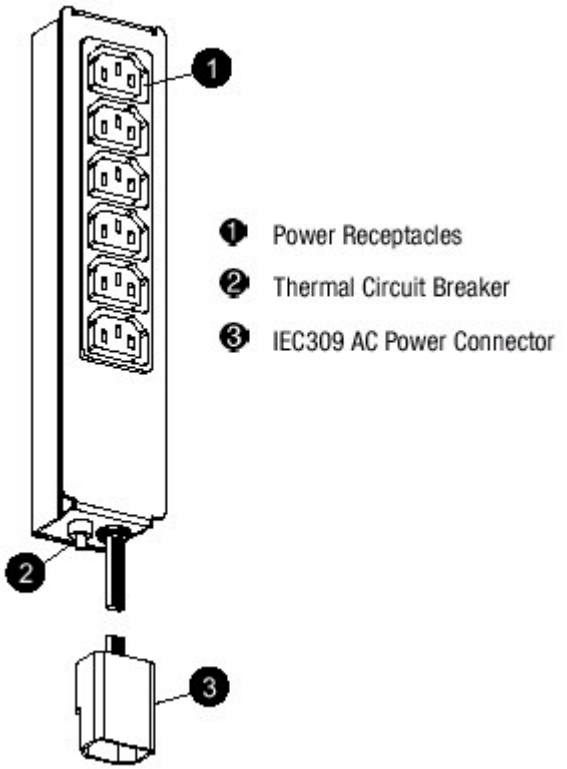


PDM

# Zero-U PDUs

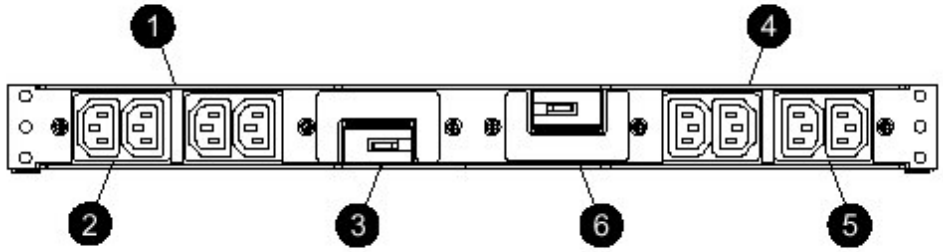


1. PDU 1
2. PDU 1 AC receptacles
3. PDU 1 Circuit Breaker
4. PDU 2
5. PDU 2 AC receptacles
6. PDU 2 Circuit Breaker

- New Zero-U PDU saves valuable cabinet space
- Dual power provides independent and redundant power paths (220-240v, 30amp)

# Zero-U PDUs

1. Unscrew the spring-loaded screws
2. Fold PDU away from the rack towards the floor

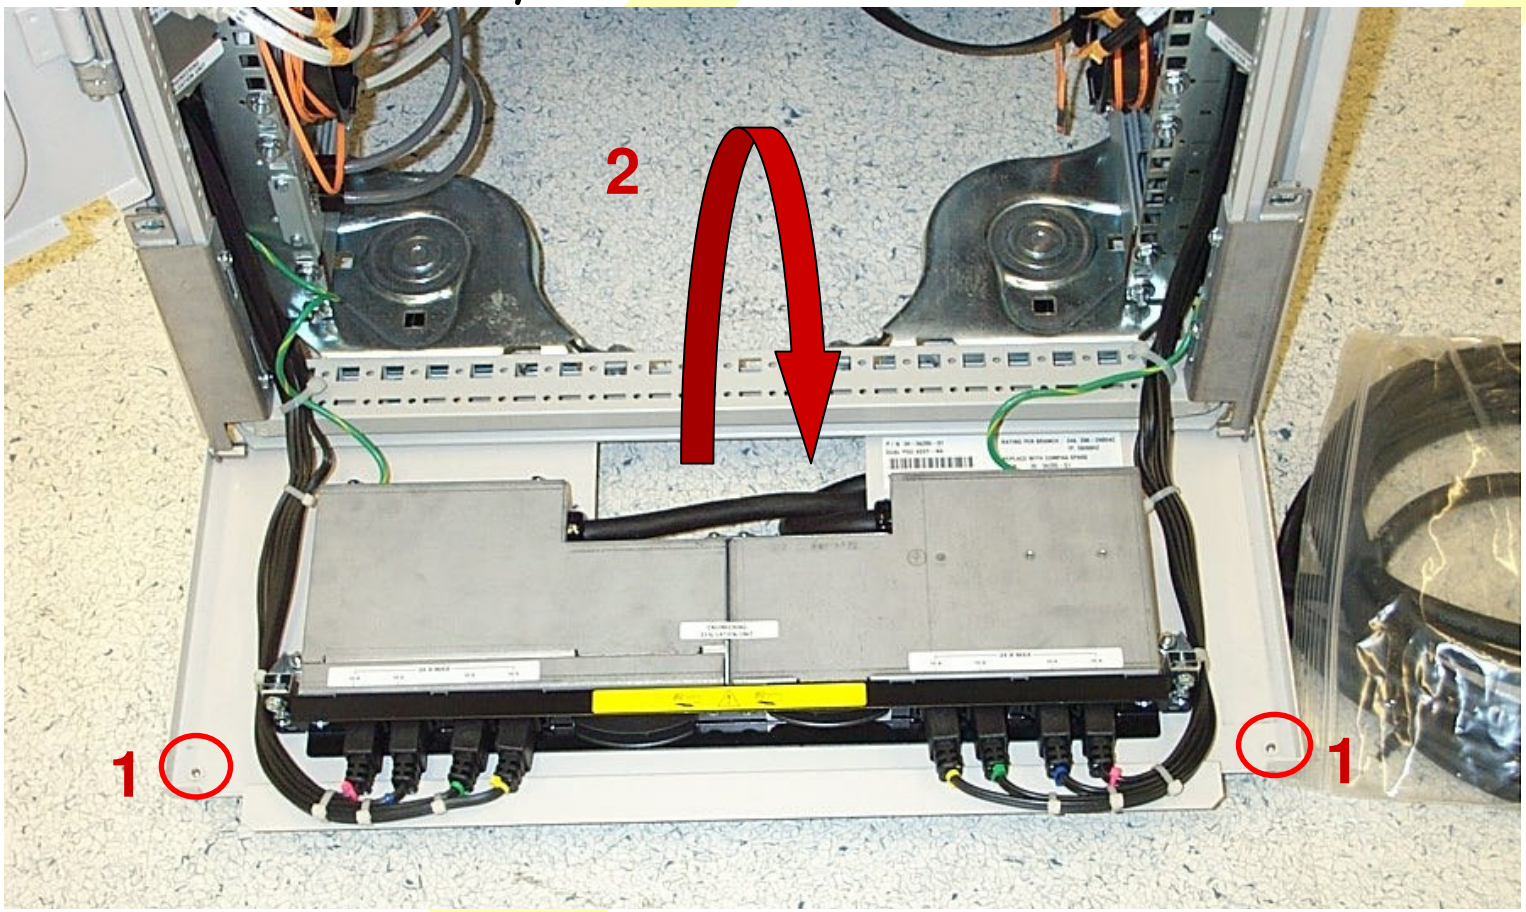

# Power Supply and Cooling Requirements

- North America
  - Single phase
  - NEMA L6-30R receptacle
  - 3 wire
  - 208V to 240V 60Hz 30A
- Europe
  - Single phase
  - 2 pole IEC 309
  - 3 wire
  - 220V to 240V 50Hz 32A
- Heat Dissipation
  - 12,708 BTUs per hour

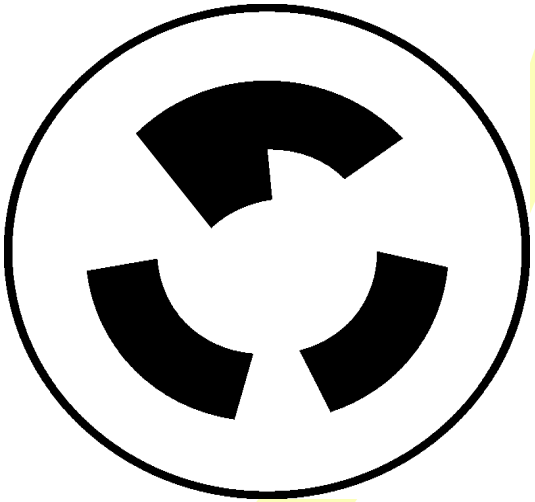

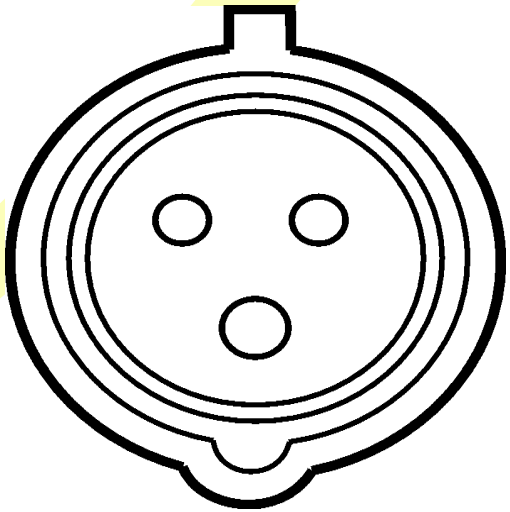

# EVA Software Overview

# Sizing HSV Disk Groups

- Hardware versus software capacities
  - Physical 1000 000 000Bytes = 1GB
  - Software 1073 741 824Bytes = 1.07GB Physical ($2^{30}$)
  - ~ 7% Variance → 1GB Physical = 0.93GB Software
- System metadata overhead — 0.2%
  - System metadata
  - MLD—HSV Element Manager metadata
  - Virtual Disk metadata
- Vraid overhead
  - Vraid0 — 0% (1 block for every 1 block usable)
  - Vraid1 — 50% (2 blocks for every 1 block usable)
  - Vraid5 — 20% (1.25 blocks for every 1 block usable)
- Snapshot working space
  - Snap — depends on rate of change of original data
  - Snapclone — same physical capacity as virtual disk
- Spare capacity
  - 2 X physical capacity of the largest physical disk X protection selected

# Virtualization Controller Software (VCS)

- Software (Firmware) for the HSVx0 controllers
- Virtualization is integrated throughout VCS:
  - Provides improved performance by spreading data volumes (LUNs or Virtual Disks) across many more disk drives or spindles
  - Load leveling of data across a LUN / Virtual Disk is automatic and helps eliminate "hot spots" which could otherwise become performance bottlenecks
  - Dynamic expansion
    - LUNs / Virtual Disks can be expanded on the fly for OSs that support it
    - Disks can be added to disk pools on the fly
  - Distributed sparing
    - Allows allocation of space per disk group to recover from physical disk failures

# VCS Versions

- V2.xxx
  - Current version for both eva3000 and eva5000
  - Snapshot and Snapclone support
- V3
  - Current version for eva5000
  - Remote replication support

# Virtualization Controller Software

- Data Protection techniques (optional purchase)
  - Two Versions of Snapshot:
    - Virtually Capacity Free (**On-Demand**) — Which can save customers a lot of disk space and money
    - Traditional (**Fully Allocated**) — Reserves exact space size as original LUN / Virtual disk
  - Virtually Instantaneous Snapclone:
    - Starts as a snapshot and becomes over time a clone
    - Gives access to the clone immediately without waiting for the clone copy "completion"
    - Can save customers time and money

# Snapshot versus Snapclone

| | Description | Pro's | Con's |
|---|---|---|---|
| Snapshot Space efficient | Pointer based Copy before Write Allocate space on demand | Space efficient (allocated on demand) | Overcommit problem |
| Snapshot Space guaranteed | Pointer based Copy before Write PreAllocate space on creation | No Overcommit problem | Space inefficient (allocated right away) |
| Snapclone | Same as Snapshot space guaranteed, but now with background process to separate VD. | No Overcommit problem Repeatable, separate VD's | Space inefficient Consumes some background process time |

# Virtual Storage Pools (virtual spare)

- Virtual Disk blocks automatically regenerated to restore redundancy

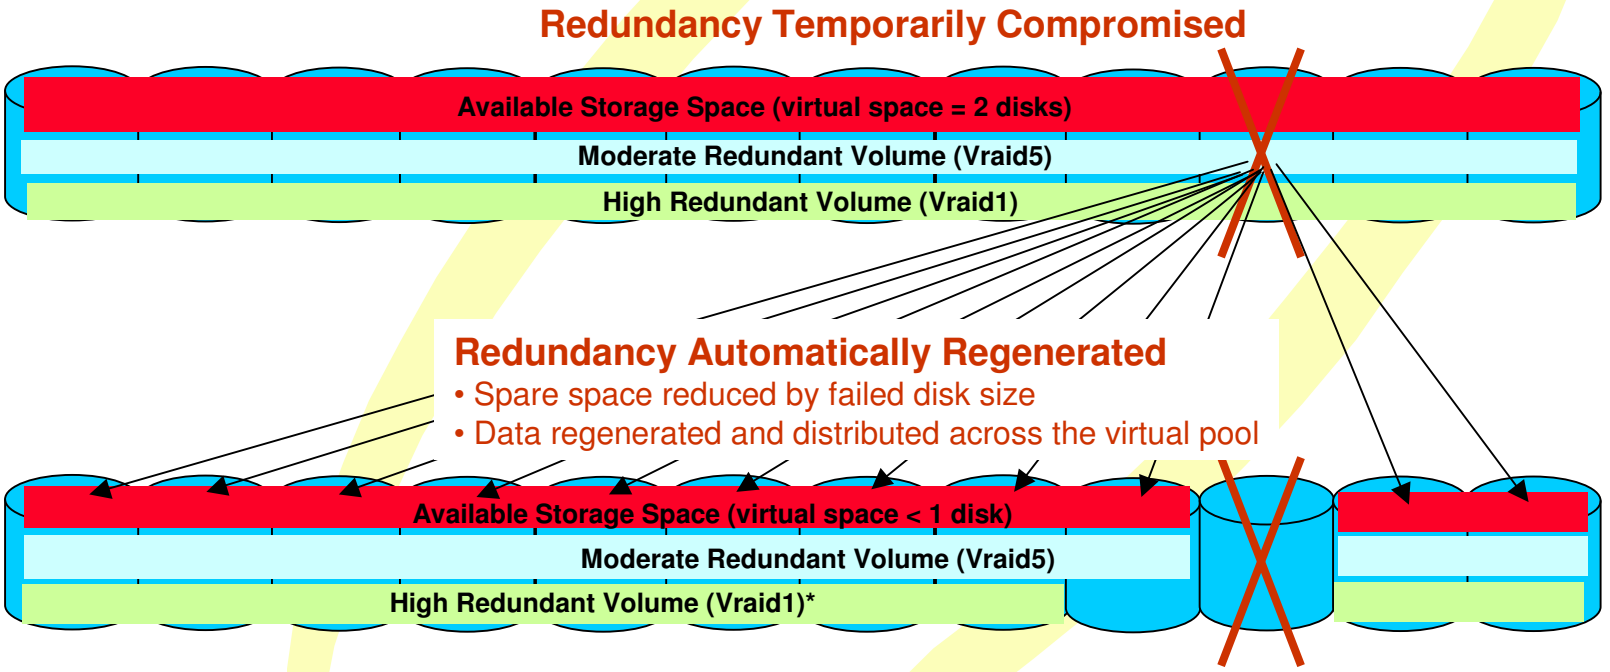


*RAID V1 uses even numbers of disks

# Distributed Sparing

- Allocated space per Disk Group to recover from physical disk failure(s) in that Disk Group
- Choices—None, Single, Double

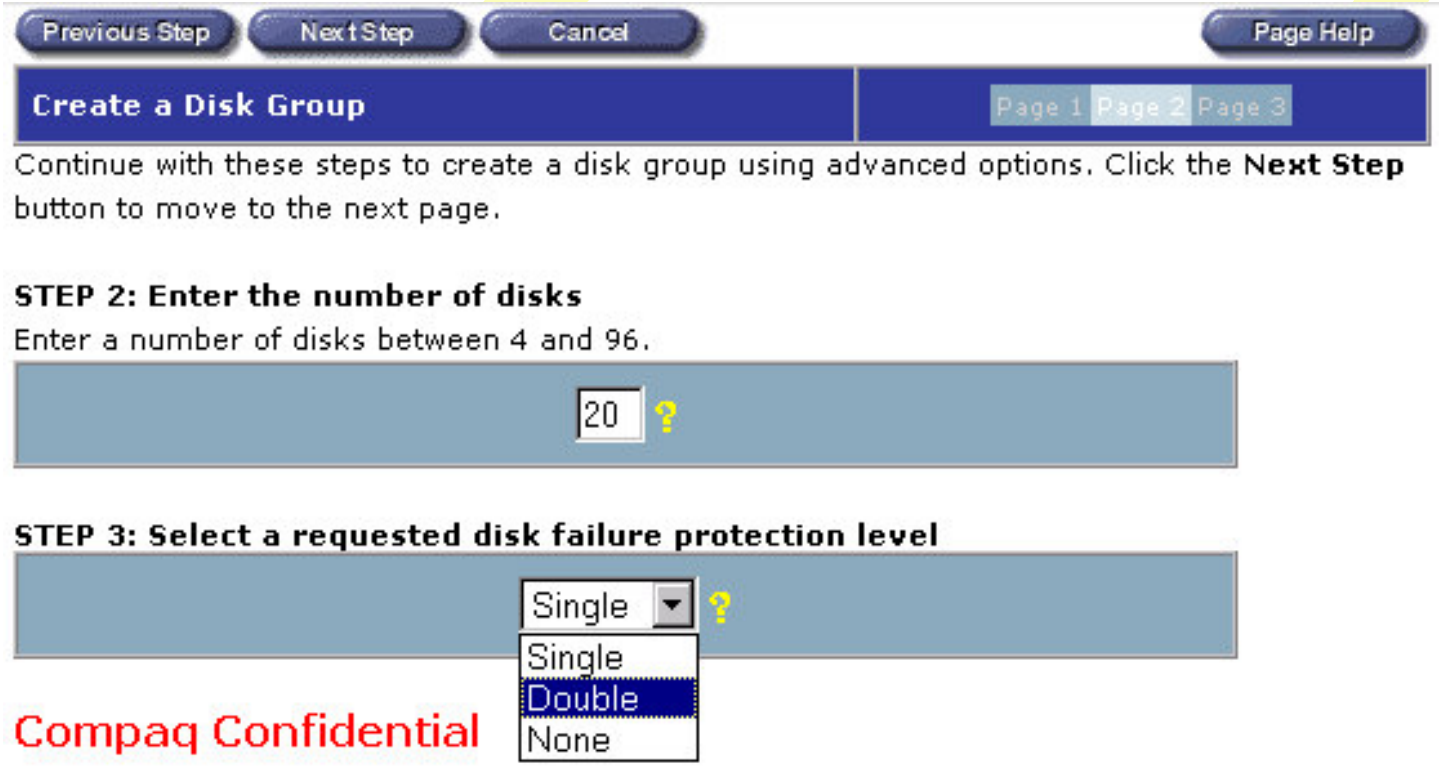

# Distributed Sparing

- No longer spare in separate spindles
- Chunks allocated, but not dedicated as spares, on all disk drives of disk group to survive 1 or 2 disk drive failures.
- Allocation algorithm
  - Single (1) = capacity of 2 * largest spindle in disk group
  - Double (2) = capacity of 4 * largest spindle in disk group

# Distributed Sparing
# HSV110 Virtualization :
## The Easy to Understand scenario

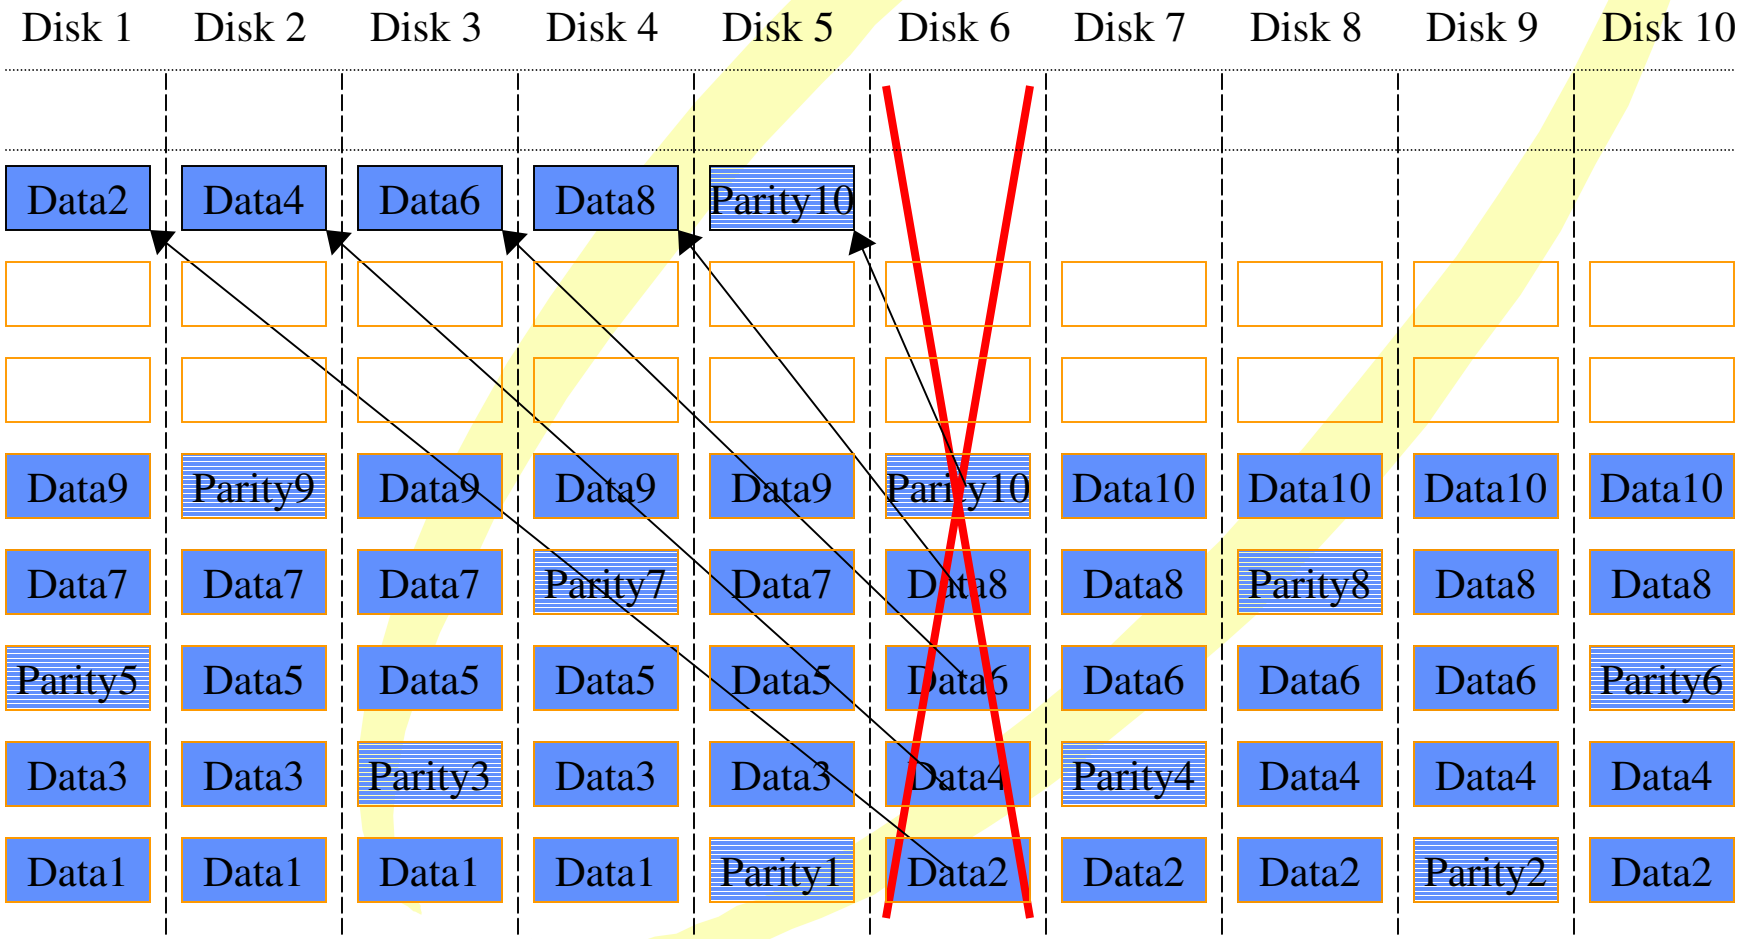

# Distributed Sparing

## HSV110 Virtualization :

### The Closer to Reality scenario

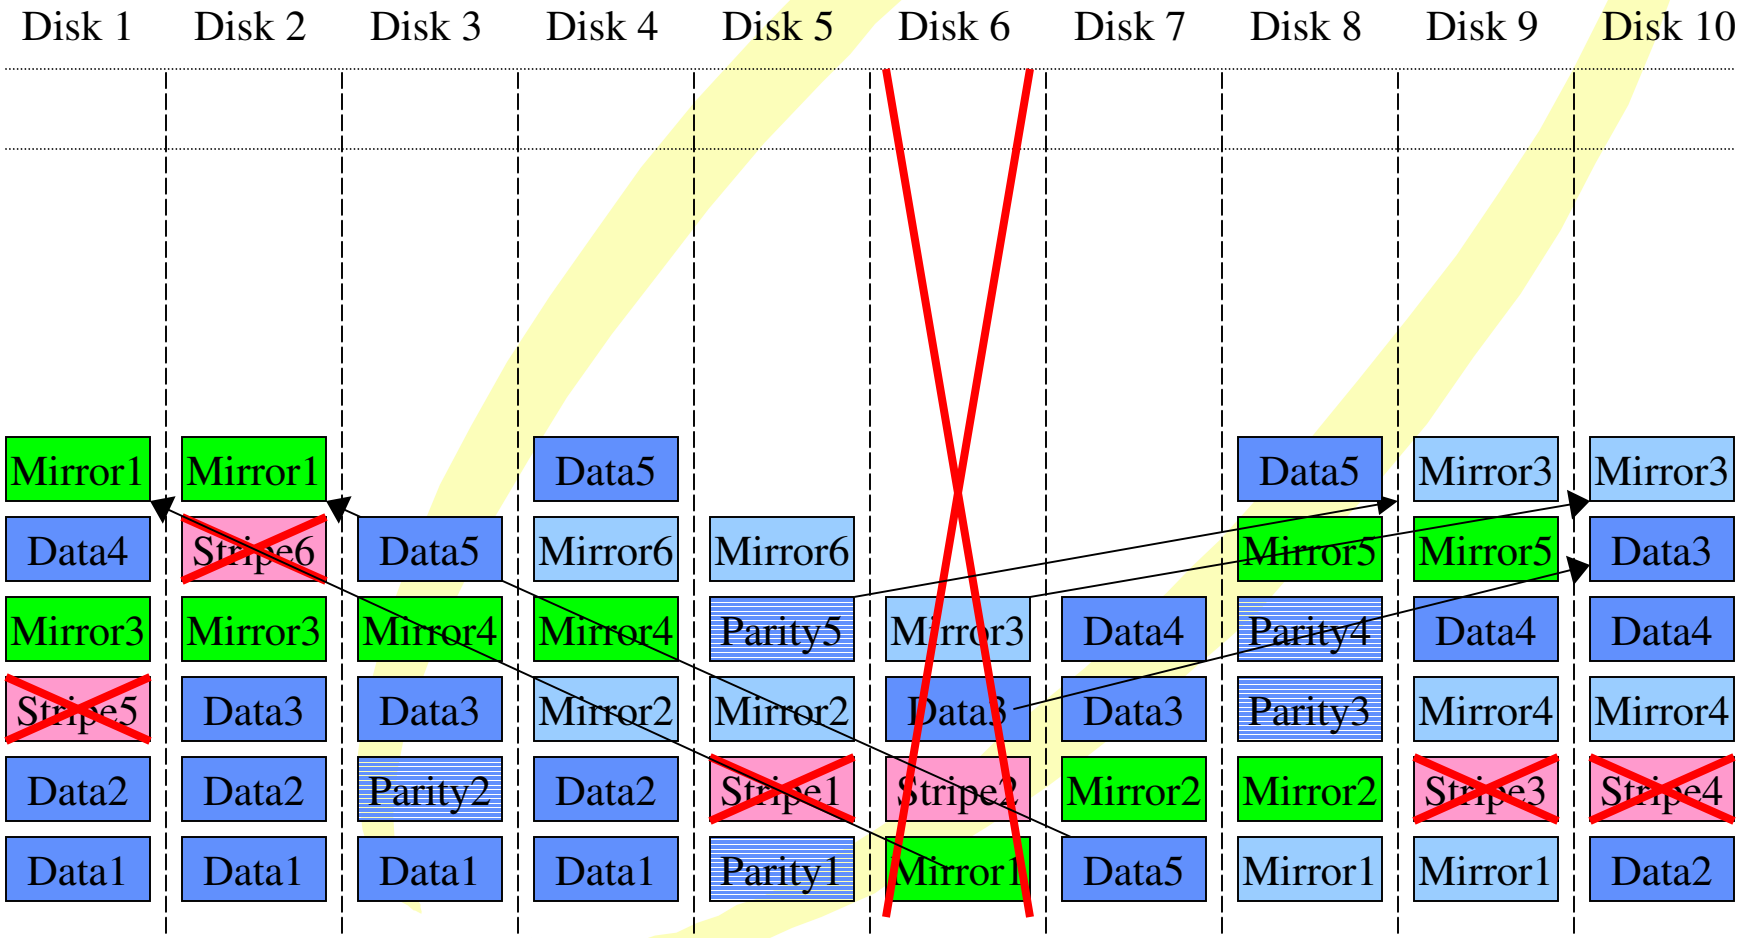

# Redundant Storage Sets (RSS)

- Reduces chance of data loss in large (> 12 physical disks) disk groups
- Not visible to user through Interface
- Complete managed by the HSV controllers
- Typical size for disk group: 8 - 12 physical disks
- If RSS equals 12 physical disks, it splits into 2 RSS of 6 disks each
- Failed disk drive recovery restricted to affected RSS only
  - Reduces access to number of disks, more efficient

- Example: Disk Group with 28 disks
  - # of disks – 6, repeat until the remainder is 12 or less
    - 28-6=22, 22-6=16, 16-6=10 stop
    - RSS configured as 6+6+6+10

# HSV110 Virtualization : Redundant Storage Sets

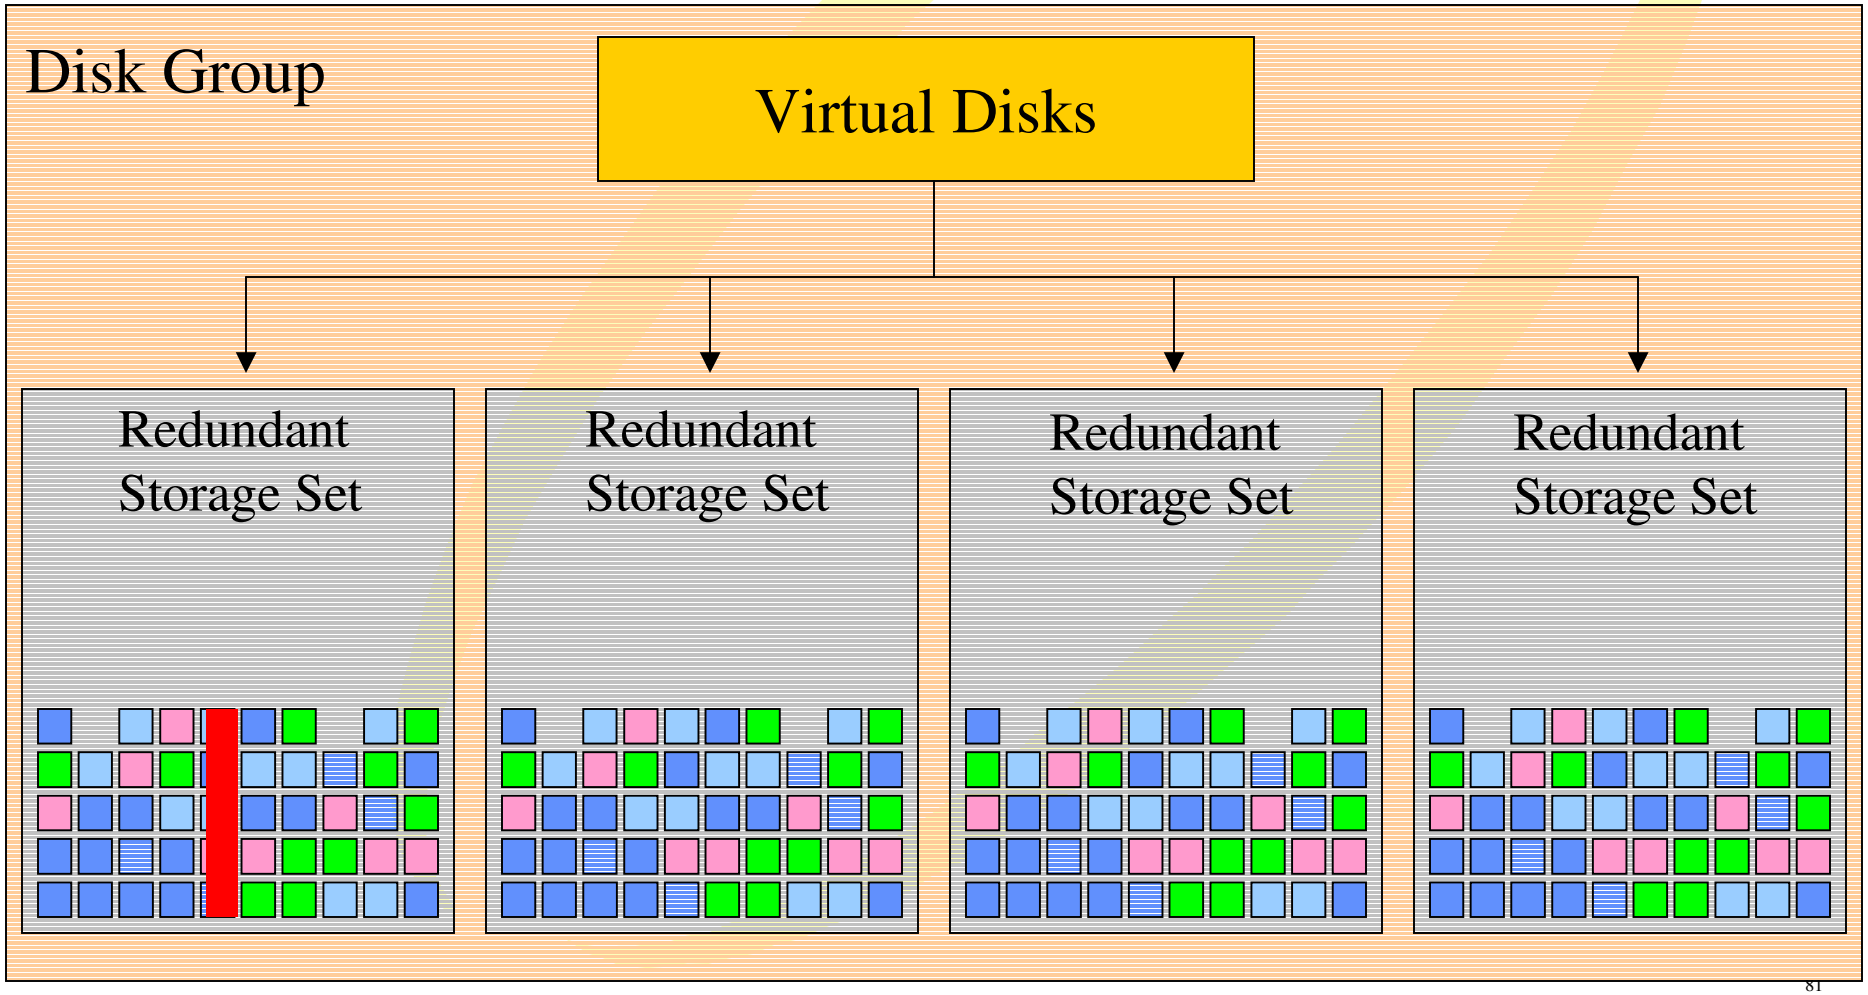

# Food for Thought

- A virtual disk is a disk drive that is imaginary, but operational. It isn't really there but it works just fine
- And then there is a snapshot which is a non-existent, but operational copy of an imaginary (but operational) disk drive
- And then there is the console LUN, which is a pretend virtual disk — an imaginary disk drive that isn't operational but acts as if it might be

# EVA Family - Supported Operating Systems

| | eva3000 | eva5000 |
|---|---|---|
| HP-UX | ✓ | ✓ |
| Tru64 | ✓ | ✓ |
| Windows NT / 2000 | ✓ | ✓ |
| Solaris | ✓ | ✓ |
| AIX | ✓ | ✓ |
| Open VMS | ✓ | ✓ |
| Linux (Intel) | ✓ | ✓ |
| Netware | ✓ | ✓ |
| Windows 2003 (.NET) | ✓ | ✓ |

# EVA OS platforms

- HP/UX
  - v11.0 and v11.i v1 and v2
  - MC Service Guard (v11.13 or later)
  - Secure path v3.0D
- IBM AIX
  - v4.3.3, 5.1, 5.2
  - HACMP 4.4.1
  - Secure Path v2.0D
- SUN Solaris
  - 2.6, v7, v8
  - VERITAS Cluster Server 3.5
  - SUN Cluster 2.2
  - Secure path 3.0C
- Microsoft Windows
  - 2003, 2000 & NT
  - MSCS and Oracle 9iRAC
  - Secure Path 4.0C

# EVA OS platforms (continued)

- Novell NetWare
  - 5.1, 6.0, 6.5
  - Novell Cluster Server
    - V1.01 (5.1), V1.06 (6.0), V1.7 (6.5)
  - Secure Path 3.0C

Novell

- Linux - x86 only
  - Red Hat 7.2, AS 2.1, SLES 7
  - SteelEye Lifekeeper 4.2, 4.3
  - ServiceGuard 11.14.02
  - Secure Path 3.0C

# HP StorageWorks continuous access EVA

- Performs remote mirroring and disaster recovery
- Enterprise-class availability solution
- Up to the last I/O data integrity
- Fast consistent application recovery

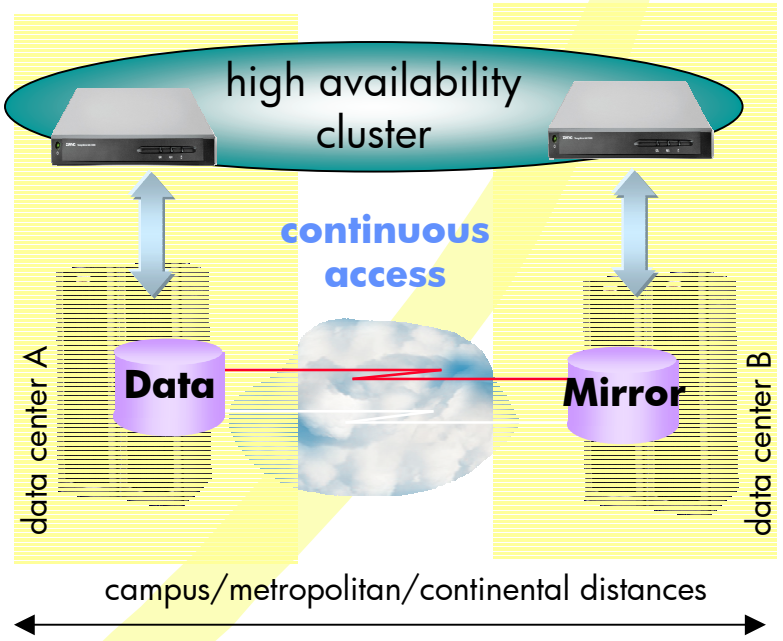

# Continuous Access Highlights

- Management solution automates mirroring functions
- Web based GUI
- Remote access
- Log reports
- Failover/fail back operations
- Engineered with robust, proven resilient technology
- Included with VCS V3, enabled via license key

# Data center – Data center Implementation Example

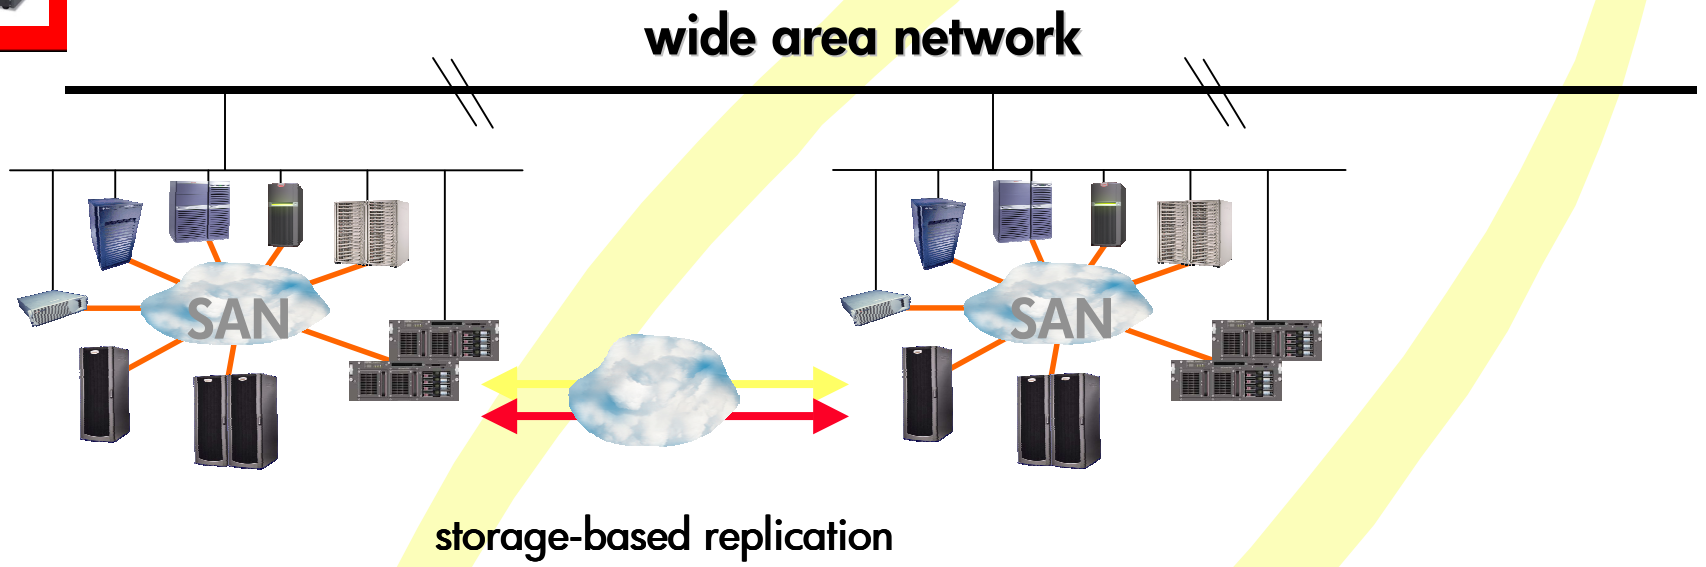


StorageWorks continuous access
data center – data center

# Data center – Data center – Branch Office Implementation Example

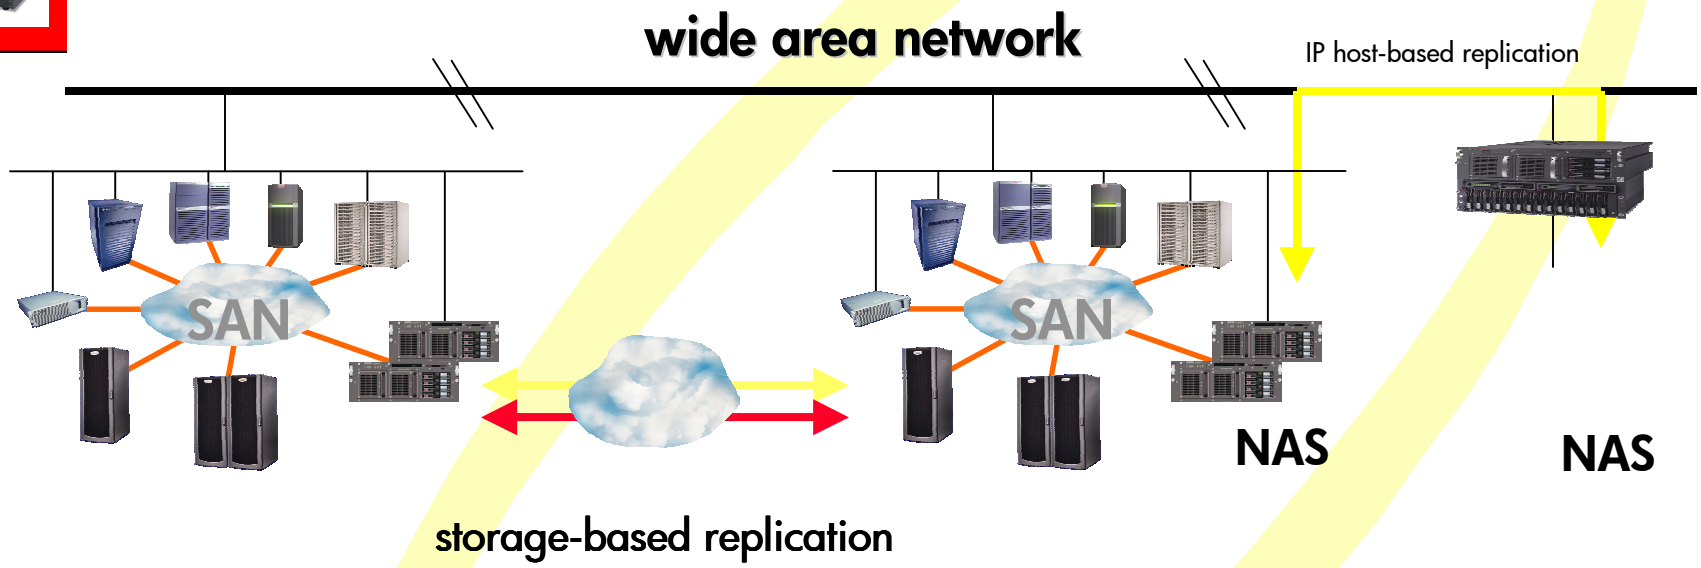


StorageWorks continuous access data center – data center

StorageWorks NAS data copy regional office - data center

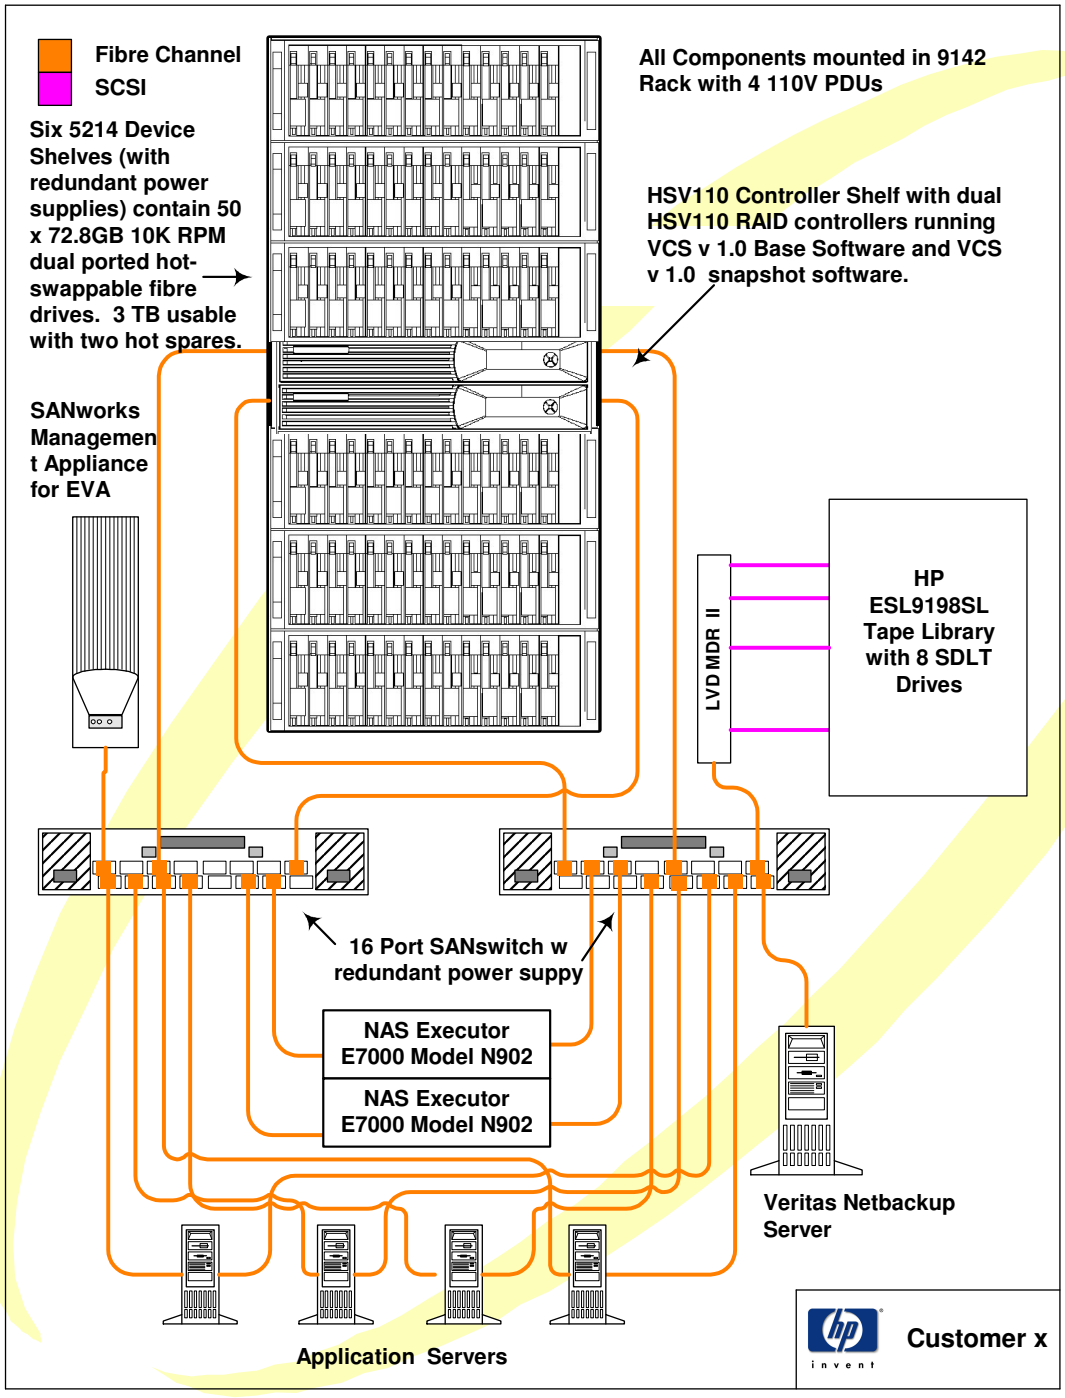
Fibre Channel
SCSI
Six 5214 Device Shelves (with redundant power supplies) contain 50 x 72.8GB 10K RPM dual ported hot-swappable fibre drives. 3 TB usable with two hot spares.
All Components mounted in 9142 Rack with 4 110V PDUs
HSV110 Controller Shelf with dual HSV110 RAID controllers running VCS v 1.0 Base Software and VCS v 1.0 snapshot software.
SANworks Managemen t Appliance for EVA
LVD MDR II
HP ESL9198SL Tape Library with 8 SDLT Drives
16 Port SANswitch w redundant power suppy
NAS Executor E7000 Model N902
NAS Executor E7000 Model N902
Veritas Netbackup Server
Application Servers
Customer x