


# Server Storage – Features & Futures

John Stewardson  
EMEA Server Storage, Networking &  
Infrastructure - April 20<sup>th</sup> 2004

© 2004 Hewlett-Packard Development Company, L.P.  
The information contained herein is subject to change without notice





## Agenda

- Technology today
  - Ultra320 is here - Finally!
  - Key Features of Smart Array 6<sup>th</sup> Generation Controllers
  - Standard Smart Array Architecture Advantages
  - Universal Storage = further consolidation
- Technology tomorrow – The next Wave!
  - Serial Futures SAS & SATA: What & Why
  - Serial Features: New Compatibility Advantages
  - Does Size Matter – Serial SCSI enables Small Form Factors
  - But It's more than the interface...

21/04/2004

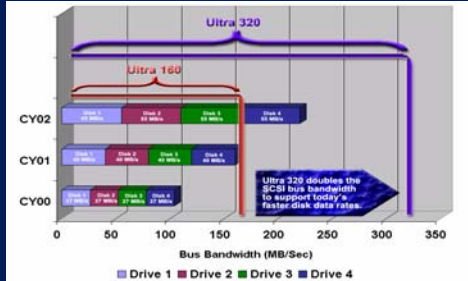
2

# What is Ultra320


- U160 (Ultra3) had reached the limits of it's bandwidth
- U320 replaces U160 as the standard parallel SCSI interface
- U320 is backwards compatible to U2 and U3 providing a clear upgrade path & ease of integration into existing infrastructure


- ✓ **Improved performance**
  - Double bandwidth of U160, up to 320MB/s per channel
  - Quick Arbitration and Select (QAS)
  - Data packetization
- ✓ **Improved signal and data integrity**
  - Active Adaptive Filtration (AAF)
  - Training pattern
  - Transmitter pre-compensation
  - Free running clock
  - Cyclic Redundancy Check (CRC)



21/04/2004 3


# Storage is not an option: The Smart Array Line-up






**smart array 5i Plus**

- low cost embedded RAID solution for ProLiant Servers
- lowest cost Battery Backed Write Cache RAID option for select ProLiant servers
- management of OS and log files




**SA-641**

- low cost entry level RAID solution
- option for servers without the embedded Smart Array 5i
- upgrade from previous generation RAID controllers
- upgradeable to BBWC
- management of OS and log files




**SA-642**

- low cost entry level RAID solution
- option for servers without the embedded Smart Array 5i
- upgrade from previous generation RAID controllers
- upgradeable to BBWC
- management of OS and log files



**SA-6402**

- high performance
- DAS beyond HDDs internal to the server
- ability to scale in the future (cache, channels)
- need for 2 int. or ext. channels
- maximum fault tolerance (ADG)
- database/exchange environments



**SA-6404**

- the 4 channel Premium controller
- high volume storage
- more than 2 channels needed for maximum data capacity
- maximum fault tolerance (ADG)
- database/exchange environments





Enterprise High Volume RAID

Workgroup and Departmental RAID

Entry level RAID

21/04/2004 4





## HP's Entry-Level U320 controllers: Smart Array 641 & 642



- Increased Performance
  - ✓ Ultra320 SCSI
  - ✓ 64-bit, 133MHz PCI-X
  - ✓ 266MHz DDR-cache
- Flexible Configuration ("on-demand")
  - ✓ **Memory option – 64MB BBWC upgrade**
  - ✓ Channels – 1 or 2
- Investment Protection
  - ✓ Upgradeable cache
  - ✓ Seamless upgrade to performance Smart Array
  - ✓ Generational Consistency between SA family controllers


21/04/2004 5

## HP's Performance U320 controllers: Smart Array 6400



- High Performance
  - ✓ Ultra320 / 64-bit, 133MHz PCI-X
  - ✓ 266MHz DDR-cache (**72 Bit wide!!**)
  - ✓ **Dual RAID / Processor (SA 6404)**
- Flexible Configuration ("on-demand")
  - ✓ Cache options – 128MB, 256MB
  - ✓ Channels – 2 to 4
- Higher Availability
  - ✓ **RAID ADG (Advanced Data Guarding)**
- Investment Protection
  - ✓ 2-4 channel upgrade (Available TODAY!)
  - ✓ Upgradeable cache (256MB +)
  - ✓ Generational Consistency between SA family controllers

21/04/2004 6



## HP Modular Storage Enclosures


### Up to 2TB per enclosure of storage\*

#### Mix and Match LVD HDD


- Ultra2, Ultra3, and Ultra320 devices on same Ultra320 SCSI bus
- Each device transfers data at its individual interface transfer rate (Requires Ultra320 controller )

#### Universal Support for:

- ProLiant, Integrity & Alpha Servers
- Smart Array Controllers Ultra2, Ultra3, Ultra320
- MSA 1000 & MA6000, MA8000, EMA12000, and EMA16000





Ultra320 single and dual I/O upgrade modules make upgrading to Ultra320 quick



\* Requires use of 146GB 1" HDD

21/04/2004 7




## Key Technologies:

### RAID ADG (advanced data guarding)

## Advanced RAID functionality for Smart Array

- Greater fault tolerance than RAID 1, 10, 5
- Lower implementation cost than RAID 1, 10



#### Can tolerate 2 simultaneous drive failures

No down time , No data loss


#### Ideal for:

- ✓ large capacity drives
- ✓ business critical data
- ✓ large RAID volumes
- ✓ remote server installations
- ✓ dense storage deployments



Available on SA-6400, MSA500, MSA1x00 Standard and SA-5302 (optional upgrade)

21/04/2004 8

## The HP advantage: Designed for the enterprise



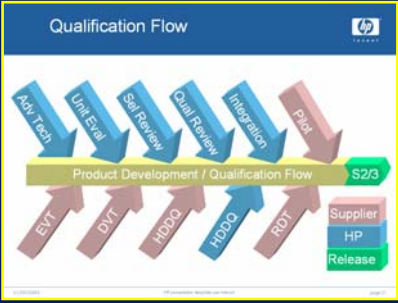
**Award Winning Design**

**Universal  
Hard Disk Drive family  
10K & 15K RPM**


- DAS, NAS, & SAN ready – re-deploy as your needs change
- Universal = rock solid spares strategy
- award winning ergonomic design allowing for single handed HDD removal & transportability
- designed for next enterprise cooling and vibration requirements

**Qualification Flow**



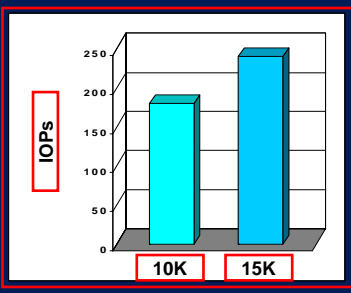
21/04/2004 9

## 15K: The NEED For SPEED!!



- **15K is more efficient:**
  - 40% less 15K rpm HDD to achieve same performance of 10K rpm HDD
  - Fewer Drives means less overall infrastructure = Less Power Required
  - Less Power means less Heat, and cost
- **15K opens the door**
  - reducing the opportunity for SCSI bus saturation in system,
  - increased data rate allows all SCSI devices to perform at optimum levels!!
- **For more performance!**
  - HP's 15,000 rpm Hard Drives can provide up to 33-50% better performance over standard 10,000 rpm drives!!

- ✓ A server with multiple drives handling large numbers of transactions benefit from high IOPS performance
- ✓ Typically 35% IOPS improvement vs 10K



21/04/2004 10





## HP – Inventing Pure Performance...

- Smart Array & ProLiant have dominated the TPC benchmark lists for years . . .
- HP Integrity Servers are now starting to lead the pack . . .
- Both ProLiant & Integrity support Smart Array & Universal Storage Enclosures for a reason...

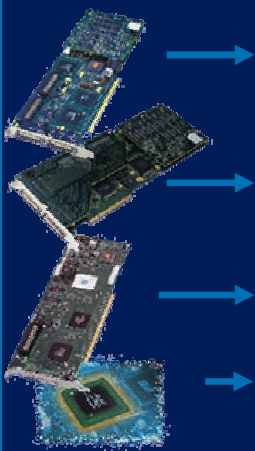




21/04/2004 11





## Upgrading Your Smart Array



**Instant Upgrade:**

1. Ensure that you are running the latest device driver and Smart Array firmware.
2. Replace the old controller and cables with the new controller and cables.
3. Reboot the server and data will be seamlessly recognized.






Smart Array 6402

- Faster performance
- Greater capacity
- Higher availability

21/04/2004 12


## Key Competitive silver bullets SUMMARY: Smart Array Advantage

- ✓ **Unmatched performance**
  - TPC Benchmark leadership
  - High Performance Architecture
  - Ultra320 & PCI-X across the board!
- ✓ **Generational Consistency**
  - Consistent configuration and management tools
  - Proven lower cost of ownership
  - Pre-failure warranty with Universal HDDs
- ✓ **Unparalleled Flexibility**
  - Data compatibility from generation to generation
  - Ease of upgrade of & between controllers
  - Support for data migration & DtS Technology with Universal HDD
  - Mix and Match LVD Technologies (Ultra2, Ultra3, Ultra320)

21/04/2004 13


## Direction: Enterprise product commonality



	# of product families	2003	2004
Hard disk	7	Pre-merger Universal	Small form factor universal SAS and SATA
Smart Array	2	Smart Array and Net RAIDCtrls & JBODs	Smart Array Controllers & JBODs
Racks	5	Consolidated to 2 Racks – 10k and HP System E Rack	Universal Rack across ESG
Power	2	Common UPS platforms	Common Intelligent Power Dist Units
Consoling	2	Common KVM Analog switch	Common IP Console Switch

*Potential to eliminate 100's of products, simplified solutions, and save millions of dollars in warehousing, logistics, P/N maintenance, and channel positioning*


21/04/2004 14



## Smart Array & MSA30 on Integrity

- SmartArray 6402 / 128MB cache for Windows A9890A
- SmartArray 6404 / 256MB cache for Windows A9891A
  - [http://www.hp.com/products1/servers/integrity/mid\\_range/rx7620/supplies.html](http://www.hp.com/products1/servers/integrity/mid_range/rx7620/supplies.html)
  - [http://www.hp.com/products1/servers/integrity/mid\\_range/rx8620/supplies.html](http://www.hp.com/products1/servers/integrity/mid_range/rx8620/supplies.html)
- Smart array 5302/128MB controller (Windows & Linux) A9825A
- Smart array 5304/256MB controller (Windows & Linux) A9826A
  - [http://www.hp.com/products1/servers/integrity/entry\\_level/rx5670/supplies.html](http://www.hp.com/products1/servers/integrity/entry_level/rx5670/supplies.html)
  - [http://www.hp.com/products1/servers/integrity/entry\\_level/rx2600/supplies.html](http://www.hp.com/products1/servers/integrity/entry_level/rx2600/supplies.html)

21/04/2004 15



## Smart Array and AlphaServers

- SA5302A 2-channel controller with 128 MB of cache
- SA5304A 4-channel controller with 256 MB of cache
- HP StorageWorks enclosure models 4354, 4314, 4254, 4214, and 4400 (MSA30)
- Supported servers
  - AlphaServer DS10
  - AlphaServer DS25
  - AlphaServer ES45
  - AlphaServer GS80
  - AlphaServer GS160
  - AlphaServer GS1280
  - AlphaServer DS20E
  - AlphaServer ES40
  - AlphaServer ES47
  - AlphaServer ES80
  - AlphaServer GS320

21/04/2004 16





## What next – coming this summer

- Completion of the 6<sup>th</sup> Generation controllers
  - Smart Array 6i
  - BBWC enabler
- Introduction of new serial bus technology to the 6<sup>th</sup> Generation: PCI-Express
- Further Qualifications of SA 6<sup>th</sup> Gen controllers
- And the beginning of the next generation...
  - Introduction of SATA JBODs into the MSA Line up

21/04/2004 17



## The next chapter . . .

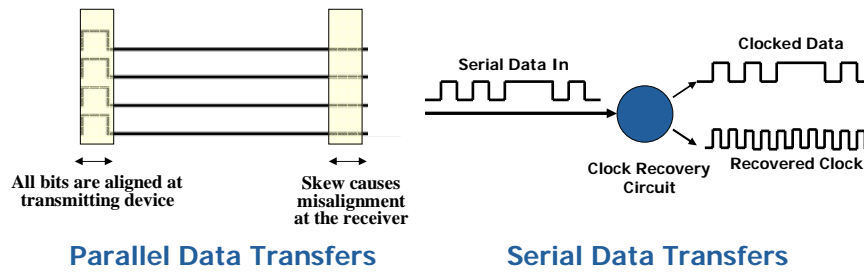
- Technology today
  - Ultra320 is here!
  - Key Features of Smart Array 6<sup>th</sup> Generation Controllers
  - Standard Smart Array Architecture Advantages
  - Universal Storage = further consolidation
- Serial Storage...
  - Serial Futures SAS & SATA: What & Why
  - Serial Features: New Compatibility Advantages
  - It's more than the interface...
  - Does Size matter ?

21/04/2004 18



## Problems with going to U640...

- With a parallel bus, all data must be arrive within in a given time period.
- Signal skew causes data to become misaligned in time at the receiver limiting the rate at which data can be transferred.
- Serial technologies uses a single data line so there is no skew.
- With serial technologies, much higher data transfer rates can be obtained per signal line.

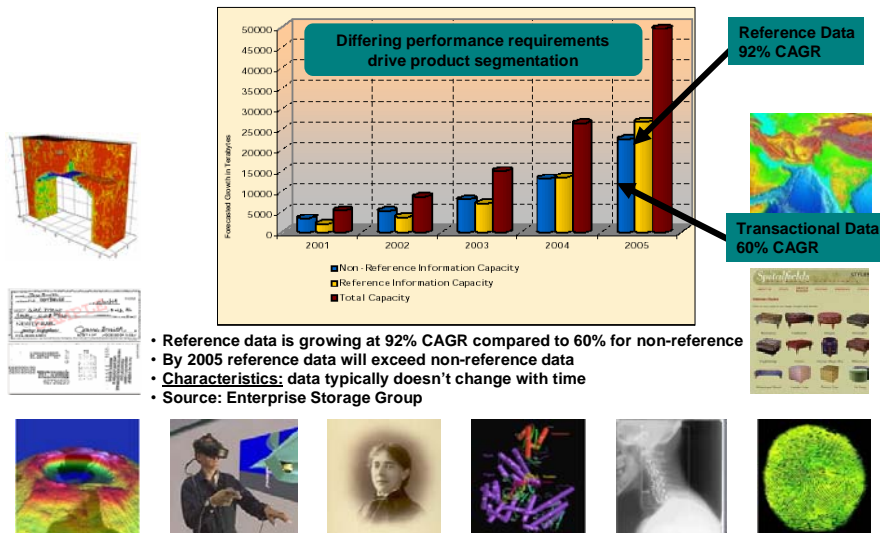


21/04/2004

19

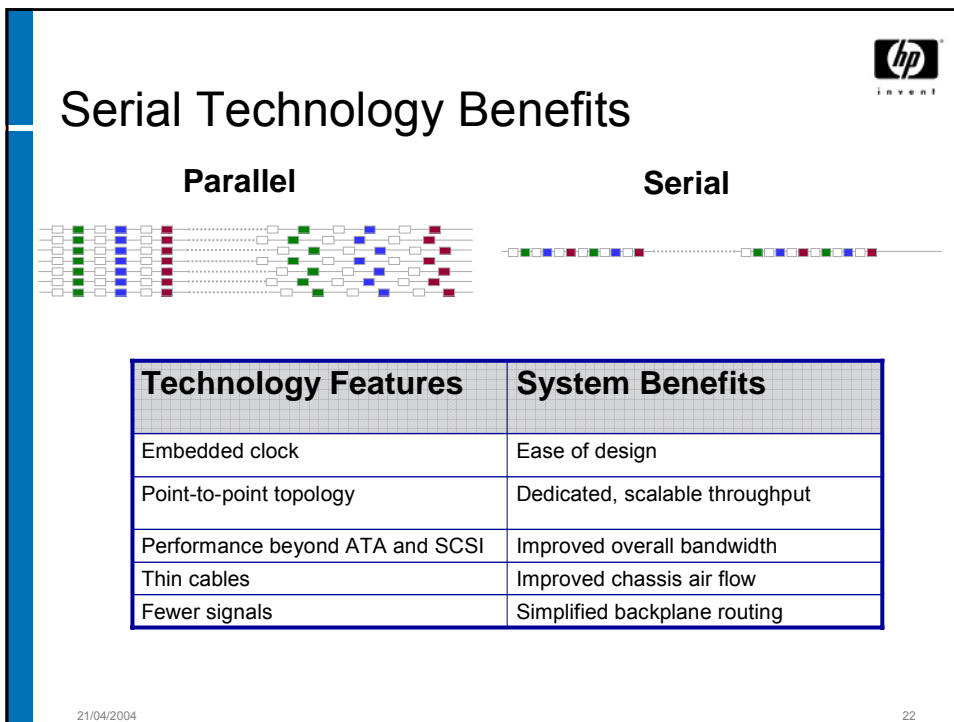
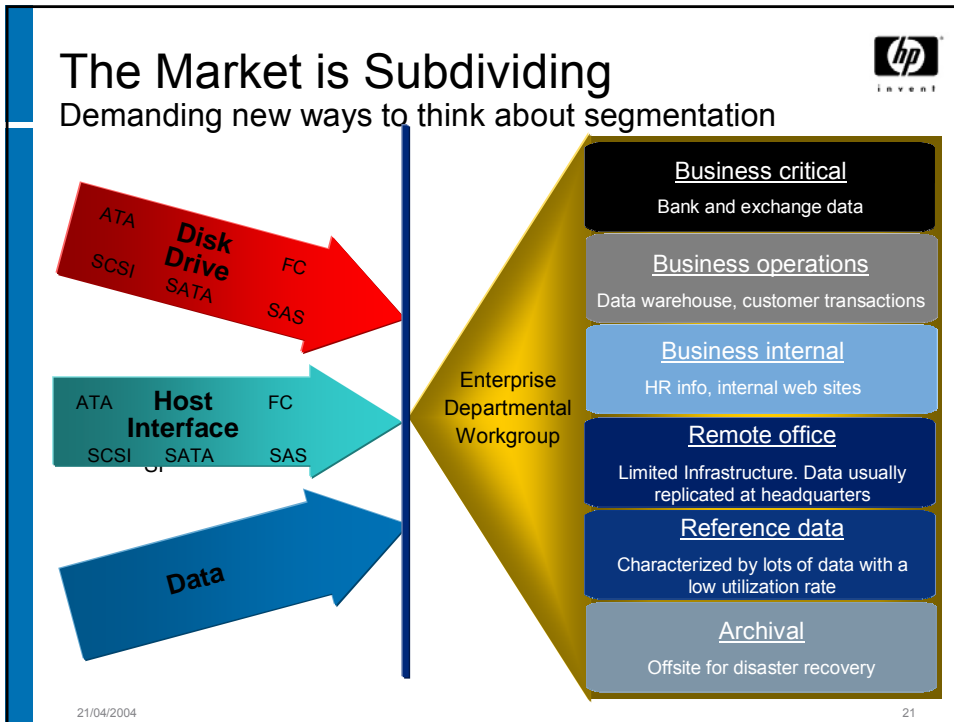



## Explosive Growth of Reference Data



21/04/2004

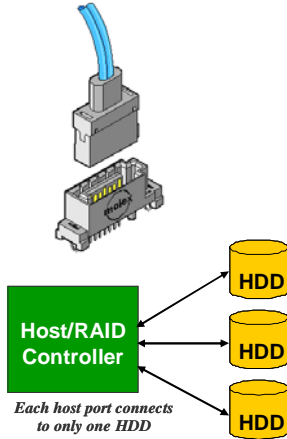
20






## What is Serial ATA?

- **Four-wire replacement for the physical layer of parallel ATA**
  - 100% SW compatible “drop in” replacement for ATA
  - Low voltage (600 mV max) differential signaling w/ 8b10b encoding
- **“Star” topology (point-to-point, no hubs)**
  - Each device gets full bandwidth
  - No bus arbitration/collision overhead
  - ATA RAID becomes simpler to implement
- **Serial ATA has attributes and features that extend its capabilities**
  - Provides additional capabilities such as hot plug and 1st party DMA



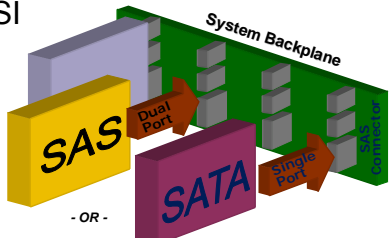
SATA connector drawing courtesy of Molex

21/04/2004
23



## What is Serial Attached SCSI?


- **Four-wire replacement for the physical layer of parallel SCSI**
  - SW compatible “drop in” replacement for parallel SCSI
  - Low voltage (1600 mV max) differential signaling w/ 8b10b encoding
- **“Star” topology (point-to-point, no hubs)**
  - Each device gets full bandwidth
  - Minimum arbitration overhead
- **SAS has attributes and features that can extend it beyond the capabilities of parallel SCSI**
  - Provides many additional superset features (such as dual-port and 16K devices per domain)
  - Mix / Match with SATA
  - Configurable Bandwidth!




- OR -

21/04/2004
24


## Why Serial Attached SCSI?




**Desktop PC  
Entry NAS  
Workstation**



**Servers  
Mid-Line Storage  
Mainstream NAS/SAN**



**Large Enterprise  
SAN & NAS**



**Drive Interface**

PC Chipset  
SAS HBA

Dual Mode SAS/SATA  
HBA & RAID

SAS RAID

FC RAID

**Drive Types**

SATA Drive  
5400 RPM  
7200 RPM

SAS Drive  
10K RPM  
15K RPM

FC Drive

**SAS supports SATA devices**

- Enables systems to be configured with SATA or SAS drives
- Provides More flexibility to channel partners and end users

**SAS was developed as a serial replacement to SCSI**


- Difficult to go beyond U320 – clock skew issues
- Need for smaller & simpler connections
- Preservation of SCSI protocol and cost structure for primary markets

**SAS includes enterprise features with parallel SCSI cost structure**

- Dual port devices
- High connectivity
- More scalable performance
- Better system interoperability
- Enhanced fault isolation

21/04/2004
25

## Technology Comparison of Interface



	SATA	SAS	Fibre Channel FCAL
<b>Performance</b>	Half-duplex	Full-duplex with Link Aggregation	Full Duplex
	1.5 Gb/s Introduction 3.0 Gb/s in 2004	3.0 Gb/s Introduction 6.0 Gb/s in 2006	2.0 Gb/s 4 Gb/s in 2004
<b>Connectivity</b>	1.5 m internal cable > 8 m external cable	> 8 m internal and external cables	15 m external cable
	One device (fan-out devices demonstrated)	>128 devices (16,384 max)	127 devices
	No peer-to-peer	Peer-to-peer	Peer-to-peer
<b>Availability</b>	Single-port HDDs	Dual-port HDDs	Dual-port HDDs
	Single-host	Multi-initiator	Multi-initiator
<b>Driver Model</b>	Software transparent with Parallel ATA	Software transparent with Parallel SCSI	Software transparent with Parallel SCSI

21/04/2004
26



## SAS Expander Devices

- SAS edge expander sets
  - Contain one or more edge expander devices (provides for increased connectivity at lower cost)
  - Connect to SAS initiators, SAS targets, SATA devices, and/or one SAS expander device via narrow or wide ports
  - May only have ports in one domain
- SAS fan-out expander devices
  - Connect to SAS initiators, SAS targets, SATA devices, and/or SAS edge expander sets via narrow or wide ports
  - May only have ports in one domain
  - Only 1 fan-out expander per SAS domain

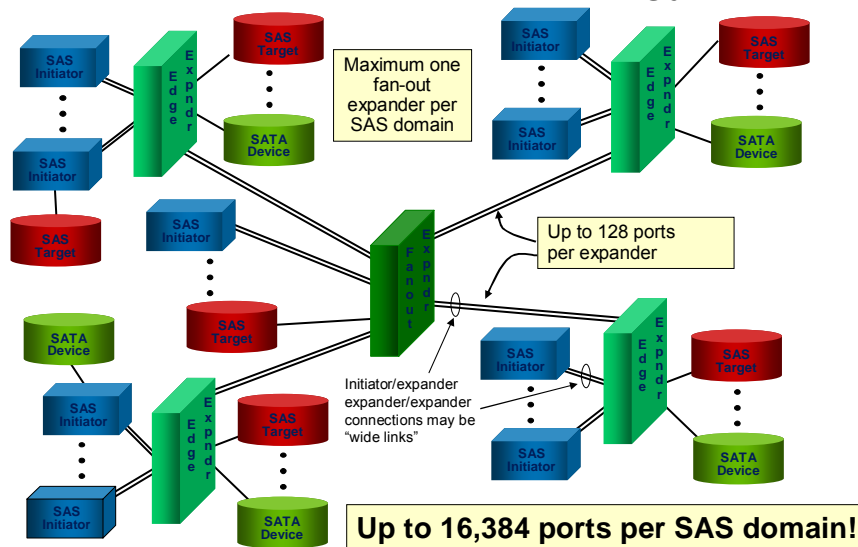


21/04/2004

27

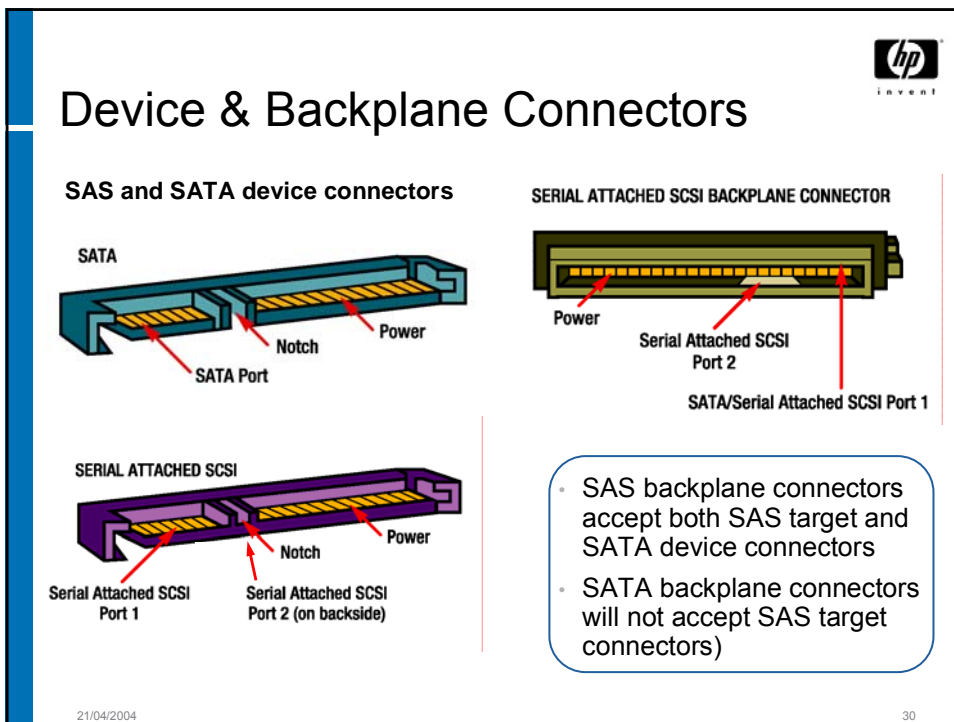
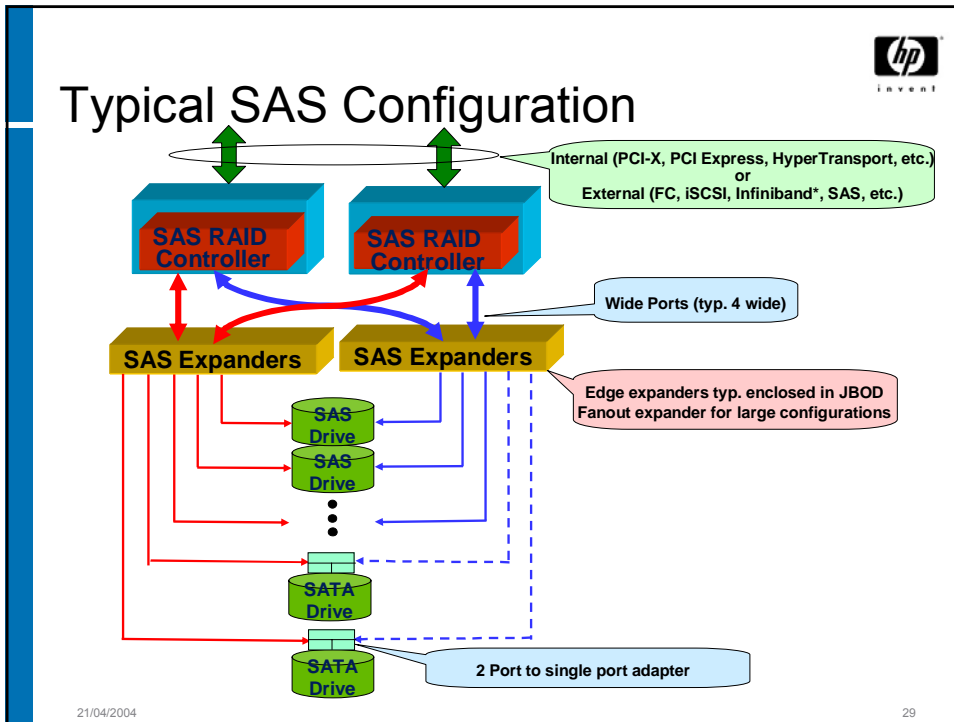


## Maximum SAS Domain Topology




21/04/2004

28

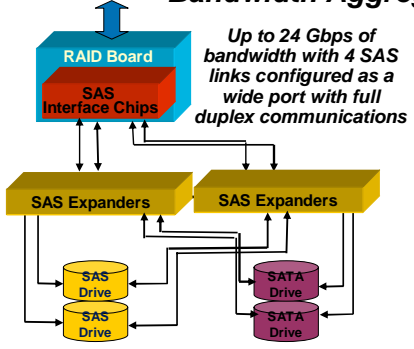


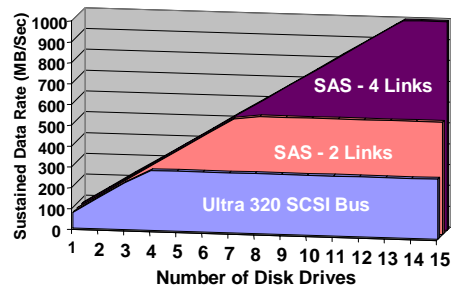
## SAS Delivers Scalable Performance



### Bandwidth Aggregation with Wide Ports

Up to 24 Gbps of bandwidth with 4 SAS links configured as a wide port with full duplex communications






Number of Disk Drives	Ultra 320 SCSI Bus (MB/Sec)	SAS - 2 Links (MB/Sec)	SAS - 4 Links (MB/Sec)	Total (MB/Sec)
1	100	0	0	100
2	100	0	0	100
3	100	0	0	100
4	100	0	0	100
5	100	0	0	100
6	100	0	0	100
7	100	0	0	100
8	100	0	0	100
9	100	0	0	100
10	100	100	0	200
11	100	100	0	200
12	100	100	0	200
13	100	100	0	200
14	100	100	0	200
15	100	100	100	300


- Four link, SAS controller provides scalable back-end bandwidth for up to 24Gbps data transfers with full duplex communications.
- Point-To-Point connections eliminate the shared bus bottleneck.
  - Bandwidth aggregated over multiple, low-cost Serial Attached SCSI links
  - Rate matching supports 1.5Gbps SATA and 3.0Gbps SAS drives for optimal performance

21/04/2004
31

## External Cable and Connector



(Photo courtesy of FCI/Berg)




- Four, full-duplex physical links per connector
- $4 \times 2 \times 3.0 \text{ Gb/s} = 24 \text{ Gb/s} = 2,400 \text{ MB/s max}$

21/04/2004
32



## Cable Configurations



**SATA 1 - Internal only**


**SATA 2 or SAS — HBA and Cabled HDDs**

**SATA 2 or SAS - HBA and BP**

**SAS - Dual / Redundant HBA and BP**

33

## Low-cost Direct Attached Clusters




Previously not possible with Parallel SCSI, or with SATA & expensive using FC HDD

SAS provides a near-cabinet interface for clusters with direct attached storage

**Serial Attached SCSI Supports Cost Efficient Clusters**

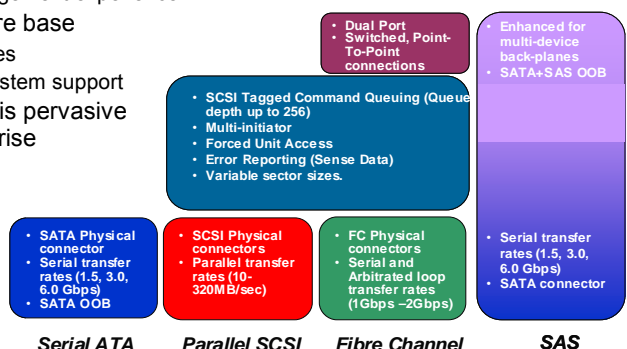
21/04/2004
34

# Enterprise Software Support




- SAS employs the SCSI command set to maintain compatibility with installed base of enterprise applications.
  - SCSI command set is robust and stable
    - 20 years of investment and enhancements
    - Extensive IT management experience
  - Entrenched software base
    - Management utilities
    - Broad operating system support
  - The SCSI protocol is pervasive in servers & enterprise storage
    - FC, iSCSI, IB...

### Extending Technology Investments




21/04/2004 35


# SAS for Mixed Online/Midline Apps




**Dual port SAS drives for main stream, enterprise storage applications**




**SATA drives integrate disk to disk backup within the RAID array or server to shorten backup and restore times**



**Dual Port SAS drives for networked file storage**

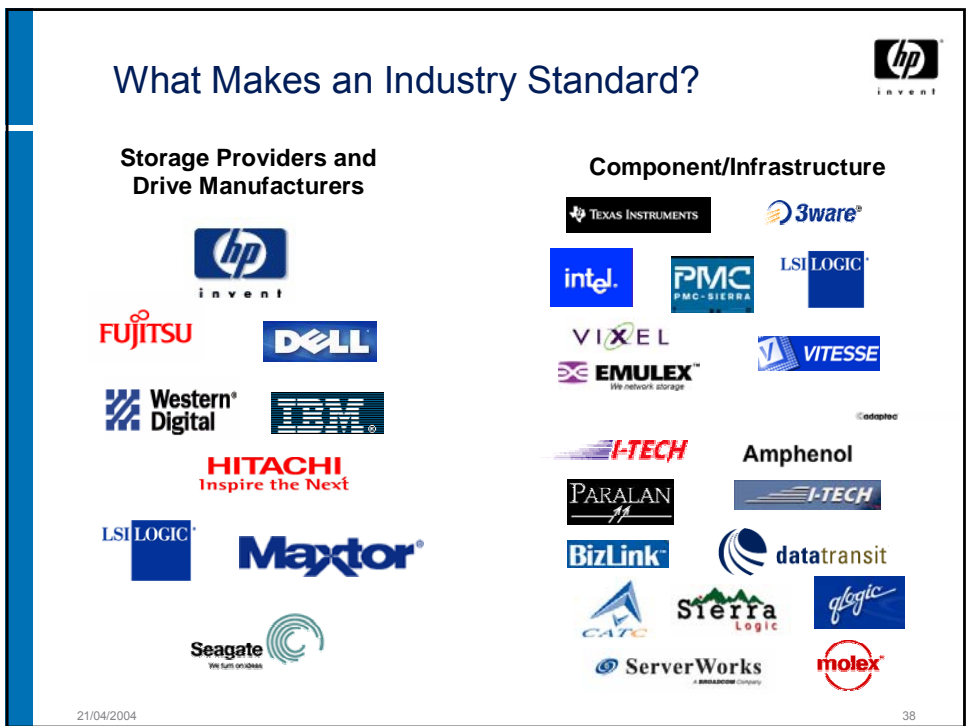
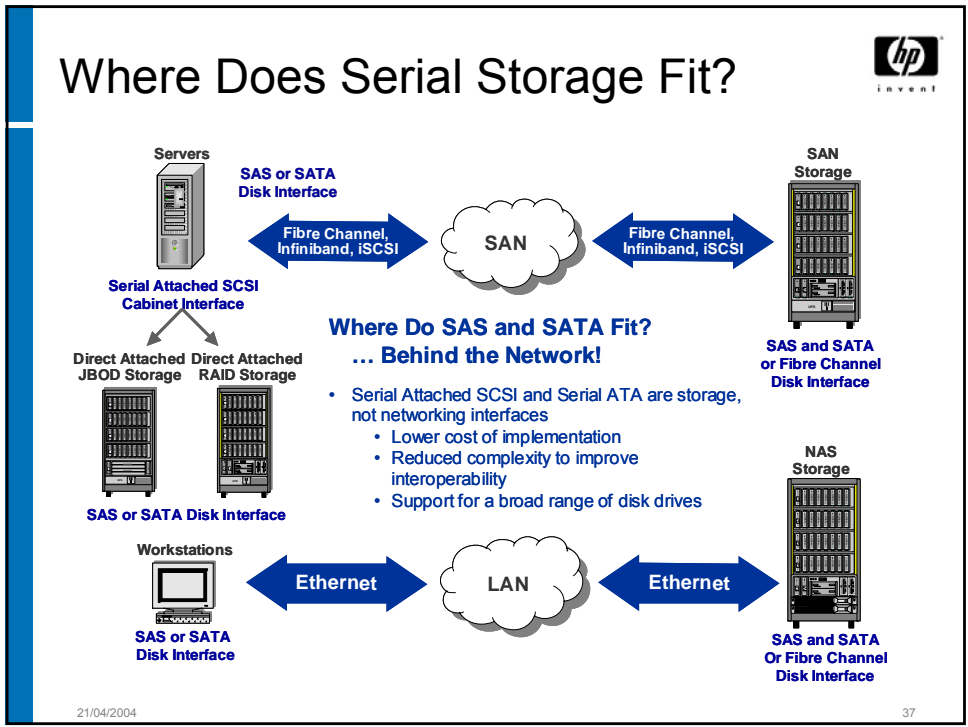



**Serial Attached SCSI will support mixed classes of disk drives within a single JBOD or RAID enclosure**



**SAS drives      SATA drives**

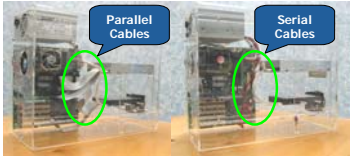
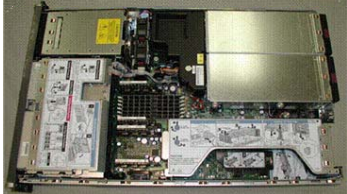

21/04/2004 36






## Why Small Form Factors?

- Product Challenges
  - dense architectures
  - thermal envelopes overloading
  - power draws overloading
  - airflow designs must improve
  - new processors run hotter
  - natural technology evolution
- Customer Requirements
  - Need for increased performance vs. capacity requirements in data center.
  - Want more spindles, not more capacity.
  - Data protection is a must
  - RAID rebuild times need to speed up (but increase with disk drive capacity)



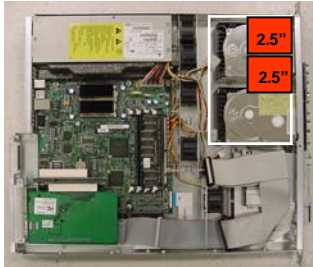
Photos courtesy of Intel

21/04/200439

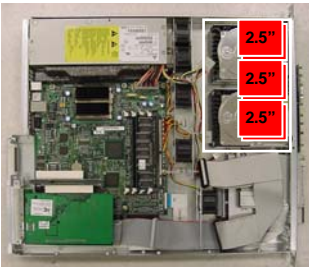


## The Benefits of 2.5" Disk Drives

**Lower power  
& Improved Airflow**




**More Drives per Server**



- 2.5" drives will consume 70% less space and 30% less power than 3.5" drives enterprise drives
- Enables either reduced power consumption and improved airflow and/or more drives per server.

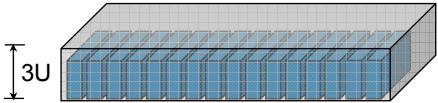
21/04/200440

## SFF = Better Performance Footprint



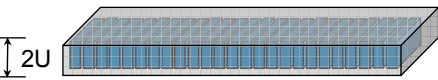
- SAS enables design of smaller form factor drives to deliver higher performance in space constrained environments

**14 x 3.5" Drives in 3U**



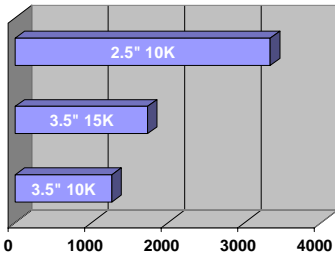
3U

**25 x 2.5" Drives in 2U**



2U


**IOPS Per U**



Form Factor	Performance	IOPS Per U (Approx.)
2.5" 10K	High	~3500
3.5" 15K	Medium	~1800
3.5" 10K	Low	~1200

21/04/2004 41

## Small Form Factor rational for suppliers: Size for performance?




→ More performance means reduction of seek times (head positioning time)

**3.5" FF**

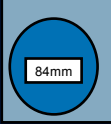
<18W (max)

7.2K rpm




95mm

10K rpm



84mm


15K rpm



65mm

?

22K rpm




48mm

- Enterprise 2.5" (18GB 15k) platters available for years
- Further development of 2.5"
  - Easier short term than going to 48mm platters
  - 2x size decrease 4x platter rigidity = better reliability

**2.5" FF**


<11W (max)

10K rpm




65mm

15K rpm



48mm

21/04/2004 42



## Summary of Server Storage


### Current Technology

- Ultra320 is here and brings loads of performances
- The summer will see the final turn of Ultra320
- The Server Bus begins to go Serial with PCI - Express

### The Next Generation

- Serial Storage is starting to roll out with SATA products
- SAS will begin about 6 Months later and allows native interoperability between SCSI & ATA HDD
- SAS brings more flexibility, better performance, and enterprise reliability to the Serial Storage revolution
- SAS also allows Small Form Factor Drives to begin rolling out.

21/04/2004 43



i n v e n t