# OpenVMS SAN/Storage Technical Directions

**Andy Schneider**
OpenVMS Business Management
Hewlett-Packard

# hp OpenVMS Storage Technology

- V7.3-2 Storage Features
- V8.2 Storage Features
- Itanium® Based Systems IO Plans
- Storage Arrays
- Tape Products
- SAN Infrastructure
- Post V8.2 Projects
- Longer Term Storage Interconnects

# Typical FibreChannel Configuration

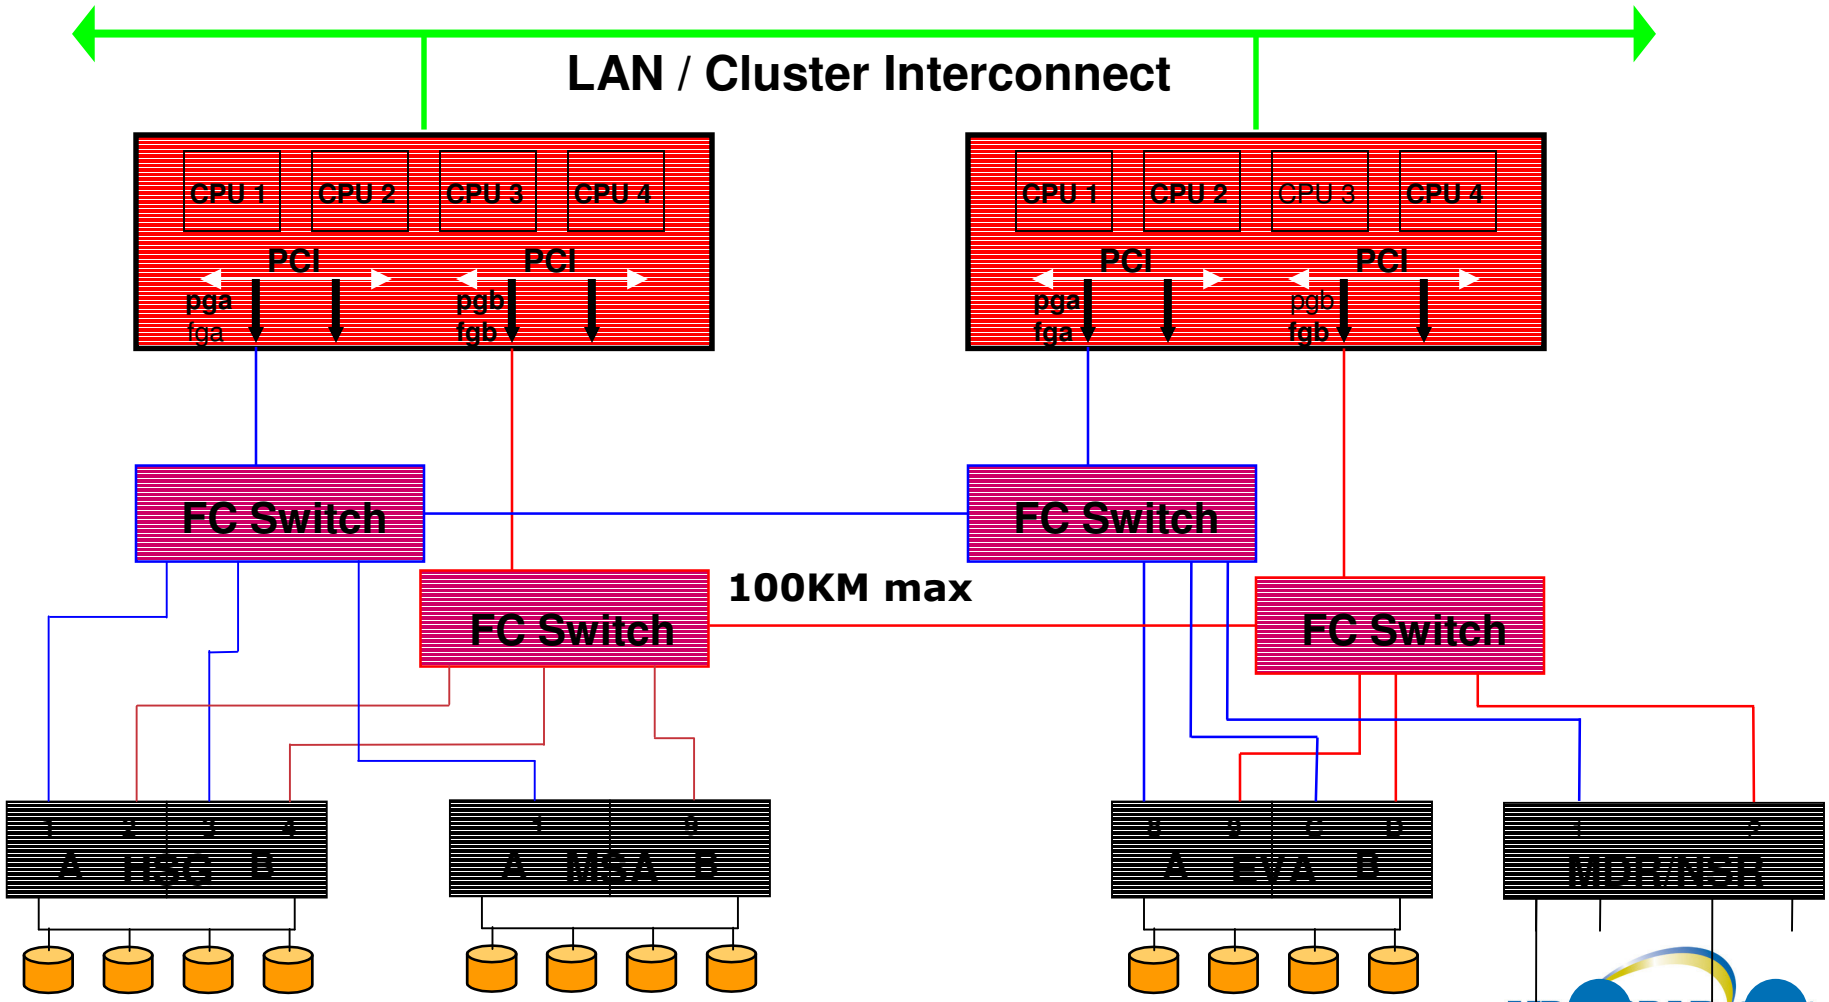

# hp OpenVMS V7.3-2 Storage Features

# hp OpenVMS V7.3-2 Storage Features

- *LP10000 Support*
- *Dynamic Volume Expansion*
- *Shadowing of dis-similar volumes*
- *Mount Verification Filtering*
- *Smart Array Fast Path*
- *Fabric Management Enhancements*

# V7.3-2 Projects

## *LP10000 Support*

- LP10000 is a new dual channel 2gb FibreChannel Adapter (FCA2404DC)
  - Consumes a single PCI slot
  - PCI-X capable
  - 40K IO/sec per port
  - 200MB/sec per port (needs PCI-X slot)
- Should ship in Q3 CY 2004
  - V7.3-2 fibre_scsi V300
  - V7.3-1 fibre_scsi V600
  - V7.2-2 fibre_scsi V600

# V7.3-2 Projects

## *Dynamic Volume Expansion (DVE)*

- Storage arrays today can dynamically expand volumes
- OpenVMS could not utilize the expanded volume without re-initializing
- OpenVMS dynamic volume expansion allows allocation of extra bitmap space at init time and then later enables volume size expansion while the device is mounted
  - $ init /limit $1$dga100: name *! Allocates 1TB bitmap*
  - $ mount/cluster $1$dga100: name *! Mount and use volume*
  - *(expand physical volume size using management tools)*
  - $ set volume $1$dga100: /size=xxxxxx *! Set new size*
- For volumes already initialized
  - $ dismount/cluster $1$dga100: *! Dismount volume*
  - $ mount $1$dga100: name *! Privately mount volume*
  - $ set volume $1$dga100: /limit *! Allocate new bitmap*
  - $ mount/cluster $1$dga100: name *! Cluster mount volume*
  - *(expand physical volume size using management tools)*
  - $ set volume $1$dga100: /size=xxxxxx *! Set new size*

# V7.3-2 Projects

## *Dissimilar Device Shadowing (DDS)*

- Prior to V7.3-2, shadowed volumes had to be identical in size
  - This made it impractical to shadow between different types of storage arrays
- Different sized volumes can now be shadowed
  - The size of the smallest volume determines the size of the shadow set
  - The shadow set can only be mounted on V7.3-2 nodes
  - The size of the shadow set can be grown using DVE

# V7.3-2 Projects

## *Quieter Mount Verification*

- In a SAN there are many reasons for “normal” mount verifications
  - Path switch by another cluster node
  - Dropped FC packets (not a norm, but it does happen)
  - Transient command timeout to a busy storage device
- These result in mount verification messages that alarm system managers
- “Quiet” mount verification will allow infrequent, immediately recovered mount verifications to be suppressed from operator logs
- Sysgen parameters (setting either of these to 0 results in previous behaviour)
  - Mvsupmsg_intvl *(default 1 hour)*
  - Mvsupmsg_num *(default 5)*
- Mount verification messages are suppressed unless that are more than “Mvsupmsg_num” mount verifications within “Mvsupmsg_num” seconds to a single device

# V7.3-2 Projects

## *Smart Array 5300 Fastpath*

- Fastpath is now available for the Smart Array 5300
  - Sysgen parameter fast_path_ports bits *(set to zero to enable fastpath)*
  - Fastpath enabled for all devices by default
    - bit 0 Controls Fast Path for PKQDRIVER (for parallel SCSI).
    - bit 1 Controls Fast Path for FGEDRIVER (for Fibre Channel).
    - bit 2 Controls Fast Path for PKADRIVER (for Ultra3 SCSI).
    - bit 3 Controls Fast Path for PEDRIVER (for LAN).
    - bit 4 Controls Fast Path for PKRDRIVER (for SMART Array 5300).
    - $ set dev pkx /pre=x *! Changes preferred CPU*

# V7.3-2 Projects

## *Fabric Management Enhancements*

- Removed FC port reset on RSCN
  - Prior versions of VMS have reset the FC port (KGPSA adapter), on receipt of fabric RSCN (most common cause would be a re-zone of a switch)
  - This caused all I/O to be aborted and resulted in many device path switches (and mount verification messages)
- PLOGI retry
  - For any device on the SAN that fails PLOGI, VMS will now re-attempt PLOGI every 5 minutes
  - Allows configuration of some class of devices without a re-boot (HSG with full host table, or an XP array which had no VMS LUNs at VMS boot time)

# V7.3-2 Projects

## *Fabric Management Enhancements*

- Fabric name table scan
  - VMS now scans the fabric name table every 5 seconds
  - This allows device configuration of a device in the event that we loose the RSCN message associated with adding a device to the fabric
- Verbose FibreChannel device configuration
  - Boot flags of 10000/30000
  - Define iogen$verbose prior to running “sysman io auto”
- New fibre_scan utility
  - $ mc sys$etc:fibre_scan
  - Scans the entire fabric for configured and non-configured devices
  - Used as a diagnostic tool for SAN configuration problems

# hp OpenVMS V8.2 Storage Features

# hp OpenVMS V8.2 Storage Features

- *Host-Based MiniMerge*
- *SCSI Immediate Error Reporting*
- *Configurable IO Queue Management*

# Host-Based MiniMerge

- CI storage used controller-assisted history log to keep track of writes. No such assists available with FC controllers
- Currently in Beta test of **host-based** MiniMerge solution, based on Write Bit Map technology (from MiniCopy)
- Will support **ALL** FC types (HSG, HSV, MSA, XP)
- Does NOT require any storage firmware assists
- Supported in H2 CY2004 on V7.3-2 via remedial kit
- Supported on V8.2 when it ships in H2 CY2004

# SCSI Immediate Error Reporting

- Some SCSI errors are reported asynchronously via “unit attention” responses on the next SCSI command
- On idle paths errors can be reported days or weeks after they happen
  - This can confuse the user because a mount verification can occur for no apparent reason
  - This can cause field service to swap a part that was previously swapped when the error was reported earlier on a different path
- VMS will now scan all paths once per hour for errors
  - Sysgen scsi_error_poll = 1 (on) by default
  - We’ll backport this new feature to V7.3-2

# Configurable IO Queue Management

- Prior to V7.3-1 VMS implemented a variable IO queue depth per LUN (3-16 based on recent IO sizes)
  - This caused severe performance problems on raid luns for large IO sizes
- In V7.3-1 we moved to a scheme where the host didn't limit IO queue depths until the storage sub-system asked us to back off (via qfull response)
  - Upon receiving qfull we issue no more IO until ½ of the outstanding IO completes (on a given storage sub-system port) and then we again allow the queue size to build until we get another qfull

# Configurable IO Queue Management

- Unfortunately due to the large number of commands that can be inflight in a SAN, we find that the existing algorithm is still too aggressive.
  - Many mount verification messages can result when the same IO gets a qfull response several times in a row
  - Performance suffers badly on the HSG when it has to return qfull responses
  - In extreme cases, the HSG can crash if it receives more IO after signaling qfull

# Configurable IO Queue Management

- In V8.2, we'll move to a default qfull algorithm that drains ½ of existing io requests and then allows the queue depth to increase by 1 enter every 5 seconds.
  - This combined with HSG ACS 8.8 will solve the HSG80 qfull problems
  - There will also be user control of IO queues
    - Hard queue limit per storage sub-system port
    - Variable queue growth limits based on time or load
    - We strongly suggest leaving the default behavior unless the customer can't upgrade to HSG ACS 8.8
- The V8.2 qfull algorithm will be backported to V7.3-2

# Living with HSG qfull problems

# Living with HSG qfull problems

- The HSG80 can queue up to 240 commands per port (4 ports on an HSG80)
  - When the command queue is exceeded you'll see performance problems, excessive mount verification messages and possibly HSG controller crashes
- The EVA can queue 2048 commands / port
- The MSA can queue 512 commands / port

# Checking qfull responses

```
SDA> fc stdt /all

PGA0 SPDT 8157A040 STDTs
------------------------

                                    | STDT  PRLI  Port  Dev   Cred | Act  Cmd  Cnf  Rst  PRLO  Cls | QF    Tgt   Ill   Seq
 STDT      FC-LA  Port Name         | Stat  Stat  I/Os  I/Os  I/Os | Sus  Sus  Pnd  Act  Pnd   Pau | Seen  Rsts  Frms  Tmo
 --------  -----  ------------------ + ----  ----  ----  ----  ---- + ---  ---  ---  ---  ---   --- + ----  ----  ----  ----
 816AC880  0000E  5000.1FE1.0001.84D1 | 0001  0001  0000  0020  0007 | 000  000  000  000  000   000 | 0913  0000  0000  00A9
 816C4840  0000F  5000.1FE1.0001.84D3 | 0001  0001  0000  0009  000B | 000  000  000  000  000   000 | 00FC  0000  0000  0278
 81A6EC00  00025  5000.1FE1.5001.A3E8 | 0001  0001  0000  0000  0000 | 000  000  000  000  000   000 | 0000  0000  0000  0000
 81A6F0C0  00026  5000.1FE1.5001.A3EC | 0001  0001  0000  0000  0000 | 000  000  000  000  000   000 | 0000  0000  0000  0000
 816DF640  00012  5006.0E80.034E.9B00 | 0001  0001  0000  0028  0000 | 000  000  000  000  000   000 | 0000  0000  0000  3D97
 817BCF40  00023  5008.05F3.0001.AF11 | 0001  0001  0000  0000  0000 | 000  000  000  000  000   000 | 0000  0000  0000  0000



PGB0 SPDT 817C6740 STDTs
------------------------

                                    | STDT  PRLI  Port  Dev   Cred | Act  Cmd  Cnf  Rst  PRLO  Cls | QF    Tgt   Ill   Seq
 STDT      FC-LA  Port Name         | Stat  Stat  I/Os  I/Os  I/Os | Sus  Sus  Pnd  Act  Pnd   Pau | Seen  Rsts  Frms  Tmo
 --------  -----  ------------------ + ----  ----  ----  ----  ---- + ---  ---  ---  ---  ---   --- + ----  ----  ----  ----
 818A8F40  00010  5000.1FE1.0001.84D4 | 0001  0001  0000  0012  0001 | 000  000  000  000  000   000 | 00DC  0000  0000  0152
 81864140  0000B  5000.1FE1.0001.84D2 | 0001  0001  0000  001D  000B | 000  000  000  000  000   000 | 0869  0000  0000  0128
 81AB3E00  00020  5000.1FE1.5001.A3E9 | 0001  0001  0000  0000  0000 | 000  000  000  000  000   000 | 0000  0000  0000  0000
 81AB4300  00022  5000.1FE1.5001.A3ED | 0001  0001  0000  0000  0000 | 000  000  000  000  000   000 | 0000  0000  0000  0000
 818D6B80  00012  5006.0E80.034E.9B14 | 0001  0001  0000  002D  0000 | 000  000  000  000  000   000 | 0000  0000  0000  0B96
 8196DA40  0001E  5008.05F3.0001.AF19 | 0001  0001  0000  0000  0000 | 000  000  000  000  000   000 | 0000  0000  0000  0000

SDA>
```

# Living with HSG qfull problems

- Balance IO across all HSG ports
  - set dev $1$dgaxxx /switch/path=FGA0.5000-1FE1-5001-A3E1
- Reduce IO load by migration hot luns to other storage arrays if possible
- Reduce DIOLM on accounts running hot IO jobs
- Backup is often the culprit behind HSG qfull problems
  - Backup will attempt 100s of concurrent IOs to try to keep ahead of streaming tape devices
  - Backup one lun at a time per HSG
  - Run backup in a captive account with diolm of 32-64 (testing shows that performance doesn't suffer)

# Integrity Server Systems IO Plans

# OpenVMS Integrity Server Systems IO Plans

- SCSI
  - A6829A U160 controller (lvd only)
  - U320 controller (lvd only)
  - No plans for multi-host SCSI
  - MSA30 storage shelf
- Host Based Raid
  - Smart array family (6400)
- FibreChannel
  - 2gb 2-channel adapter
    - A6826A (qlogic)
    - Storage (V8.2) & Cluster Interconnect (V8.3)
    - Full boot support
  - Storage Arrays
    - HSG / EVA / MSA / XP
  - Tape Libraries
    - ESL / MSL
      - Connect via MDR or NSR variants

# Storage Arrays

# Storage Products
# EVA 5000

- 4 2gb host ports
- 4 2gb FC disk ports (2 redundant loops)
- 240 disks / 35TB max
- 1GB mirrored cache
- Multi-level snapshots
- Dynamic LUN expansion
- CA/DRM for OpenVMS
- SSSU host based control utility

# Storage Products
# EVA 3000

- 4 2gb host ports
- 2 2gb FC disk ports(1 redundant loop)
- 56 disks / 8TB max
- 1GB mirrored cache
- Multi-level snapshots
- Dynamic LUN expansion
- CA/DRM for OpenVMS
- SSSU host based control utility

# HP StorageWorks xp128

| xp128 | |
|---|---|
| Max Disk Drives | 128 |
| Max Capacity | 18 TB raw<br>16 TB usable |
| Max Aggregated Crossbar Throughput | 7.5 GB/s total<br>5 GB/s data<br>2.5 GB/s control |
| Max Sequential Data Transfer Rate | 1.1 GB/s sustained<br>2.4 GB/s peak from cache |
| Max Random IOPS from Cache | 375,000 |
| Max LDEVs | 8192 |

Enterprise Performance, Capacity, and Reliability in a Small Footprint

# HP StorageWorks xp128

Start at the size you need, scale as your needs grow, staying in the same small footprint

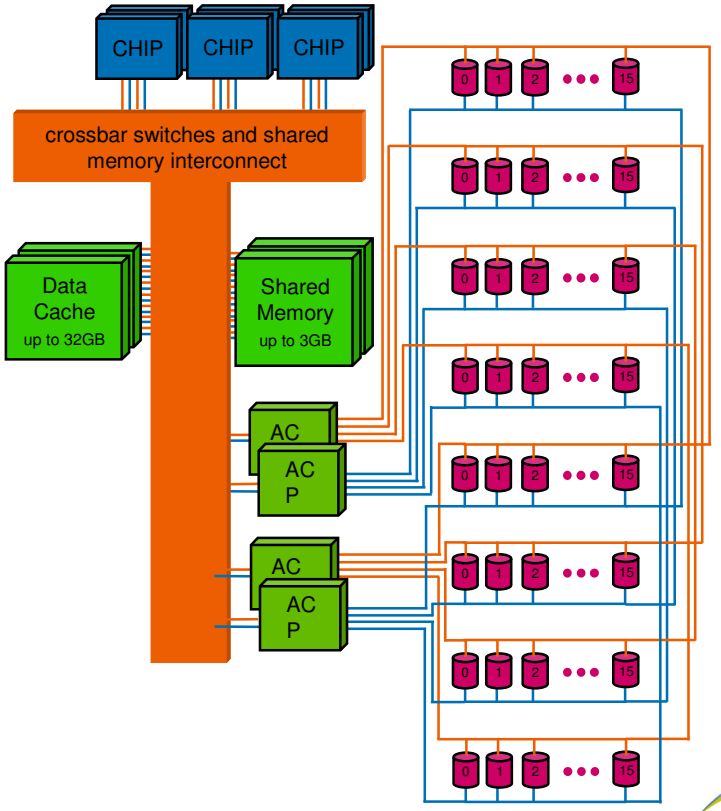

# HP StorageWorks xp1024

| xp1024 | |
|---|---|
| Max Disk Drives | 1024 |
| Max Capacity | 149 TB raw<br>129 TB usable |
| Max Aggregated Crossbar Throughput | 15 GB/s total<br>10 GB/s data<br>5 GB/s control |
| Max Sequential Data Transfer Rate | 2 GB/s sustained<br>3.2 GB/s peak from cache |
| Max Random IOPS from Cache | 500,000 |
| Maximum LDEVs | 8192 |

Unprecedented Performance, Scalability, and Reliability for the Enterprise

HP WORLD 2004
Solutions and Technology Conference & Expo

# HP StorageWorks xp1024

Start at the size you need, scale to the highest usable capacity of any array on the planet

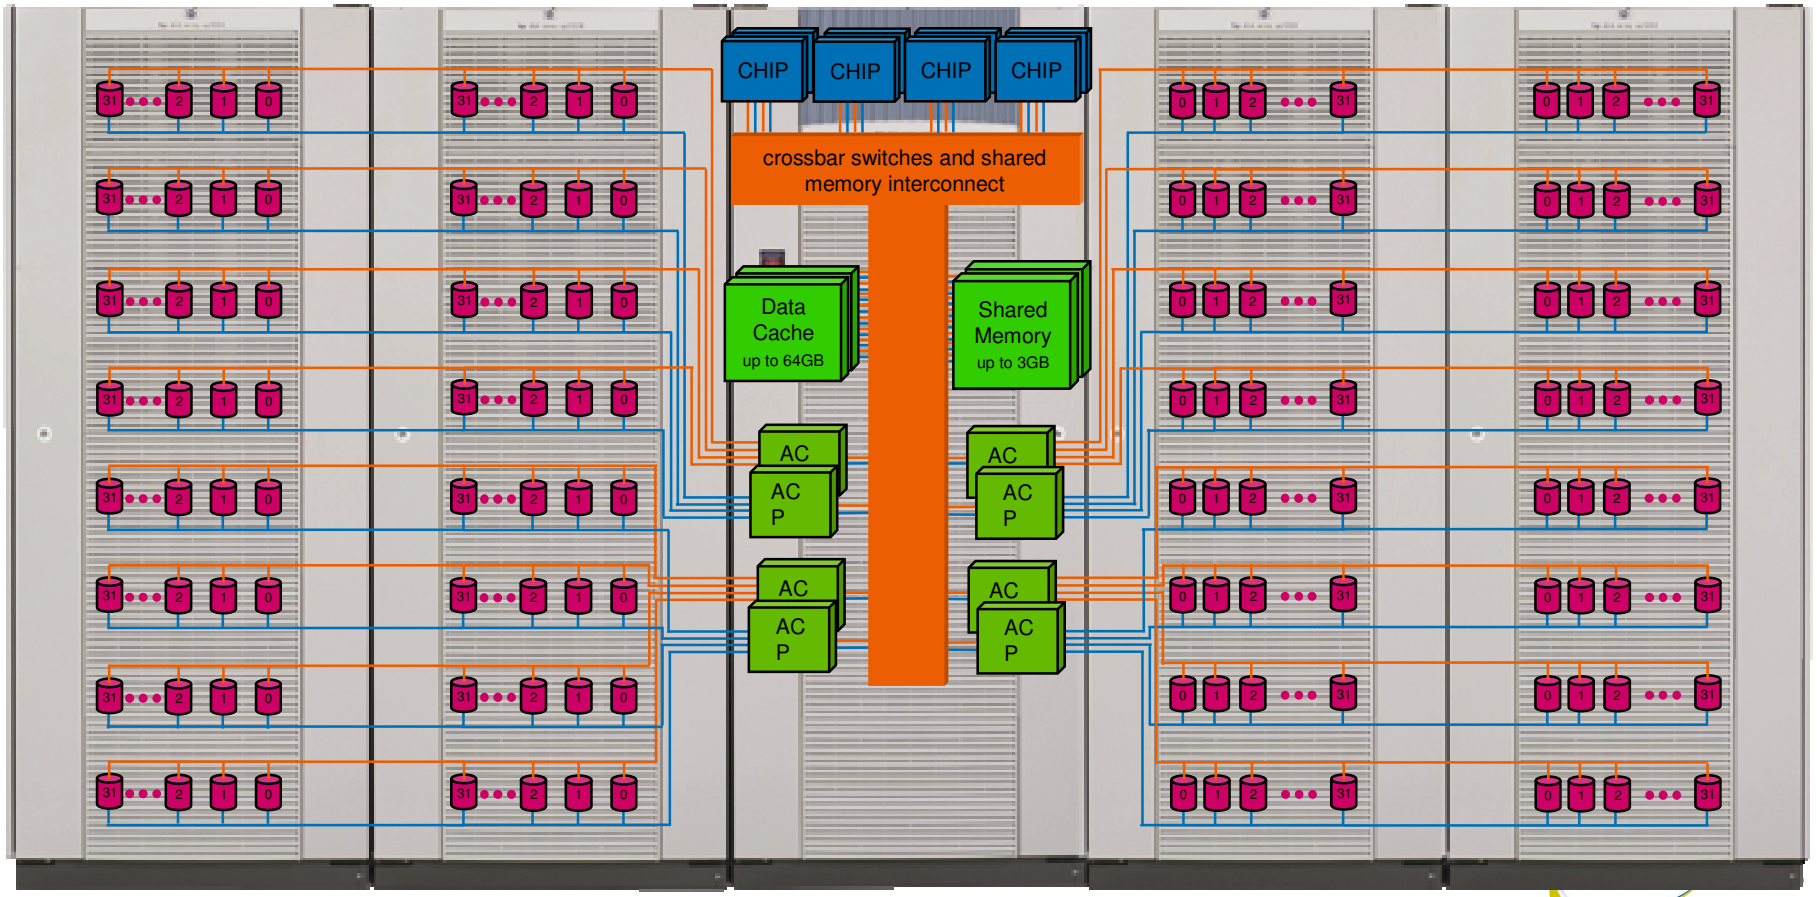

# Storage Products

## *Modular Storage Array 1000*

- 2gb FibreChannel front-end
- 4 U160 SCSI backend ports
- 4u rackmount with 14 drives
- 28 additional drives with
  2 external storage shelves
- Works in existing SANs
- Low cost 2 node clusters
  *with embedded 3 port FC-AL hub*
  *(V7.3-2 only, Release TBD)*
- Supported with
  - V7.2-2
  - V7.3
  - V7.3-1
  - V7.3-2

# Storage Products

## *Modular Storage Array 1000*

- V4.24 firmware supports all OS
- Minimum Tima kits
  - OpenVMS V7.2-2 VMS722_FIBRE_SCSI-V0400
  - OpenVMS V7.3 VMS73_FIBRE_SCSI-V0500
  - OpenVMS V7.3-1 VMS731_MSA1000-V0100
- FC-AL release in future

# MSA1000 SAN Solution

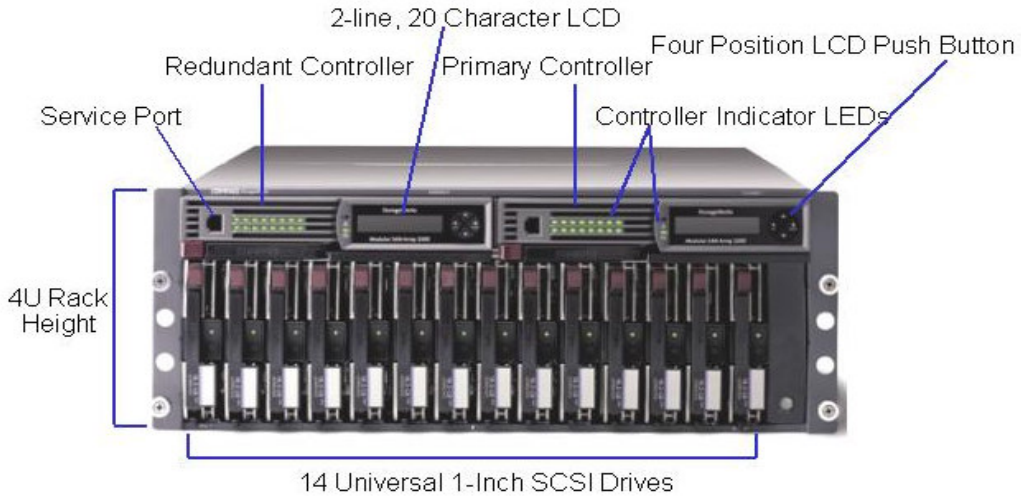


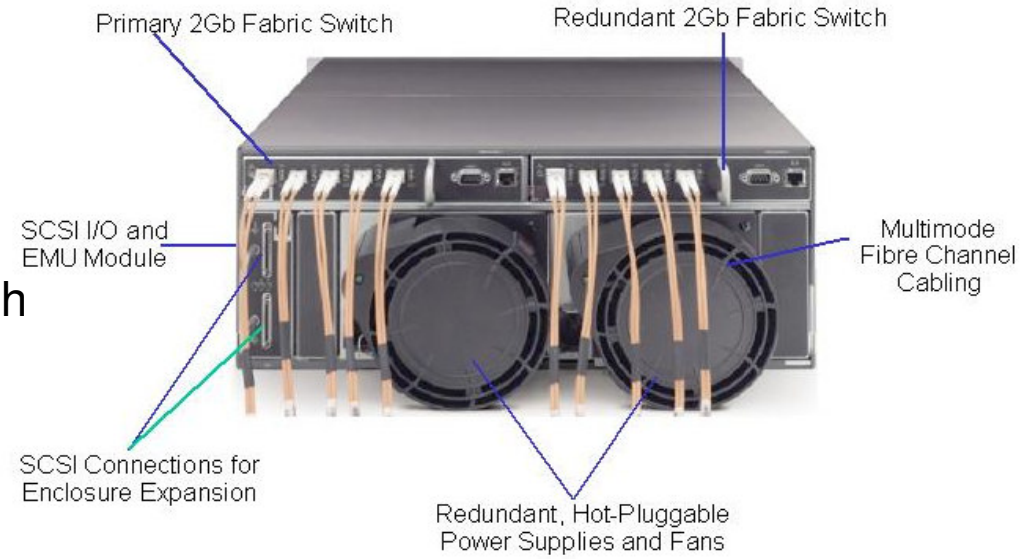


**FC Connection Options**

Single 2Gb port for external switch

3 port 2Gb Hub (Future for VMS)

8 port 2Gb switch (Brocade)

# Storage Products

## *Modular Storage Array 1000*

- High Performance Architecture
  - 200MB/sec throughput
  - 25,000 I/O per second
- Redundant controller support
  - Active/Standby in initial product
  - Active/Active in future
- RAID 0, 1, 0+1, 5 and ADG
- LUN Masking (SSP)
- 2Gb/1Gb auto-sense host ports
- Dual Cache Modules
  - Upgradeable to 512MB (per controlle
- Serial line config and management

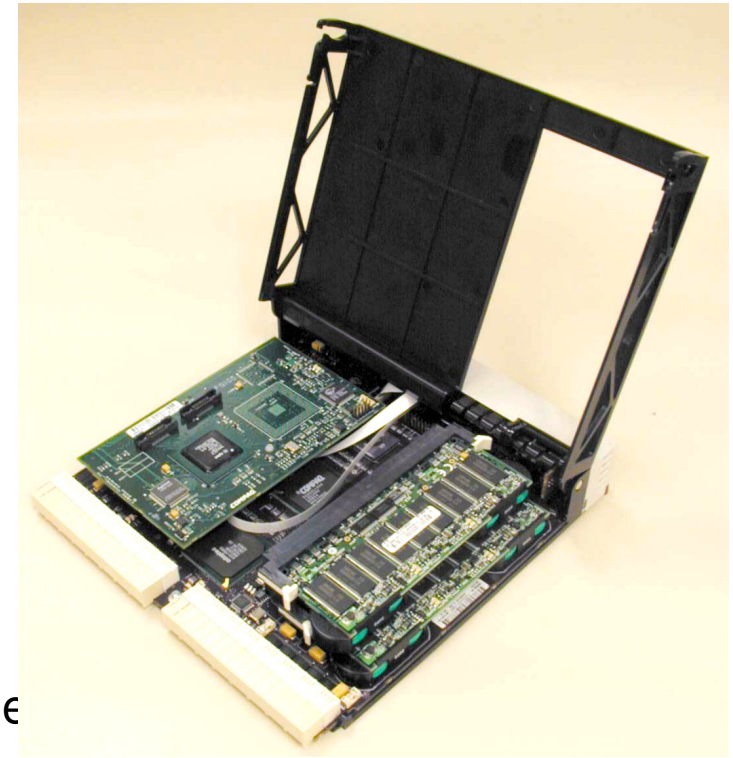

# Low End FC-AL Based Cluster

(Not Yet Available)

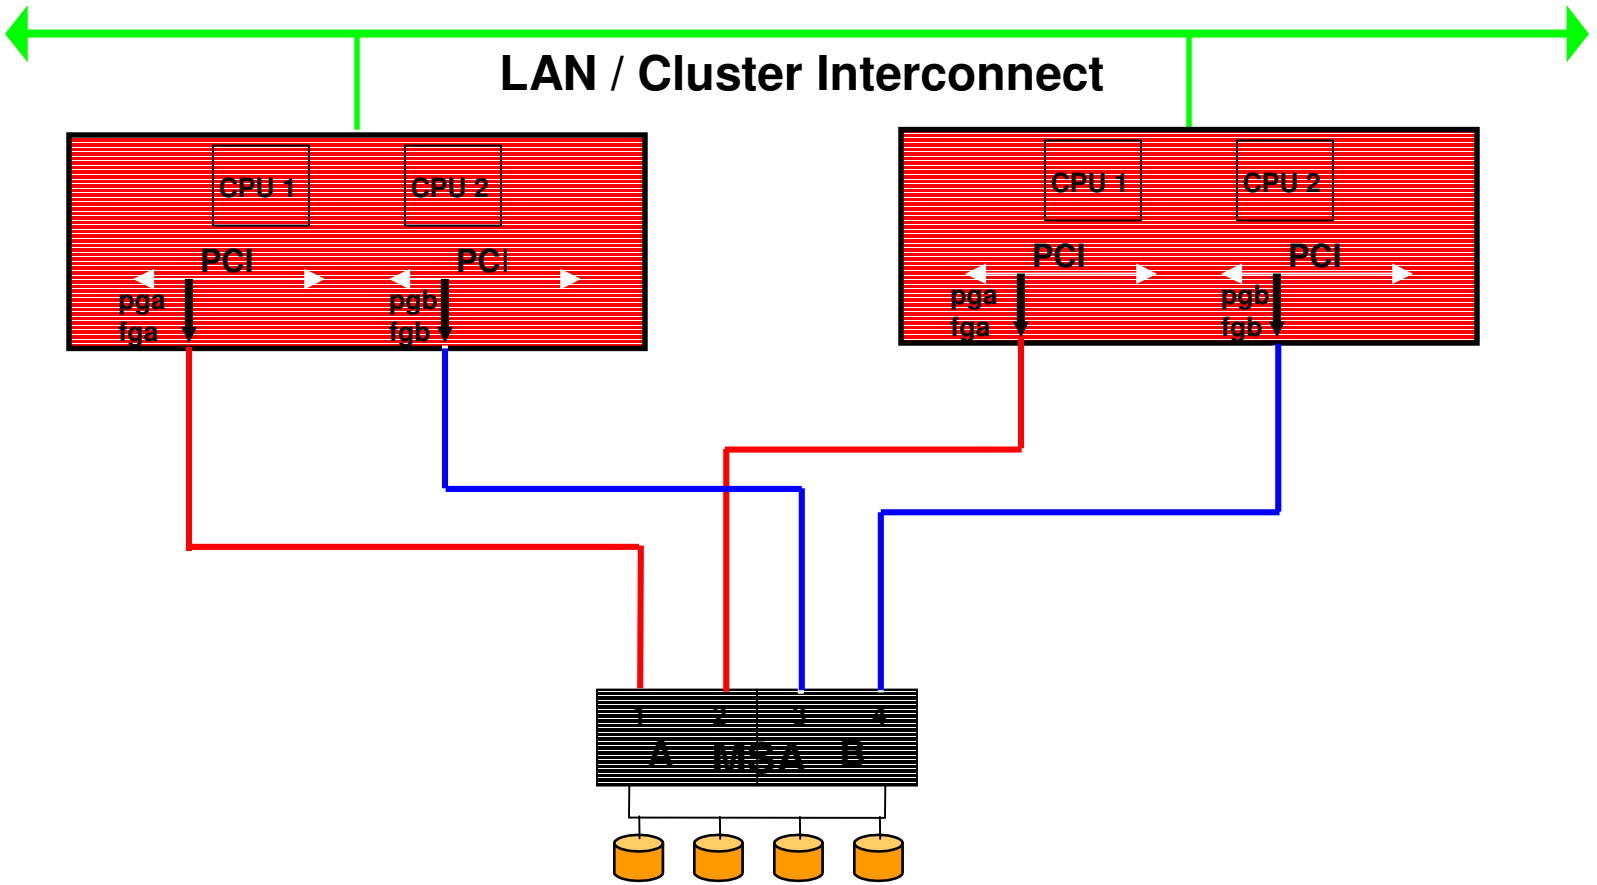

# Storage Products

## *Smart Array 5300*

- Smart Array 5300
  - 2/4 U160 SCSI Channels
  - Up to 56 drives (4TB)
  - Raid0/1/5/ADG
  - Up to 256MB cache
- Supported on V7.3-1
  - 300MB/sec
  - 20K io/sec
- Fastpath with V7.3-2
- **Shadowing:**

  *Doesn't support forced error commands so full shadowing support is an issue. Shadowing works fine but a member will be ejected if an unrecoverable disk error occurs on one member and an error cannot be forced on the shadow copy*

- Smart Array 5400 (U320), due in late 2004

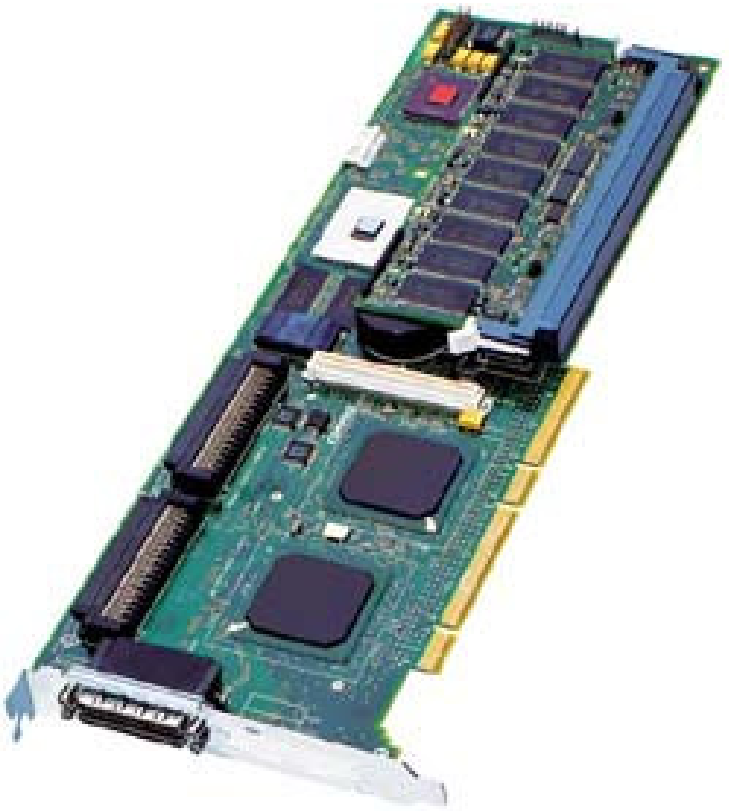

# HSG / MSA / EVA Volume Construction

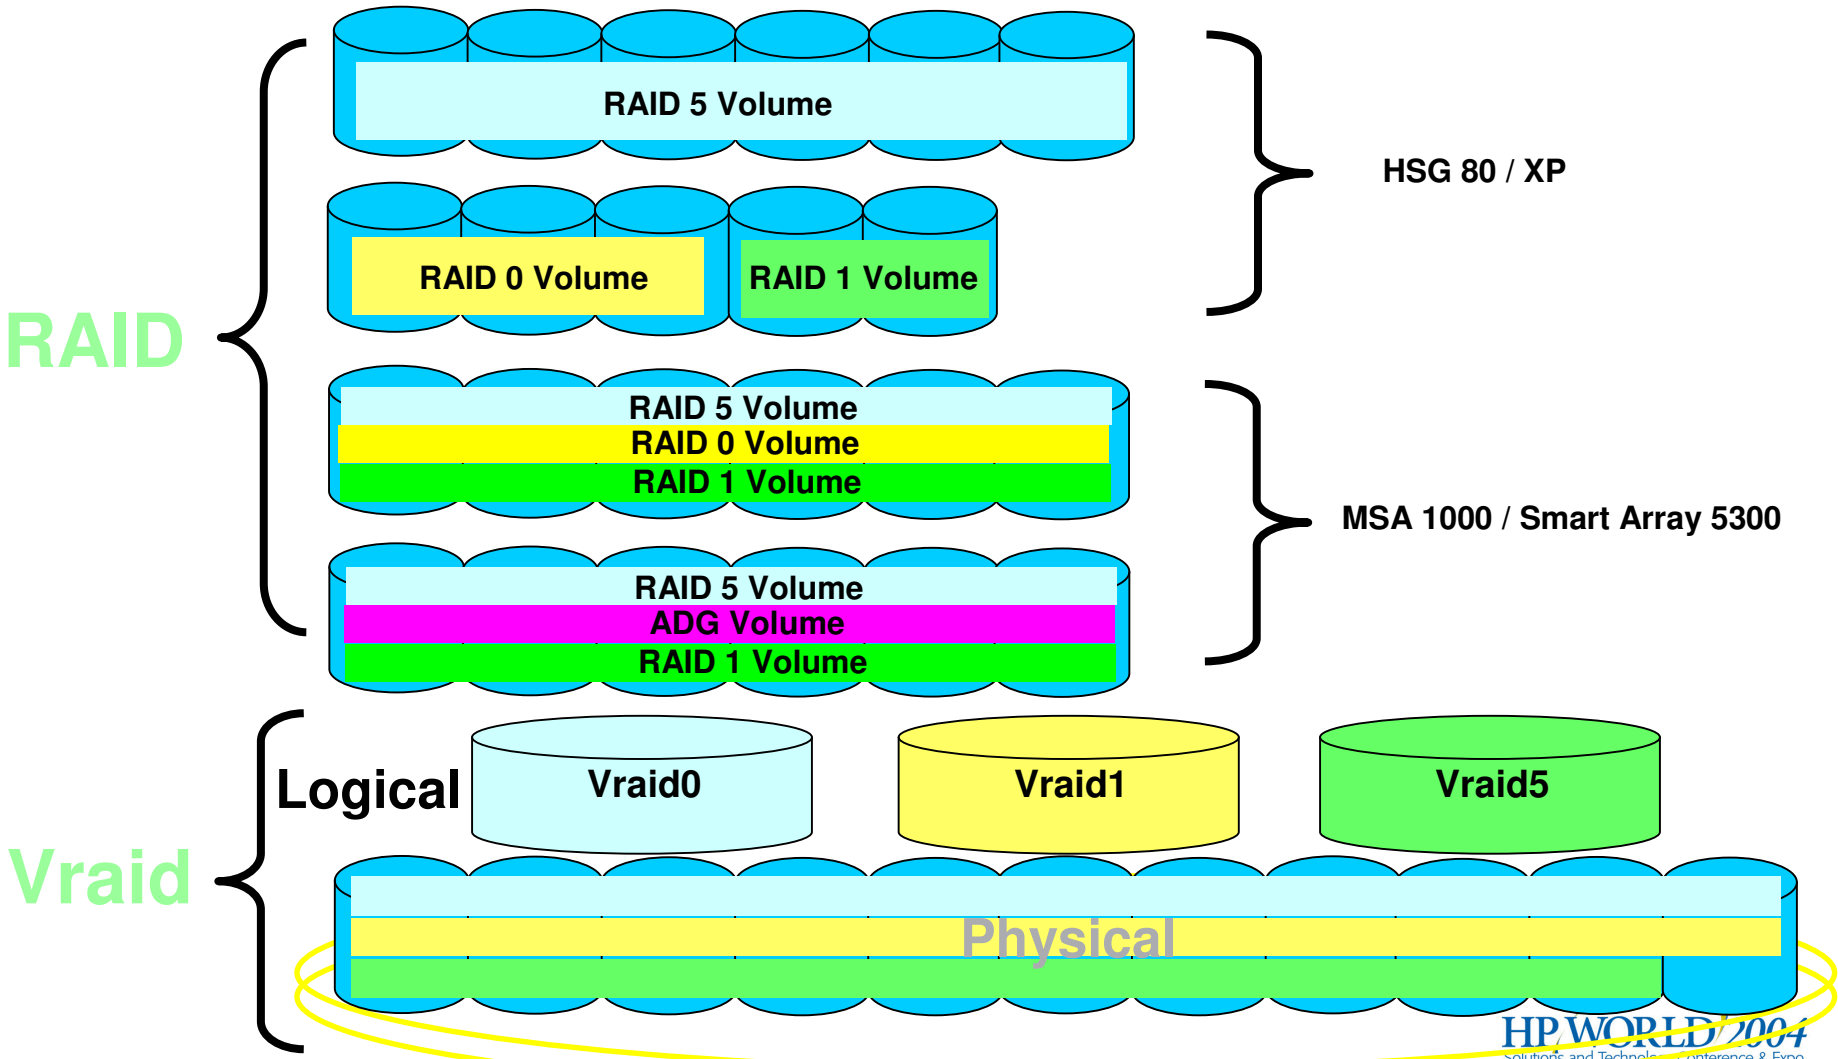

# Storage Positioning

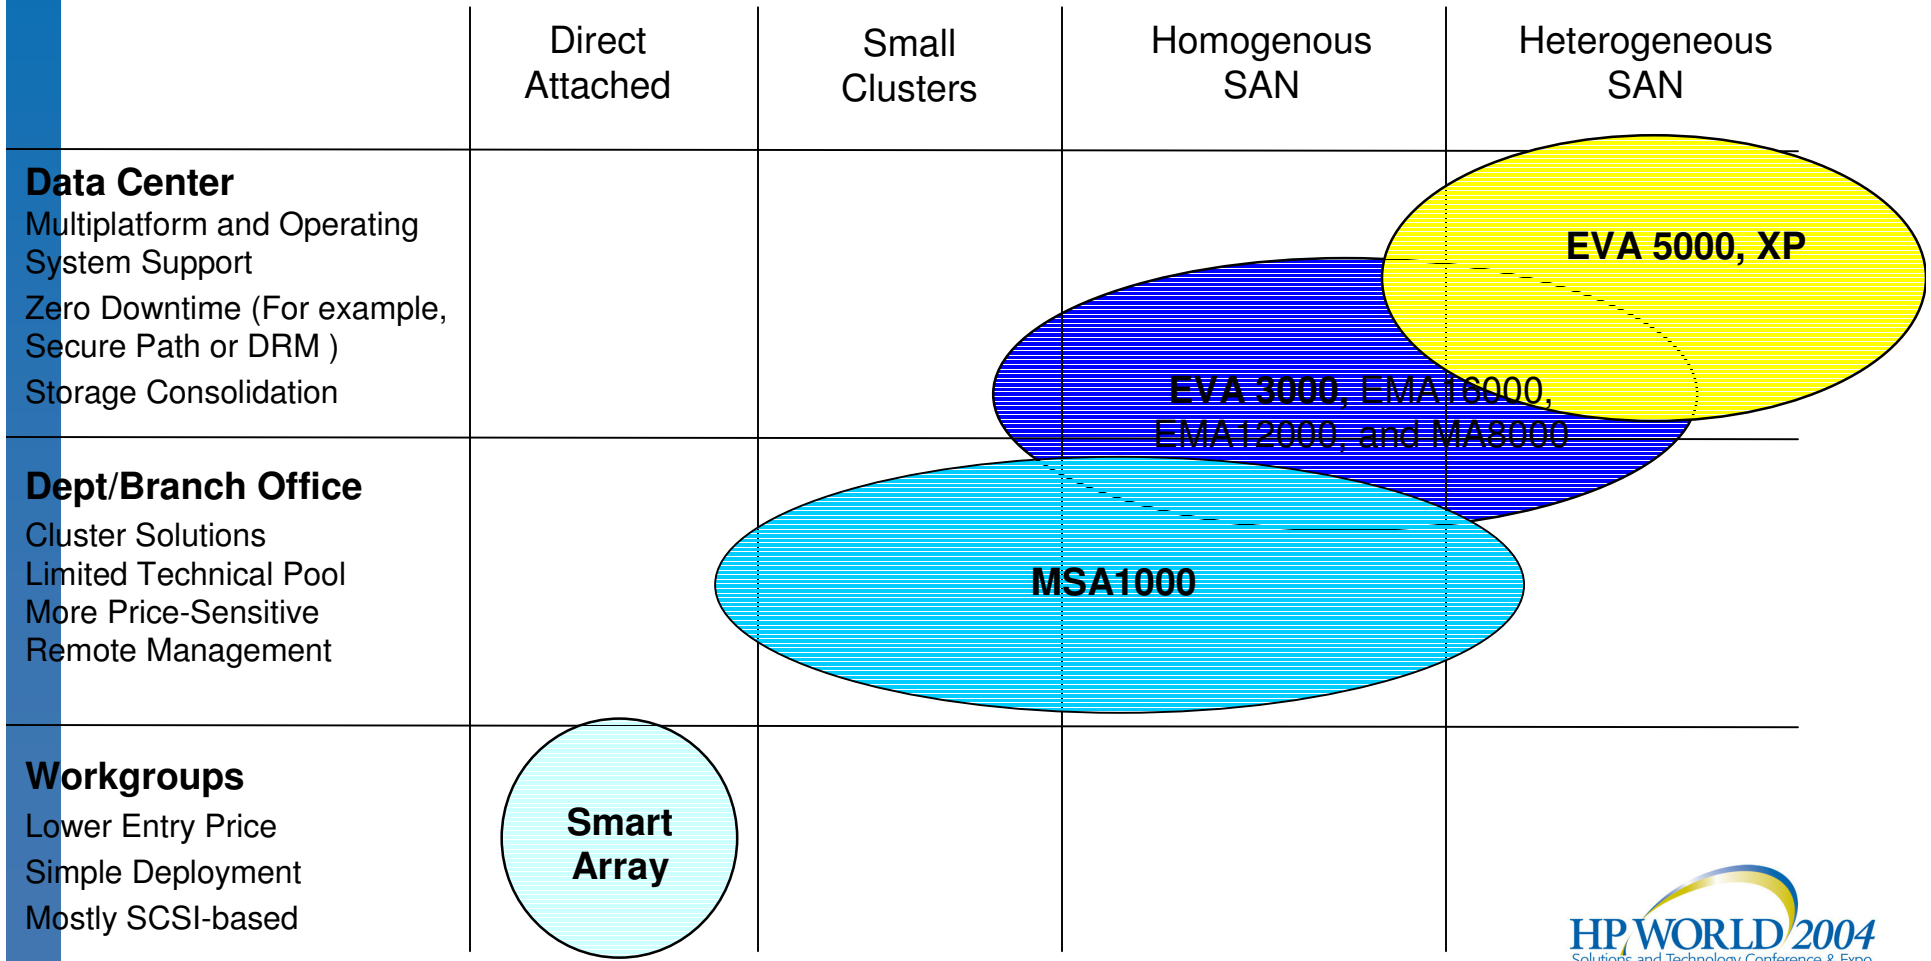

# StorageWorks Array Specs

| Array | MSA1000 | VA7110 | VA7410 | MA / EMA | EVA3000 | EVA5000 | XP 128 | XP 1024 |
|---|---|---|---|---|---|---|---|---|
| controllers | 2 | 2 | 2 | 2 | 2 | 2 | 6 | 8 |
| host interface | FC | FC | FC | FC | FC | FC | FC | FC |
| host ports | 1 to 16 | 2 | 4 | 4 | 4 | 4 | 24 | 32 |
| drive interface | SCSI | FC | FC | SCSI | FC | FC | FC | FC |
| drive enclosure density | 14 | 15 | 15 | 14 | 14 | 14 | 32 | 32 |
| max. raw capacity | 6 TB | 6.5 TB | 15 TB | 6.1 TB/ma8k 18.3TB/ema12k | 8TB | 35TB | 18 TB | 149 TB |
| disk sizes | All Ultra3 / Ultra 320 drives (18 to 146 GB) | 146GB/10k 73GB/15k,10k 36GB/15k,10k | 146GB/10k 73GB/15k,10k 36GB/15k,10k | 146GB/10k 73GB/15K,10k 36GB/15k,10k | 146GB/10k 73GB/10k 36GB/15k,10k | 146GB/10k 73GB/10k 36GB/15k,10k | 146GB/10k 73GB/10k 36GB/15k,10k | 146GB/10k 73GB/10k 36GB/15k,10k |
| max. no. of drives | 42 | 45 | 105 | 42/ma8k 126/ema12k | 56 | 240 | 128 | 1024 |
| peak iops from cache | 30 | 16,5k | 31k | 21k | 105k | 168k | 375k | 500k |
| peak iops from disk | ND | ND | 10k | ND | 13k | 55k | 27k | 57k |
| peak sustained throughput | 202 MB/s | 160 MB/s | 330 MB/s | ND | 337 MB/s | 525 MB/s * | 2400 MB/s (front-end) | 3200 MB/s (front-end) |
| cache capacity | 512MB | 2 GB | 2 GB | 512MB | 2GB | 2GB | 32 GB | 64 GB |
| snap/clone | Virt. Replicator | Business Copy | Business Copy | Snapshot | BC | BC | BC | BC |
| RAID configurations supported | RAID 0, 1, 3/5, 0+1, partitioning | RAID10, 5DP, AutoRaid | RAID10, 5DP, AutoRaid | RAID 0, 1, 3/5, 0+1, partitioning | Vraid0, Vraid1, Vraid5 | Vraid0, Vraid1, Vraid5 | RAID1 (2+2), 1 (4+4), 5 (3+1), 5 (7+1) | RAID1 (2+2), 1 (4+4), 5 (3+1), 5 (7+1) |
| replication | | CASA | CASA | DRM (sync + async) | CA (3Q2003) | CA (2Q2003) | CA (sync + async) | CA (sync + async) |

# Storage Arrays Roadmap


MSA1000

**Entry-Level:** Industry's lowest cost SAN storage


eva3000

**Mid-Range:** Enterprise-class for the mid-range


eva5000

**Enterprise-modular:** Enterprise-class that's powerfully simple


XP

**Enterprise-frame:** Exceptional availability, scalability, and throughput

Open VMS & TRU 64

Integrity Server dual-path support
-HP-UX 11.23
-Windows Server 2003

Blade server support

300 GB disk drive
Low cost drive (SATA)

HP-UX

iSCSI

Choice of rack

Integrity Server dual-path support
-HP-UX 11.23
-Windows Server 2003

Sync & Async Continuous Access:
- More copy sets
- Fan in/out
- Improved management

24x7 Foundations Services
Phone home support
Critical Services support
Pro-active support services
SMI API for management

300 GB disk drive

Low cost drive solution

Integrity Server dual-path support
-HP-UX 11.23
-Windows Server 2003

Enhanced Continuous Access
HP-UX MetroCluster
Host-based management
Direct connect to host
Increased heterogeneity

Integrity Server dual-path support
-HP-UX 11.23
-Windows Server 2003

Industry standard API for management

300 GB disk drive

iSCSI

Performance enhancements

Expanded cache-memory capacity

today | 2nd half - CY2003 | CY2004

HP WORLD 2004
Solutions and Technology Conference & Expo

# Tape Products

# Tape Libraries

- ESL 712e / 630e
  - Up to 24 SDLT / LTO drives
  - 712 LTO cartridges / 630 SDLT cartridges
  - Embedded FC ports (1:2 drives)
  - Dedicated FC port for robot
  - VMS testing underway now
- ESL 9595/9322/9326
  - Up to 16 SDLT / LTO drives
  - Up to 4 2gb FC ports
  - Up to 595 cartridges
    - 95TB of uncompressed data
    - 900GB/hour uncompressed backup rate SDLT 160/320
    - 1.7TB/hour uncompressed backup rate LTO-2
- MSL 5052
  - 4 SDLT drives / 52 cartridges
- MSL 6060
  - 4 LTO drives / 60 cartridges
- MSL 5026
  - 2 SDLT drives / 26 cartridges
- MSL 6030
  - 2 LTO drives / 30 cartridges

# Network Storage Router M2402

- 1U Fibre Channel-to-SCSI router
  - 2Gb FC Support
  - 4 module slots
    - 2x 2gb Fibre Channel
    - 4x LVD/SE SCSI
    - 4x HVD SCSI
  - Web Based Management
  - Embedded product for tape libraries (E2400-160)


Power Side


Port Side

# Network Storage Router N1200

- 1U Fibre Channel-to-SCSI router
  - High Performance - 2Gb FC Support
    - 200 MB/s of information throughput
  - 1 2Gb Fibre Channel port
  - 2 U160 LVD SCSI ports
  - Web Based Management
  - Embedded product for tape libraries (E1200-160)

# LTO Tape Drives

- VMS will never support Ultrium 1 drives
  - No support of odd byte transfers
- VMS supports the Ultrium 2 (LTO460) drives
  - Testing completed on all current AlphaServer systems
  - Supported in both direct-attach SCSI and behind the FC Bridges (NSR, MDR)
  - Now supported in **ESL/MSL libraries**

# Recent Backup Performance Measurements

| Drive | No compression | Highly compressed |
|---|---|---|
| SDLT 110/220 | 12MB/sec | 49 MB/sec |
| STK 9940A | 10MB/sec | 33 MB/sec |
| SDLT 160/320 | 18MB/sec | 49 MB/sec |
| STK 9940B | 18 MB/sec | 55 MB/sec |
| Ultrium 2 | 32MB/sec | 60MB/sec |

# Tape and Optical roadmap

**Tape Drives**

Available with Ultrium LTO, SDLT, DLT and DAT technologies

**Autoloaders**

Entry level

Available with LTO, SDLT, DLT1

**Tape Libraries**

MSL series

ESL series

Mid-range:

Enterprise features for the mid-range

Available with LTO, SDLT

Enterprise

Mission-critical: redundancy, high-availability

**Optical**

Exceptional

Long-term archiving solutions

5th generation DDS – DAT72 and DAT72h (hot plug) drives

Ultrium 3

DDS 6

SDLT 600

Ultrium Tape WORM

Filling the niche for entry level libraries support on both Windows and HP-UX environments with full ProLiant support

SDLT 600 Autoloader

DAT 72 Autoloader

Integrity Server support:
HP-UX
Windows
Linux
Open VMS

ESL9000 series
Extended tape library architecture (ETLA)
Embedded interface & controller
Tape mirroring
Path fail-over
Secure-path

MSL5000/6000 series
ETLA
Full mixed LTO/SDLT media support

ESL e series

ESL/MSL series with SDLT600

Enterprise library with Ultrium 460 native fibre drives

Ultra Density Optical (UDO) drive
•integration into optical libraries

Today

2nd half - CY2003

CY2004+

# SAN Infrastructure

# SAN Infrastructure product portfolio

## Infrastructure alternatives that span the enterprise

Business enterprise


core 2/64


director 2/64 director 2/140


MDS-9509

Business mid-range


switch 2/32


edge 2/32


MDS-9216

Business entry


switch 2/16


switch 2/8 EL


blade 2/8


edge 2/24


MDS-9120 & 9140

B-Series M-Series C-Series

# New SAN Features

## *Bigger SAN Switches*

- SAN Switch 2/32
  - 32 port monolithic switch
    - 4 port 8Gb/sec ISL trunking available
- Core Switch 2/64
  - 2 x 64 port blade type switch
    - 16 ports per blade
    - Total of 128 2gb ports
    - 4 port 8Gb/sec ISL trunking available
- SAN Director 2/128
  - 128 port blade type switch
    - Non-blocking
    - 16 ports per blade
    - 4 port 8Gb/sec ISL trunking available

# New SAN Features

## *Cisco Switches*

- OpenVMS now qualified with the full line of Cisco Fibre switches
  - 9120/9140
    - 20/40 port edge switches
  - 9216
    - 16 ports + 1 plug in (16/32 ports)
  - 9506/9509
    - Director class switches with 4/7 plug in modules
    - Up to 224 ports


# SAN Infrastructure roadmap

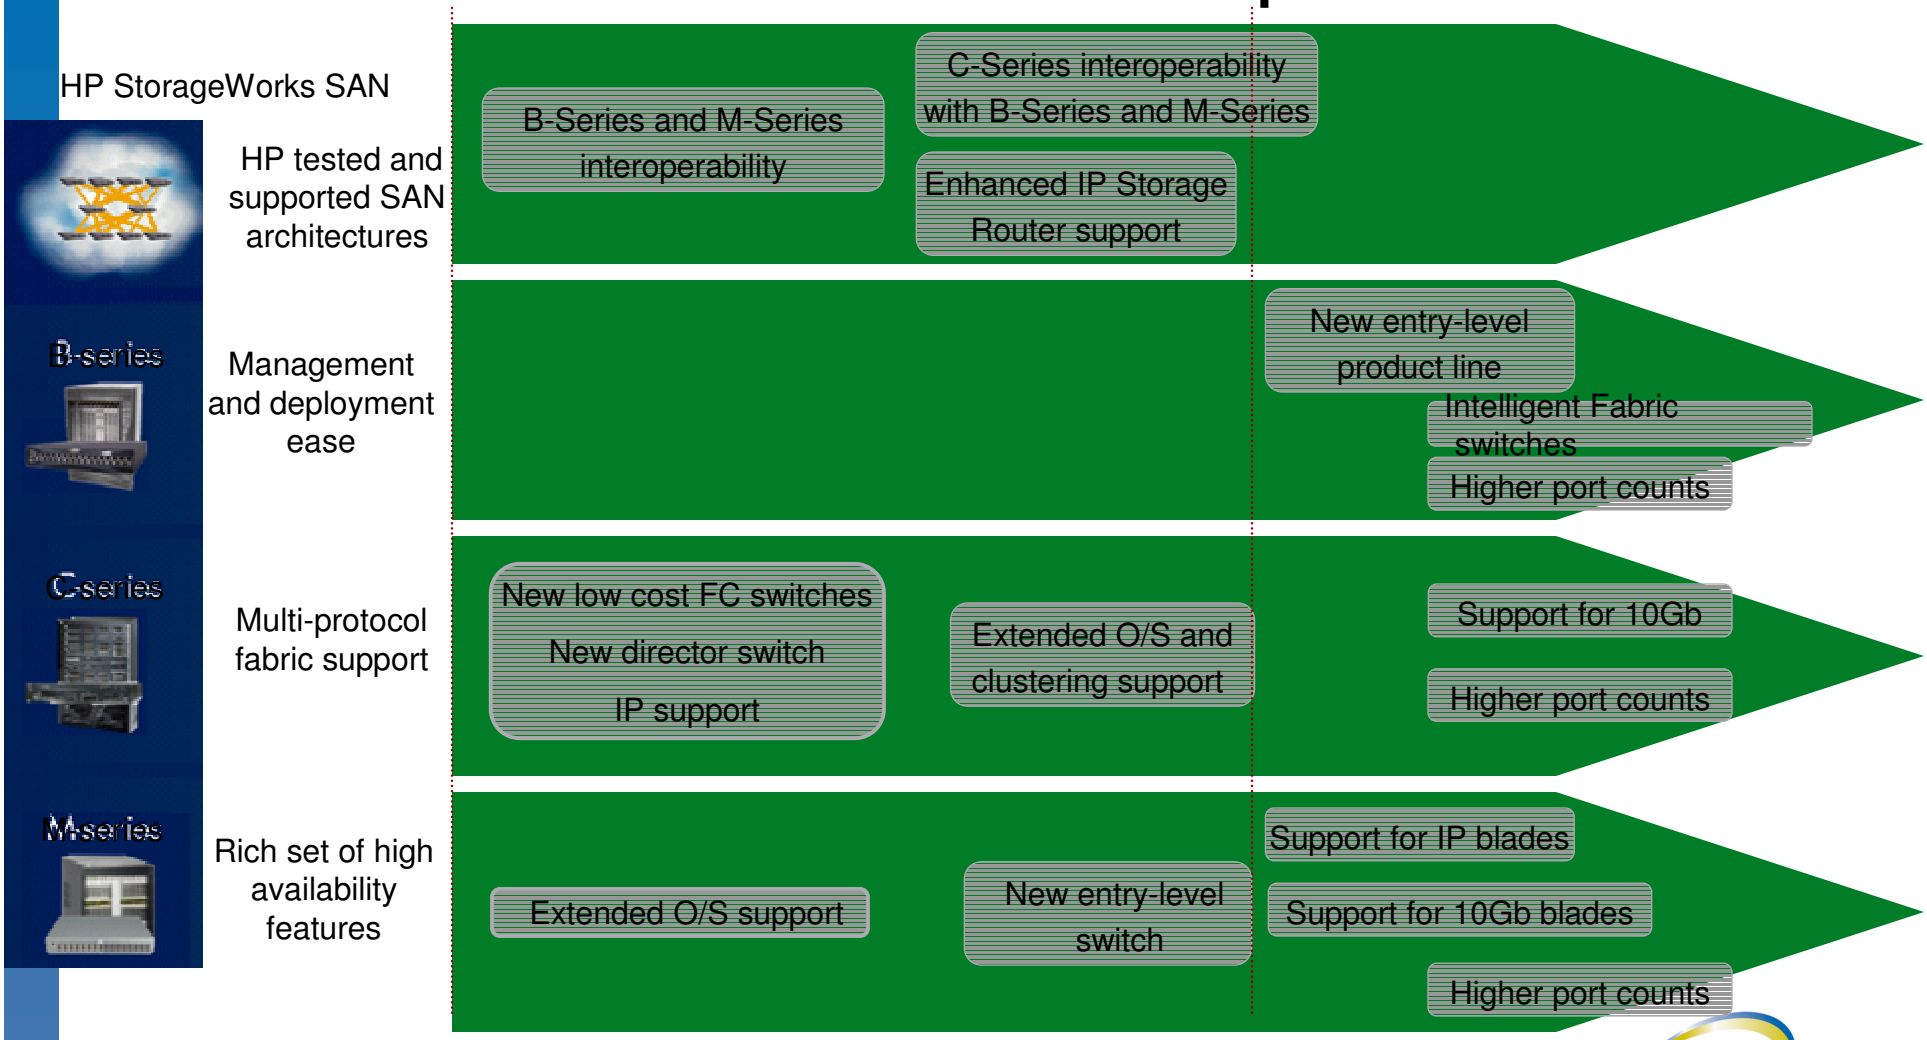

# Post V8.2 Projects

# Post V8.2 Projects

## *Fibre Channel Cluster Interconnect*

- Development project underway to send SCS traffic over FibreChannel
  - Provides LAN over FC
  - Use PEDRIVER to provide SCS communication
  - Goal is to provide stretch clusters without requiring additional cluster interconnect
    - May also have “cleaner” failure characteristics because SCS and storage will fail as a single unit
    - CI/MC class latency is not a goal (and not possible)
  - It is a non goal to provide general TCP/IP or DECNET over FC links (but we’ll probably get it for free)
  - Prototype currently being tested
  - Release commitment is OpenVMS V8.3

# Post V8.2 Projects

## *SATA / SAS Connectivity*

- Next generation host attached storage is moving towards SATA/SAS (serial ATA / serial SCSI)
- Currently support is planned for Integrity based platforms only
- Expect both motherboads based controllers as well as PCI based controllers

# Longer Term Storage Interconnects

# Long Term Storage Interconnects

- 4Gb FibreChannel
  - 2005
  - Gaining momentum for SAN usage
- 10Gb FibreChannel
  - 2006???
  - Very expensive infrastructure costs at first
  - Probably most interesting as inter-switch links
- iSCSI
  - Industry has stagnated some in recent years
  - Has some promise as a low-cost way to connect PCs to a SAN
  - Host performance overhead is the main issue today
  - OpenVMS will track development and consider support when native iSCSI storage is available and TCP/IP Offload Engines (TOEs), are available.

# HP WORLD 2004

Solutions and Technology Conference & Expo

Co-produced by:

interex

shared knowledge • shared power

encompass

AN HP USER GROUP

RECOMMENDED TRAINING VENUE FOR THE

**HP Certified Professional**