

ProLiant Server Storage *the future*

Lonnie Pope
Systems Manager
Product, Development
ISS Server Storage



agenda

- Disk Drive directions
 - Ultra 320
 - The transition to serial
 - Small Form Factor
 - Roadmap
- iSCSI
- The system bus (PCI-X / PCI Express)
- q&a

Disk Drive Directions ultra320



U320 Transition

- U320 Hard drives began shipping in 4Q02
- Entry level array controllers (SA641 / SA642) began shipping in 2Q 03
- Mainstream array controllers (SA6402 /SA6404) will ship in 3Q03
- U320 JBOD (4414/4454) will ship in 3Q03

Disk Drive Directions the transition to serial



background

- Three interfaces exist today, each serving unique segments
 - serial fibre channel
 - parallel SCSI
 - parallel ata
- Due to signal integrity challenges all hdd interfaces are moving to serial
 - U320 has been very difficult for the industry to bring to market
 - U640 would have been an even more formidable challenge
- FC is serial already
- PATA was going to SATA
- The question 2 years ago was what will replace pSCSI
 - fibre channel is too much, serial ata is too little
 - sSCSI is the answer

What is Serial ATA ?

- **Serial ATA (SATA)** is the next evolution and replacement of the **Parallel ATA** interface currently in development with **product announcements slated for late 2003 or early 2004**.
- **SATA is a 1.5 Gb/s serial point to point** architecture with a primary focus on low cost.
- **SATA is driven by cost** and coupled with desktop, notebook, entry level servers, and low cost non-mission critical server storage solutions.

What is Serial Attach SCSI ?

- **Serial Attach SCSI (SAS)** is the next evolution of the SCSI interface with product announcements slated for 2H04.
- SAS will provide enterprise storage solutions cost competitive with parallel SCSI
- **SAS** inherits attributes from parallel SCSI (command set), Fibre Channel (frame formats), and Serial ATA (physical characteristics).
- **SAS** uses expanders (low cost switches) to extend the point to point physical characteristics of Serial ATA and doubles the data rate to 3.0 Gb/s.

breaking the silos

desktop PC
workstation



PC Chipset

SAS HDDs
connect
via HBA

S-ATA

HDD

mainstream servers
DAS deployment



SAS HBA / Chipset / RAID

one design
accommodates
SATA or SAS
for
a customized
solution

SAS

HDD

storage area
networks



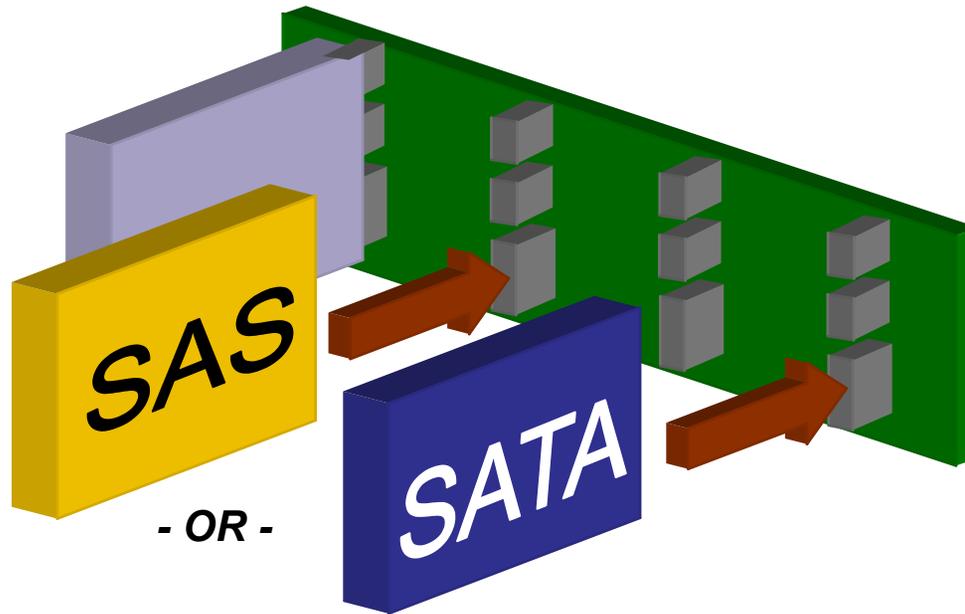
FC HBA / RAID

FC

HDD

serial attached SCSI helps to bind FC, SCSI, & SATA solutions

Enabling Customer Choice



A "dual mode" backplane can accommodate either SAS or SATA disk drives

- SATA/High-Capacity disk drives can be used to enable "near-line" or tape augmentation applications
- SAS/High-Performance disk drives can be used to enable "on-line" and performance-oriented applications

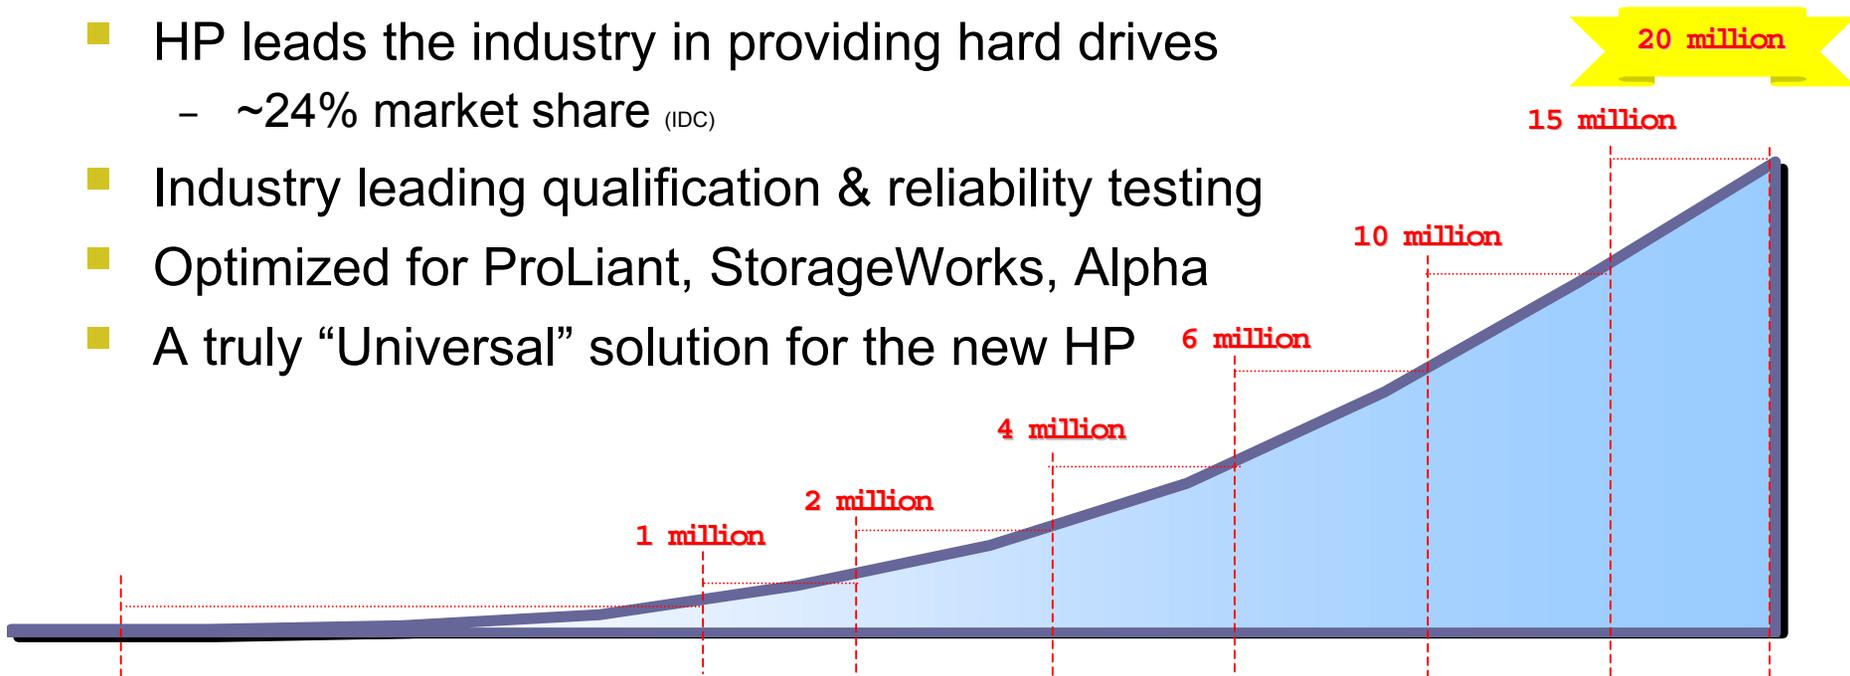
HP will provide customers the ability to deploy and redeploy a server or storage solution, and customize the internal storage between SAS & SATA

why is this important for the customer?

- choice drives the adaptive infrastructure
 - standardize on server or storage platform
 - customize storage by application and customer requirements
- minimize costs (acquisition and TCO)
 - service pools of hardware can be minimized/standardized
 - allow customers to standardize infrastructure
 - Standardize on electrical and physical connection scheme
- exceeding the customer requirement
 - SAS will enhance 'abilities from p-SCSI (reliability, performance)
 - SATA will offer lowest cost with desktop device mechanics

hp leading storage provider

- HP leads the industry in providing hard drives
 - ~24% market share (IDC)
- Industry leading qualification & reliability testing
- Optimized for ProLiant, StorageWorks, Alpha
- A truly “Universal” solution for the new HP



10 Years – 20 Million Hard Disk Drives

SATA HDD attributes

Same reliability as P-ATA

- ATA: 600K MTBF >5000 hrs/yr with desktop workload
- SCSI: 1.5M MTBF 8760 hrs /yr with enterprise workload
- desktop device or mechanism leveraged by the enterprise for cost

Hot plug issues

- no design, electrical or test specification
- reliance on vendor resources for HDD qual
- emerging spec caught between the desktop and the enterprise requirements (enhanced functionality without cost adder)

cost

- 15% to 20% cost premium to PATA + cost of drive carrier
- backplane and drive cage
- solution will have to absorb extra cost to offset rotational vibration

Will ATA dominate the enterprise?

ATA opens new doors, but does not cannibalize SCSI
ATA usage is targeted at tape augmentation, and low duty cycle environments

Be careful over-positioning reliability of ATA devices – this technology is a desktop class design based on strict cost controls

Mainstream and mission critical storage will continue to use enterprise class disk drives (serial SCSI) based on features:

higher availability

- dual-port disk drives with redundant controllers

- higher data integrity & MTBF specifications

- fast RAID re-builds

- better error handling

better performance

- more IOPs per drive bay

- higher sustained data rates

Does a hard drive matter?

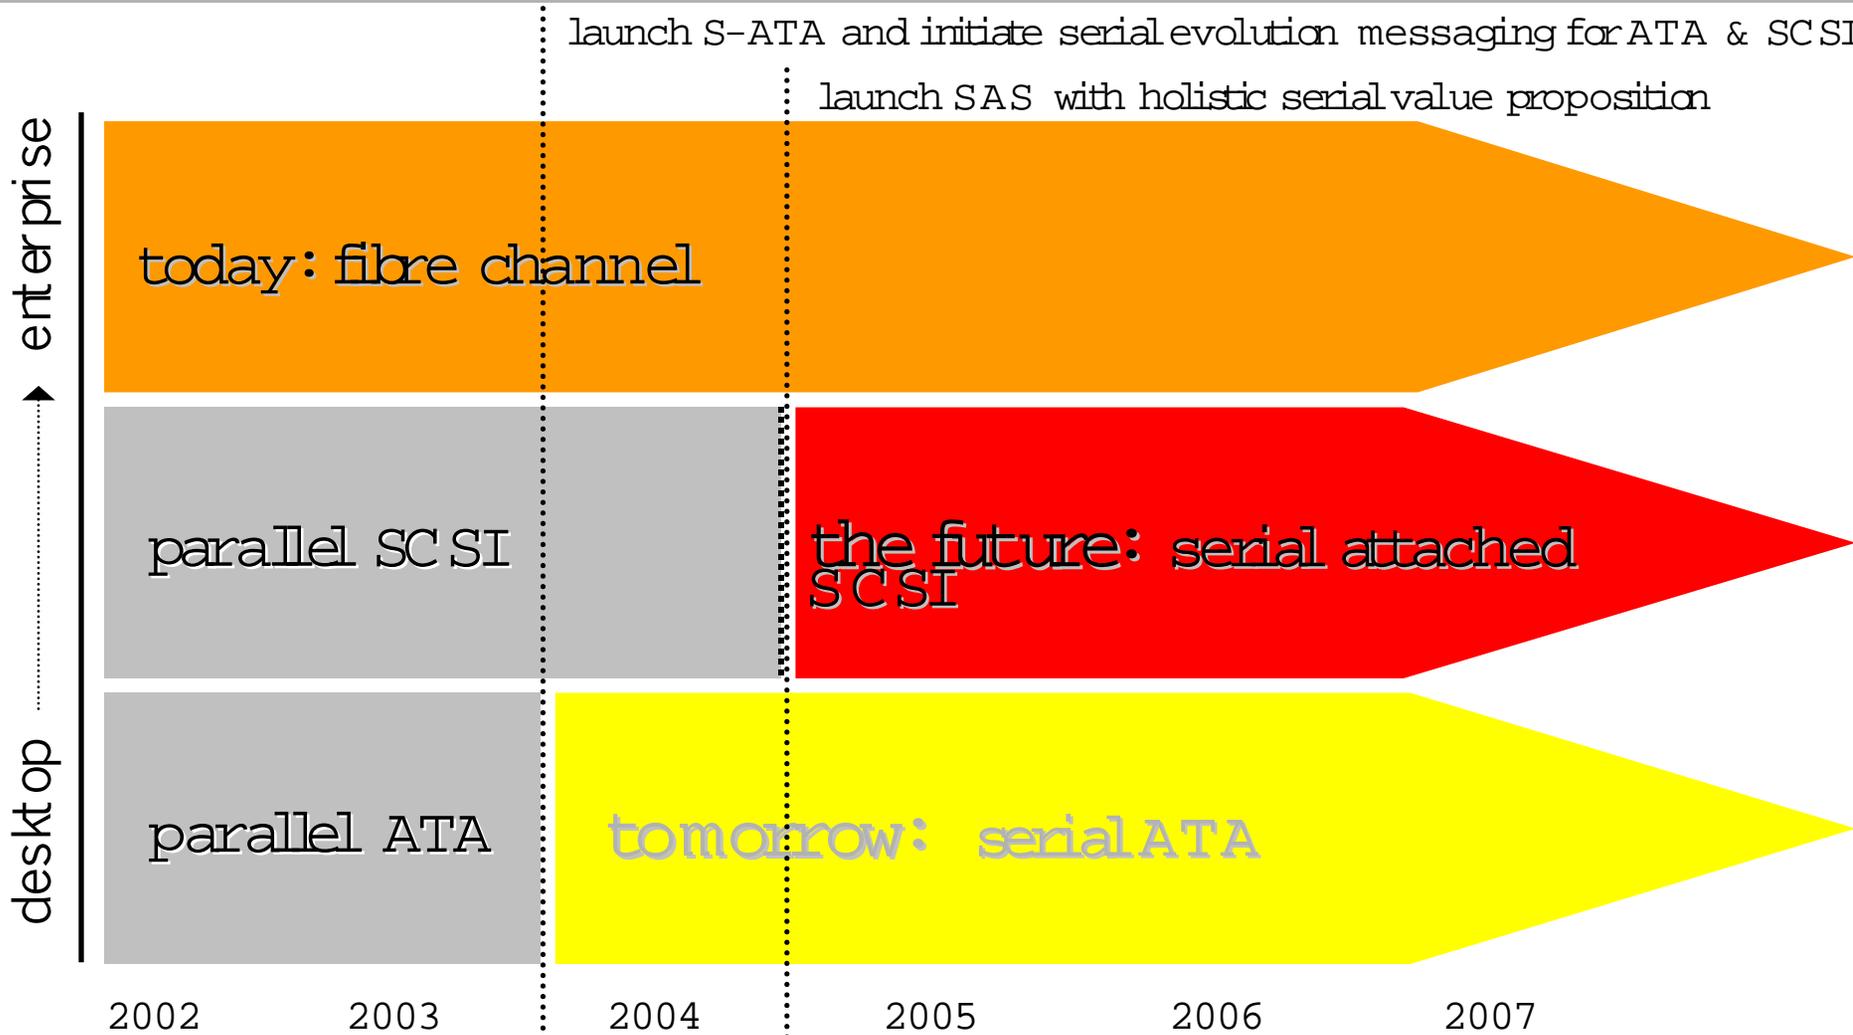
core product attributes:

- it's where your data lives!
- server optimized
- external Storage optimized
- support full 'solution' migration
- industry leading qualification process
- most rigid test and Q/A
- pre-failure warranty
- not every hard drive is created equal

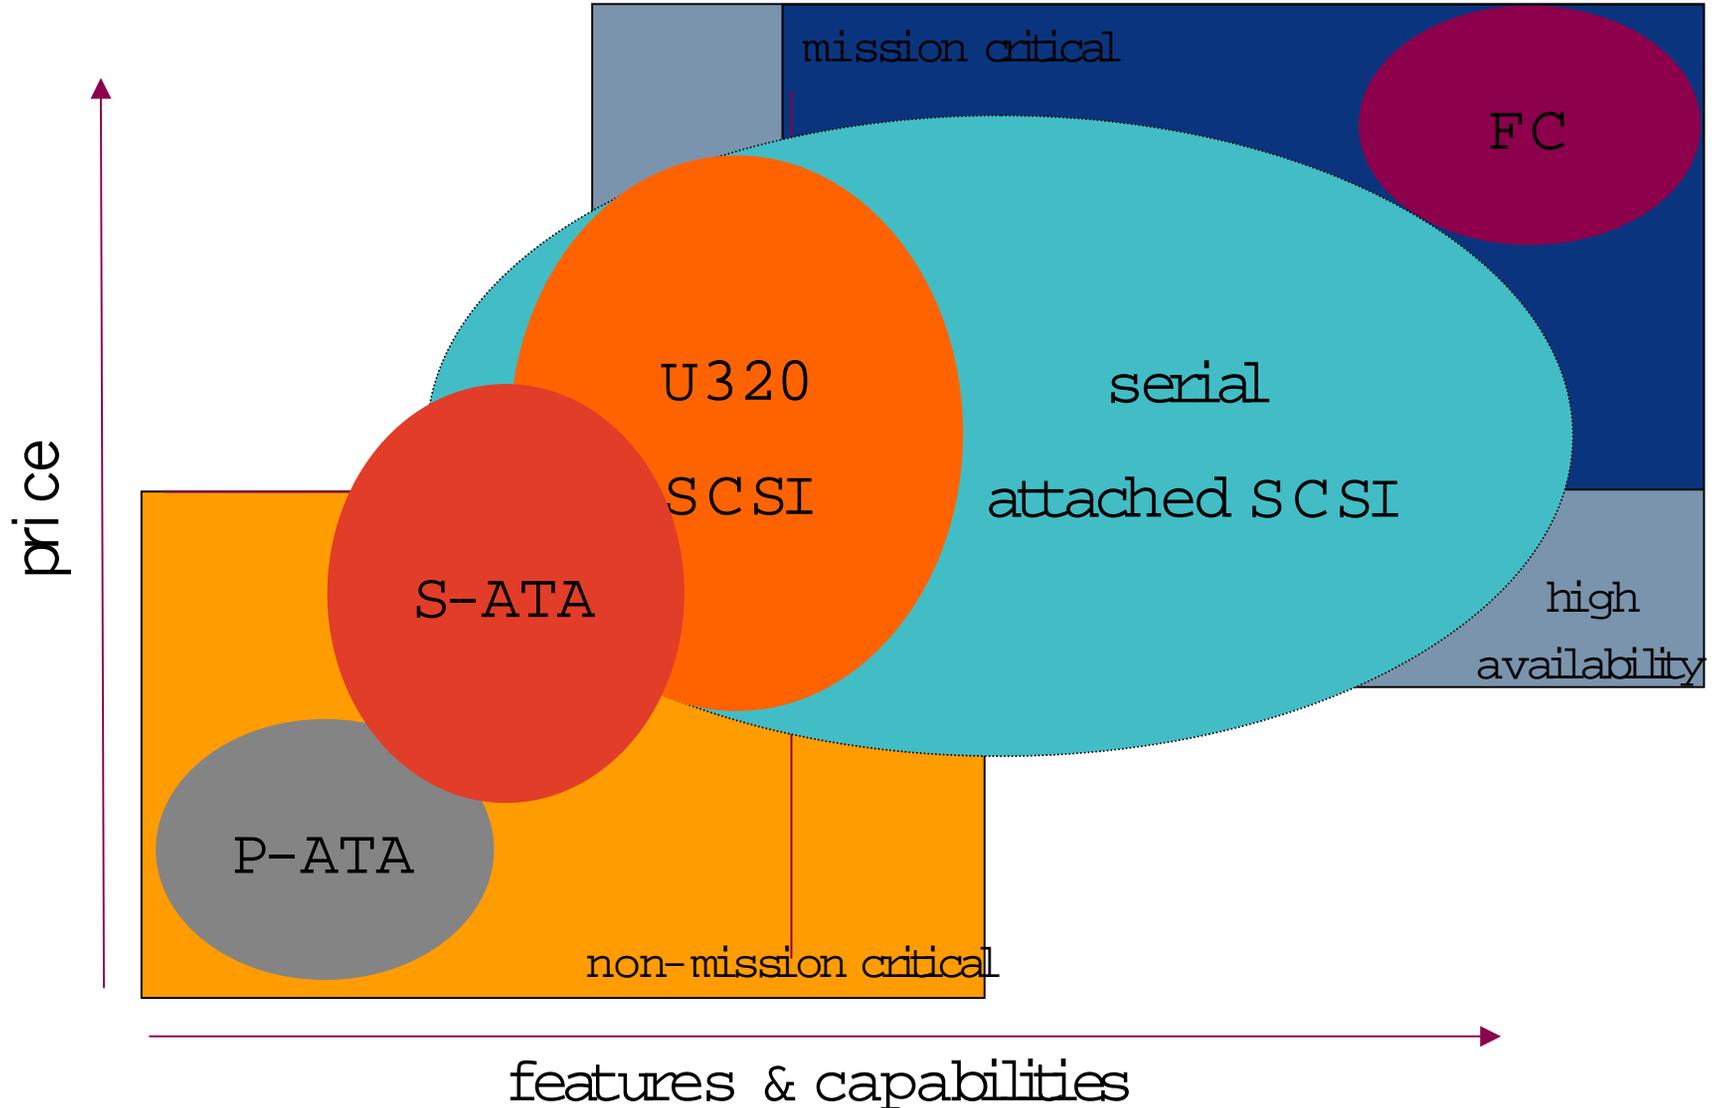


do you care about performance, scalability, availability, investment protection, reliability or TCO?

Serial technology in the datacenter



Interface Futures



HP mix in the enterprise by interface

separate interface from device

serial ATA

minimum feature set

extremely cost sensitivity

strategic business linkage to SAS

positioning: **non-mission critical**

target: entry level

applications: low cost nearline tape replacement, repositories



P-ATA
(1 % mix)



S-ATA
(erodes P-ATA & incremental)



P-SCSI
(87 % mix)



SAS
(replacing P-SCSI)

FC
(11 % mix)

serial attached SCSI

evolution of SCSI with enterprise features

universally compatible with S-ATA

positioning: **24x7 operation**

target: DAS, NAS

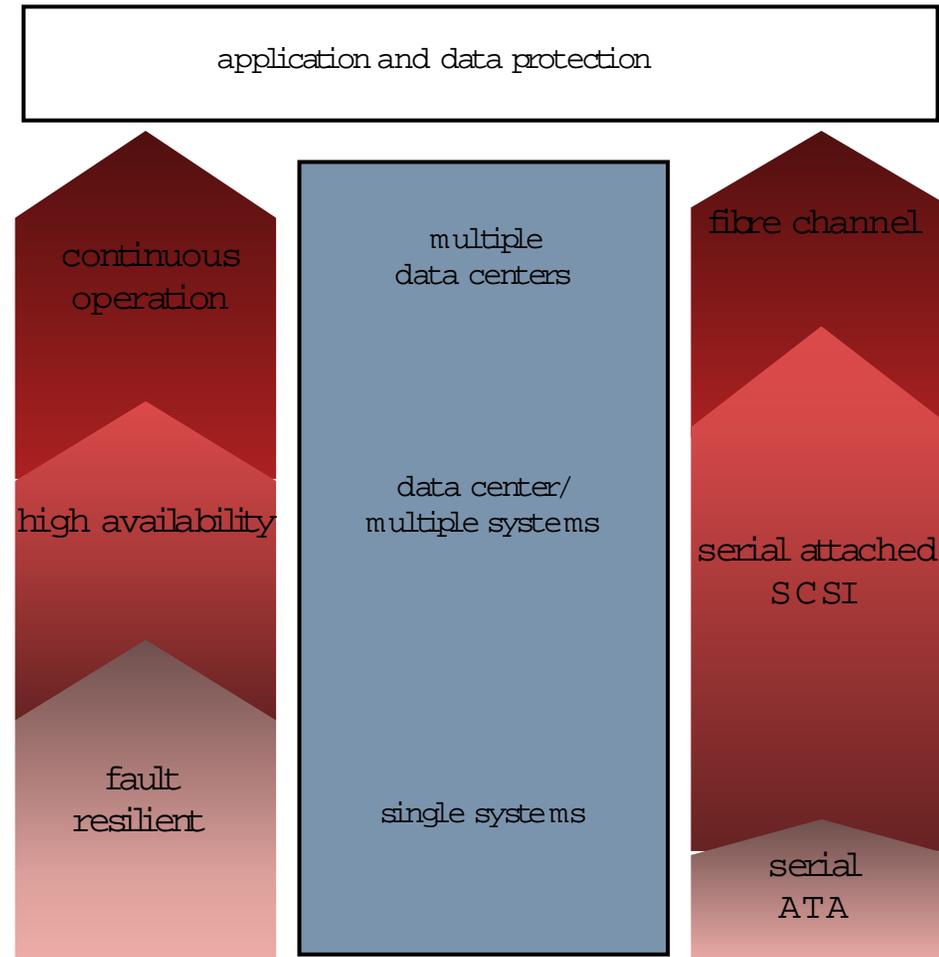
Volume: 87 % of enterprise HDDs

Units: increasing over time with SFF

Continuous, Secure Operations

business continuity in the enterprise

- tier one infrastructure (FC or SAS)
 1. intelligent fault resilience
 2. built-in security and proactive self-management
 3. application and data integrity
 4. optimized for planned and unplanned downtime
- tier two infrastructure (SAS)
 1. Core infrastructure for servers and storage
 2. Response time is important in doing business
 3. Uptime is critical
- tier three infrastructure (SATA)
 1. Low cost for non-mission critical
 2. Target bulk/reference/nearline
 3. Rebuilds won't hurt user response requirements
 4. Offset or stage to tape for archiving



serial comparative chart

SATA	Key Decision Factors	SAS
	<ul style="list-style-type: none"> ■ <u>Highest performance capability</u> <ul style="list-style-type: none"> - Full duplex, higher bandwidth, port aggr.. extensive command queuing, rich command feature set 	✓
✓	<ul style="list-style-type: none"> ■ <u>Lowest cost per gigabyte</u> <ul style="list-style-type: none"> - Low acquisition cost 	
✓	<ul style="list-style-type: none"> ■ <u>Best performance for transactional data</u> <ul style="list-style-type: none"> - High frequency immediate random-access type data; Database, on-line purchase, bank transactions, CRM etc. 	✓
✓	<ul style="list-style-type: none"> ■ <u>Most cost effective for reference and sequential-type data</u> <ul style="list-style-type: none"> - Low frequency access, streaming and sequential data; File-sharing/Email/Web/Nearline/Backup and archival data 	
✓	<ul style="list-style-type: none"> ■ <u>Greater Scalability</u> <ul style="list-style-type: none"> - Greatest physical device addressing range; Connection to out-of-box storage, long (8-meter) cable length 	✓
✓	<ul style="list-style-type: none"> ■ <u>First to volume disk interconnect</u> <ul style="list-style-type: none"> - Specification completed in 2001; Products started shipping in 2002 	
*	<ul style="list-style-type: none"> ■ <u>Highest reliability and availability</u> <ul style="list-style-type: none"> - More extensive error recovery techniques ; Multi-initiator (simultaneous access) + Dual-active port support 	✓
*	<ul style="list-style-type: none"> ■ <u>Simplest configuration setup</u> <ul style="list-style-type: none"> - Ubiquitous drivers in the desktop (enterprise still requires drivers for SATA) – No requirement for vendor drivers; Motherboard chipset integration 	*
✓	<ul style="list-style-type: none"> ■ <u>Greatest device flexibility</u> <ul style="list-style-type: none"> - Support for Serial Attached SCSI and Serial ATA devices 	✓
✓	<ul style="list-style-type: none"> ■ <u>Power management</u> <ul style="list-style-type: none"> - Saves power by automatically powering-down when not in use 	

serial ATA positioning

Serial-ATA is the best choice for desktop and non-mission critical storage where price is the primary selection factor

Serial-ATA 2.0 feature and performance enhancements build on the current trend of desktops, servers and networked storage embracing ATA as a viable disk interface

Serial-ATA is rapidly evolving to be the industry standard disk interface for the desktop, providing incremental opportunities for cost-sensitive server and networked storage segments

reality – no cost adder is acceptable since the same physical device is leveraged by the desktop market. This leverages a desktop hard drive.

serial attached SCSI positioning



serial attached SCSI is the best choice for mainstream and enterprise server and storage market segments

building on the foundation of the world's most popular enterprise interface, serial attached SCSI's enhancements, quality and reliability standards, and universal compatibility will meet the requirements of enterprise storage segments

serial attached SCSI is evolving from 20 years of SCSI infrastructure; enabling better performance, scalability and distance for mission critical enterprise solutions

reality – will continue the dominant DAS build-out for servers with enhanced capabilities of p-SCSI, and add incredible flexibility in system design

fibre channel positioning

fiber channel is *today's* best choice for the most demanding storage applications that require the highest performing, highest availability, most scalable networked disk storage systems

fibre channel drives pioneered serial disk technology in 1994 and continue to provide "Best in Class" features to support 24 by Forever storage systems and solutions

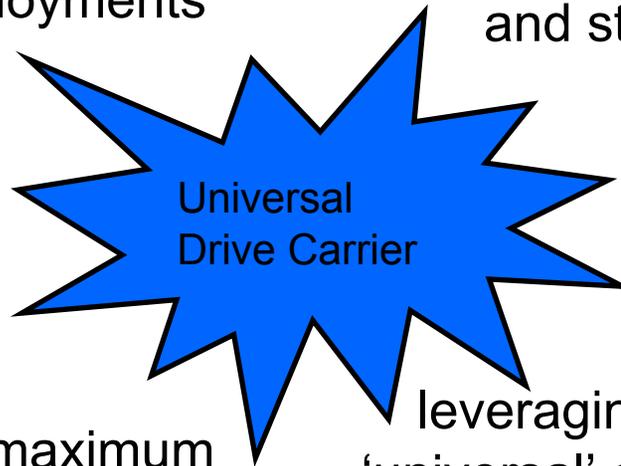
fibre channel drive architecture provides high performance (2Gb/sec, full duplexing, multi-drive loops) high levels of RAS (dual-porting and hot-pluggability) and high scalability (hundreds of drives)

reality – **will continue the dominant SAN build-out for storage though feature set is now similar to SAS**

serial positioning summary

customers requiring the best price advantage for desktop, server and storage deployments should choose serial ATA

customers requiring performance, reliability and software consistency for server and storage environments should choose serial attached SCSI

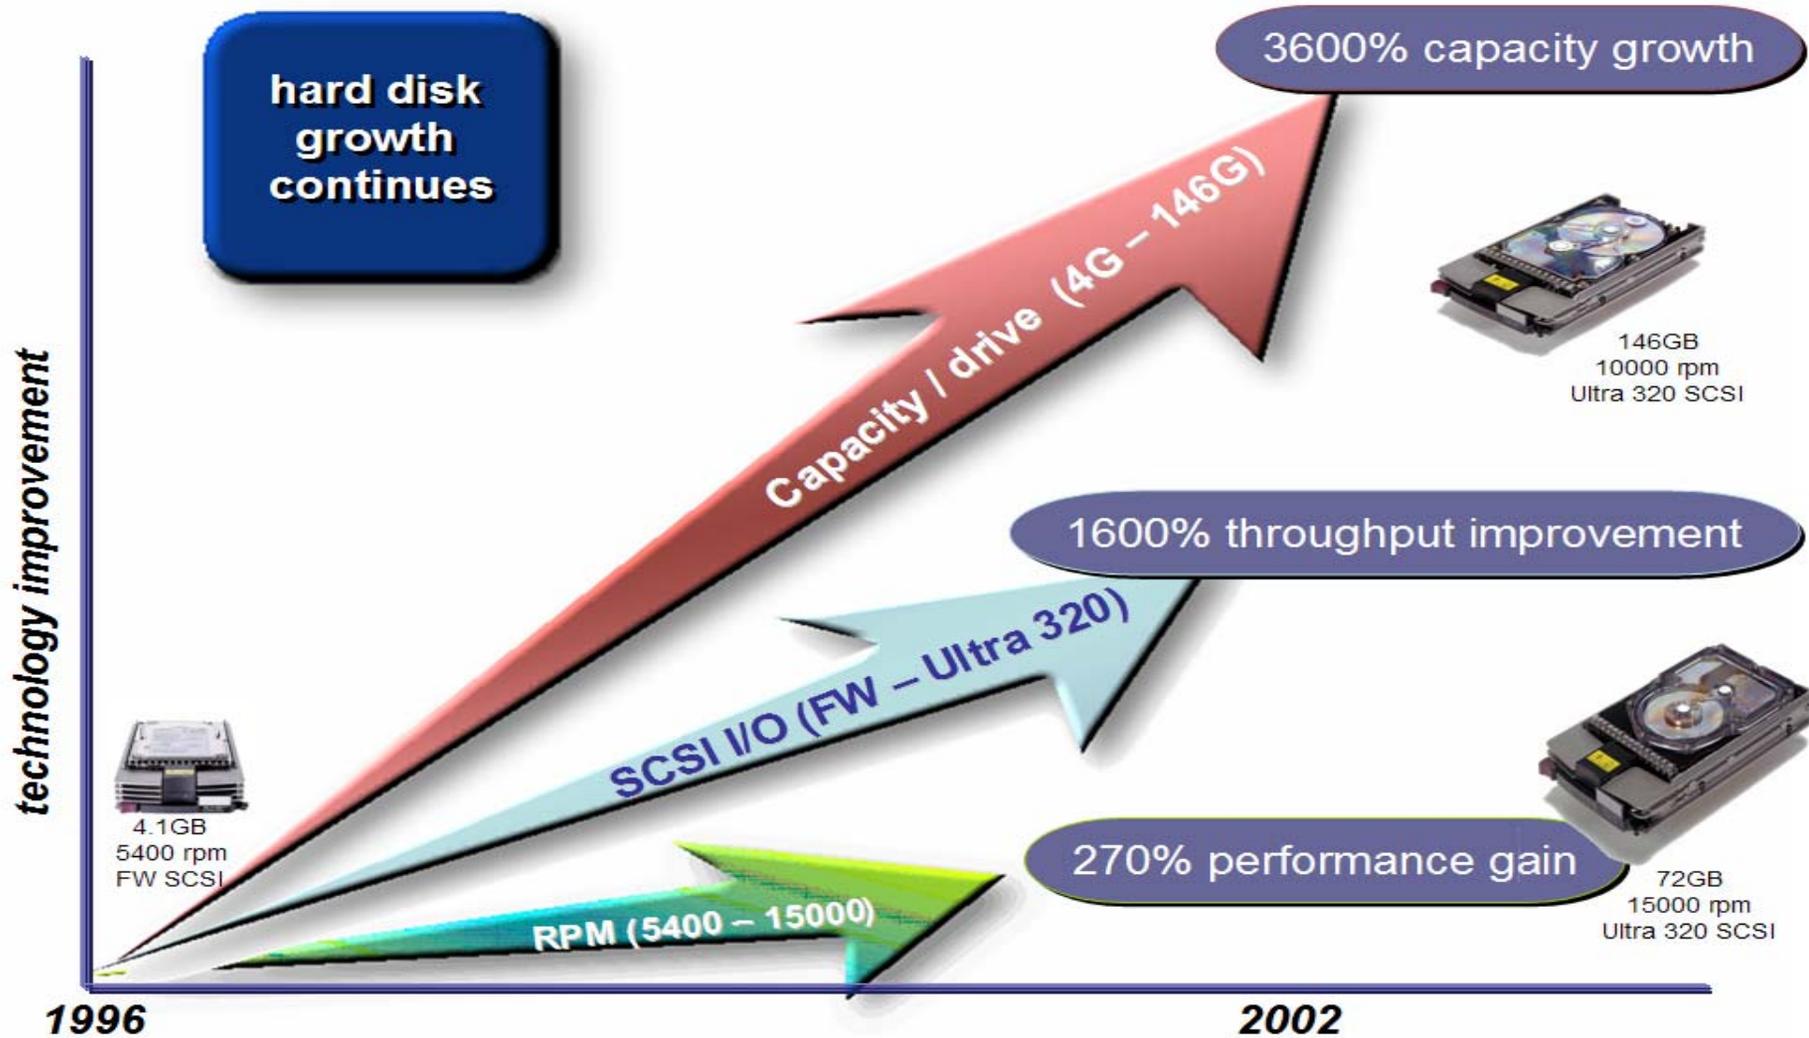


customers requiring maximum performance, reliability and scalability for SAN-based storage applications should chose Fibre Channel

leveraging industry standards in a 'universal' architecture, serial ATA and serial SCSI allow customers unprecedented flexibility in choosing the right class of storage that meets their requirements, while fibre channel continues to build out the SAN based infrastructure

Disk Drive Directions Form Factor





small form factor enterprise harddrives

choice of enterprise interface

- serial SCSI for server & storage deployments
- native Fibre Channel for SAN environments

easy to deploy

- more manageable capacity points per spindle
- familiar SCSI & FC protocols
- HP's pre-failure warranty
- optimized for density

high performance

- most I/O per U
- optimized for data intense environments

enabling solution breakthroughs

- enterprise reliability in blade environments
- enterprise performance in dense server applications
- addresses thermal and power issues in server and storage enclosures, providing flexibility to system and blade designers (2p & 4p)



why small form factor

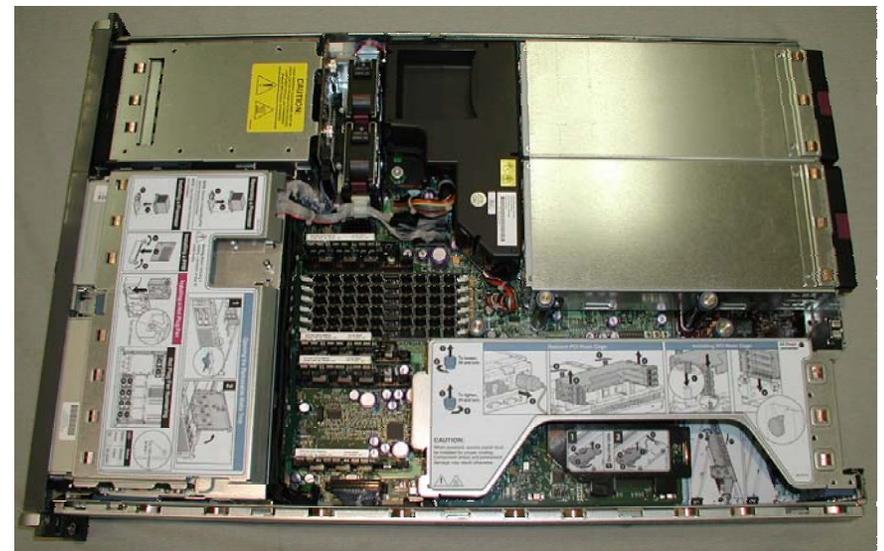
product challenges

- dense architectures
- thermal envelopes overloading
- power draws overloading
- airflow designs must improve
- new processors run hotter
- natural technology evolution



customer requirements

- excess capacity in the datacenter
- sweet-spot 2 capacities behind
- core requirements not addressed:
 - data protection – shrinking windows
 - data access – performance increase needed
 - hard drives fail too much today
- Customers want to be changed as little as possible – combine SFF & SAS for painless transition



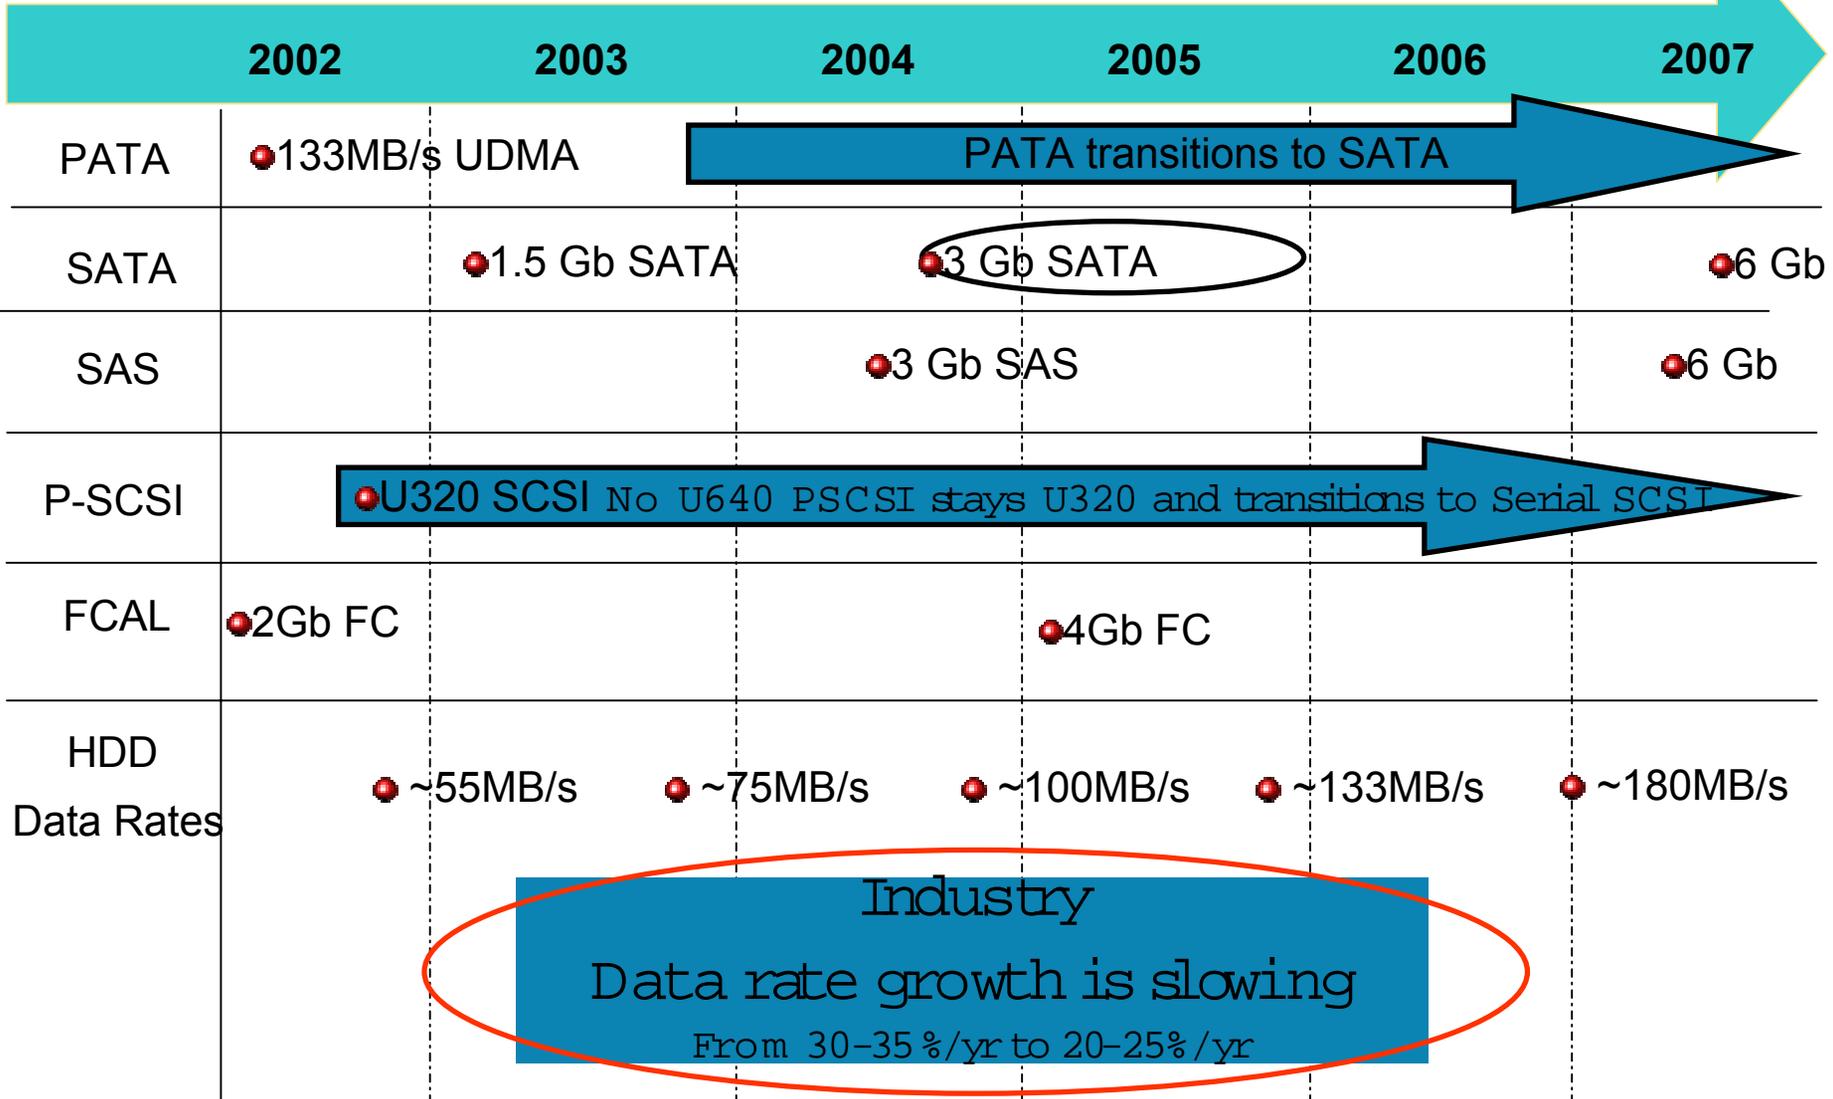
Disk Drive Directions Roadmap



Industry items to note

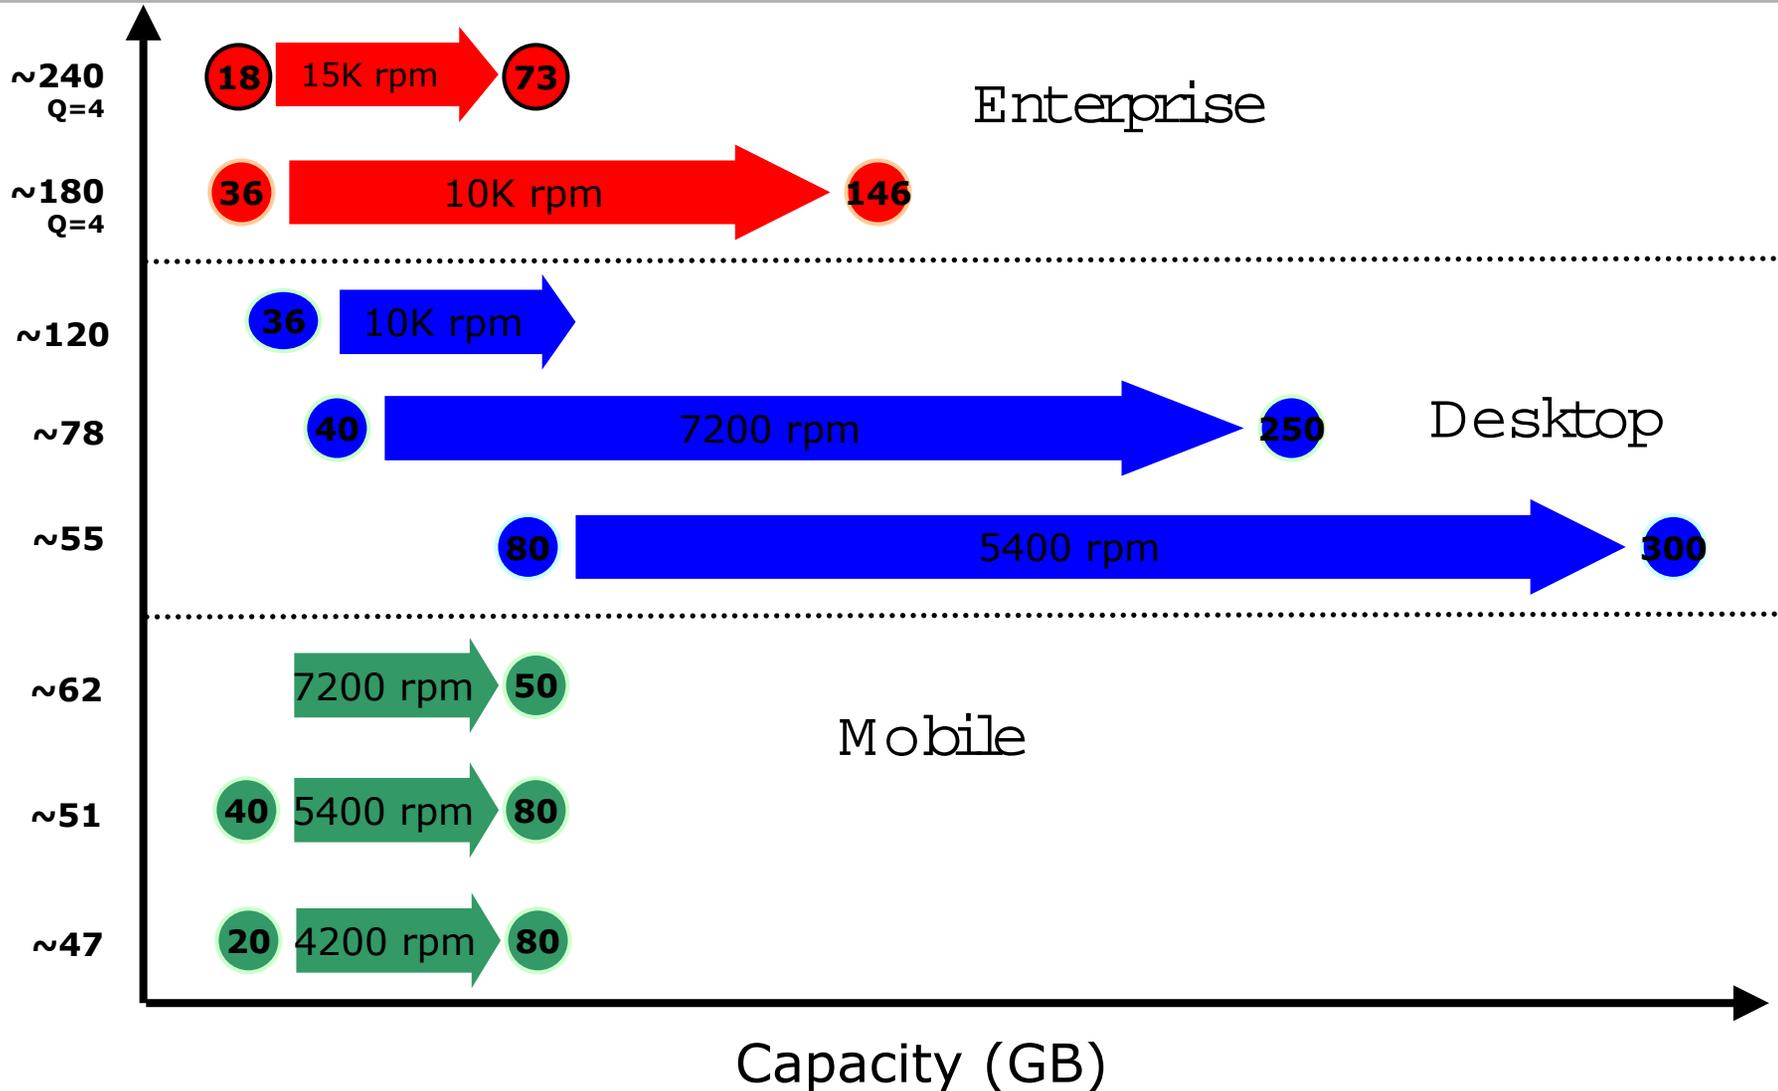
- A real density growth continues to slow. Capacity doubling cycles are increasing to 15 - 18+ months.
- Suppliers are struggling with 80G/p. Heads and media are the challenge.
- Greater than 80G/p schedules are very questionable. Expect slips. Industry targeting 120GB/p - 166GB/p in 1Q04
- 2005 looking like beginning revolutionary tech year. (micro-actuators, perpendicular recording).
- All enterprise HDD suppliers on board for SAS product 2H04 aggressive, 2005 with confidence.
- U640 SCSI is dead.
- SATA schedules slip.

HDD Interface roadmap



Performance differences

Performance (random IOPs)



iSCSI



HP IP Storage Strategic Direction



- HP will design and develop IP based network storage solutions that take advantage of customers' existing Ethernet knowledge, expertise, and infrastructure.
 - HP's IP based storage solutions will focus on the customer's need for agility and efficiency for increasing business effectiveness.
 - These IP based storage solutions will incorporate FCIP, iFCP, iSCSI, and RDMA technologies which integrate seamlessly and complementarily with existing Direct Attach Storage (DAS), Network Attached Storage (NAS), and Fibre Channel Storage Area Networks (FC SAN).
 - HP will work with industry standards groups and industry partners to develop and deploy adaptive storage infrastructures and solutions which best meet customer needs.

Today's Storage Environment

- Today's IT managers have storage solution choices which include FC SAN, NAS, and DAS.
 - SAN and NAS have become pervasive in IT environments and are expected to become the dominant data storage methodology in the future.
 - The Fibre Channel SAN has increasingly come to represent enterprise storage. SANs use FC infrastructure to create robust, consolidated, and manageable environments. 2Gbit FC is being deployed in data center environments.
 - NAS uses Ethernet fabric with the most recent 1Gbit infrastructure being deployed in most IT environments. Until now, NAS has concentrated on moving data in the form of files, using TCP/IP.
 - DAS, including both online and nearline storage, has been dominated by parallel SCSI. Customers have had to manage greater numbers of individual instances of DAS storage with associated activities of data protection through backup, replication, and disaster recovery solutions.

The Changing Storage Environment

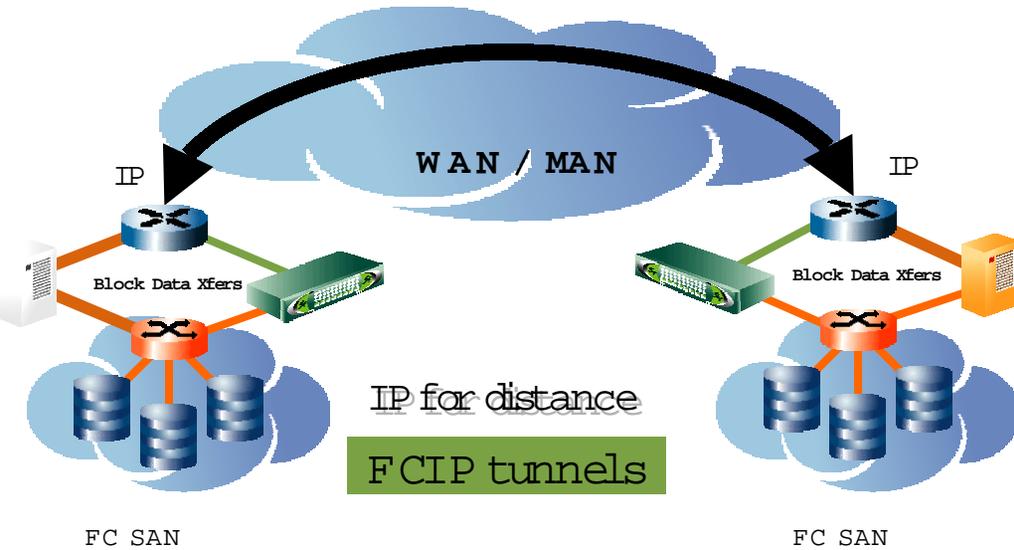


- Customers will expect increasing ease and flexibility, more cost-effective remote data access, and remote management of their IT environment from a single management interface (pane). In other words, customers are seeking adaptive storage infrastructures.
- With the advent of block data movement over IP, storage solutions will have the opportunity to leverage the ubiquitous Ethernet infrastructure in IT environments.
- Although FC infrastructure will continue to decline in cost and increase in interoperability and manageability, there are evolving opportunities for which IP based block storage solutions will become both advantageous and complementary to FC SANs.
- HP will continue to work with the IETF and other standards groups to ensure IP based storage standards are optimized for future needs of customers.
- Finally, HP will continue to invest in technology development and solution delivery that maintains industry leadership and provides effective, robust, and cost effective solutions for customers.

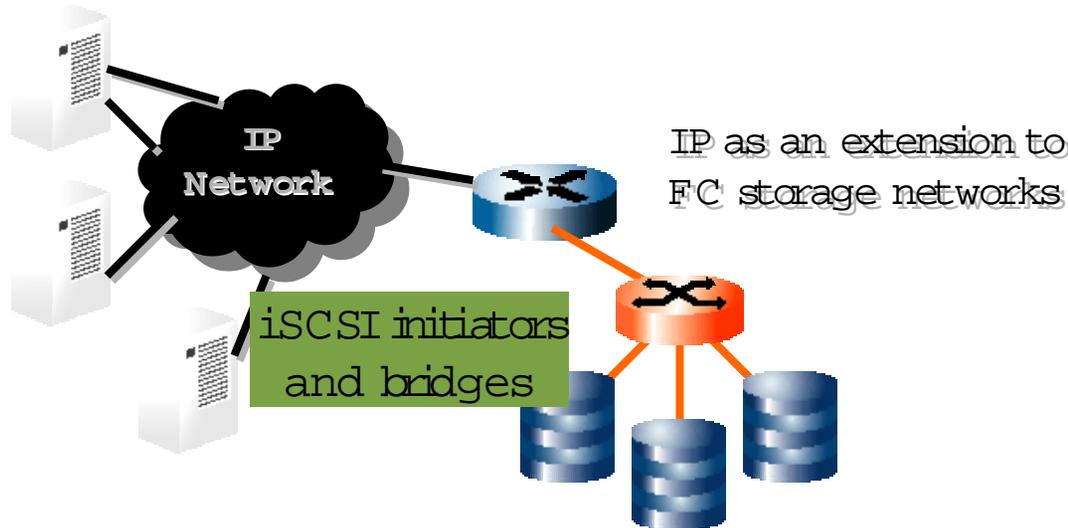
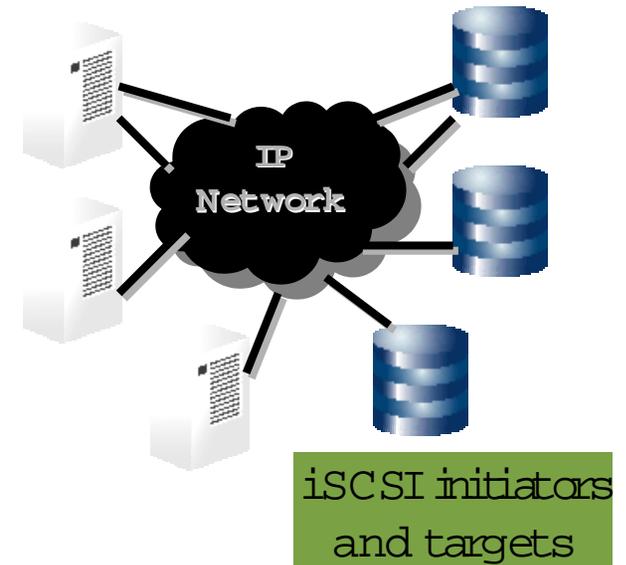
Where IP Storage is Going

- The first opportunities for customers to take advantage of IP based storage solutions are in the following areas:
 - Linking SAN islands which are geographically separated will become more critical for data assurance and accessibility across the enterprise.
 - Linking servers with the FC SAN without FC Host Bus Adaptors (HBAs) will become more common as iSCSI to FC bridges offer reliable and cost effective connectivity.
 - Native iSCSI storage targets resident on an Ethernet network will begin to offer basic functionality at even lower total costs than today's typical DAS environment.
- Future opportunities for IP based storage solutions will be enabled by the deployment of RDMA and ultimately the convergence of block and file data management in enterprise environments.

HP IP storage direction



IP as a storage network infrastructure



iSCSI specification approved

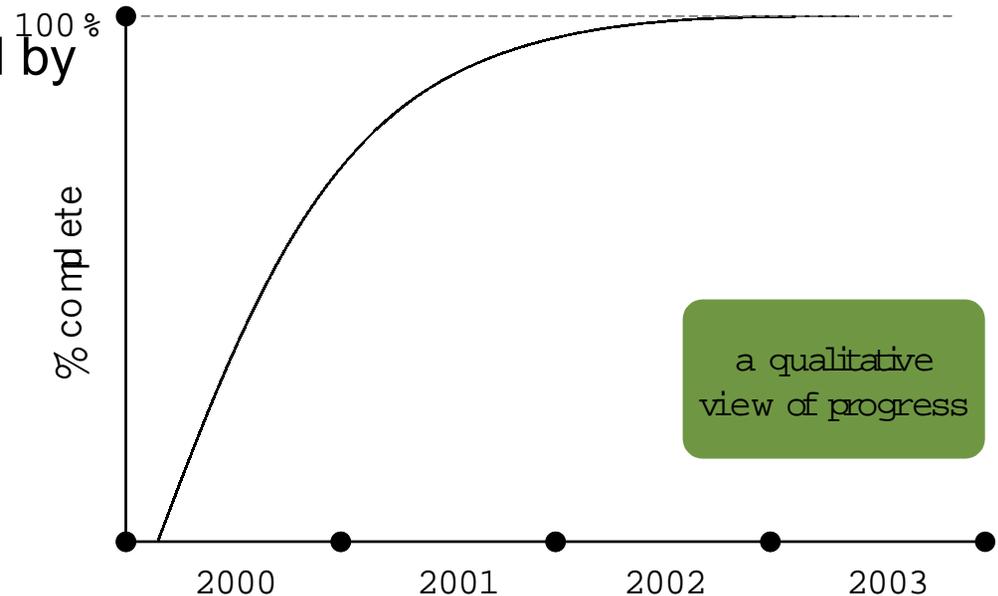
- Three years after being brought to the IETF
 - IETF approved iSCSI "1.0" in February 2003
 - RFC number to be assigned by the IETF

■ <http://www.ietf.org>

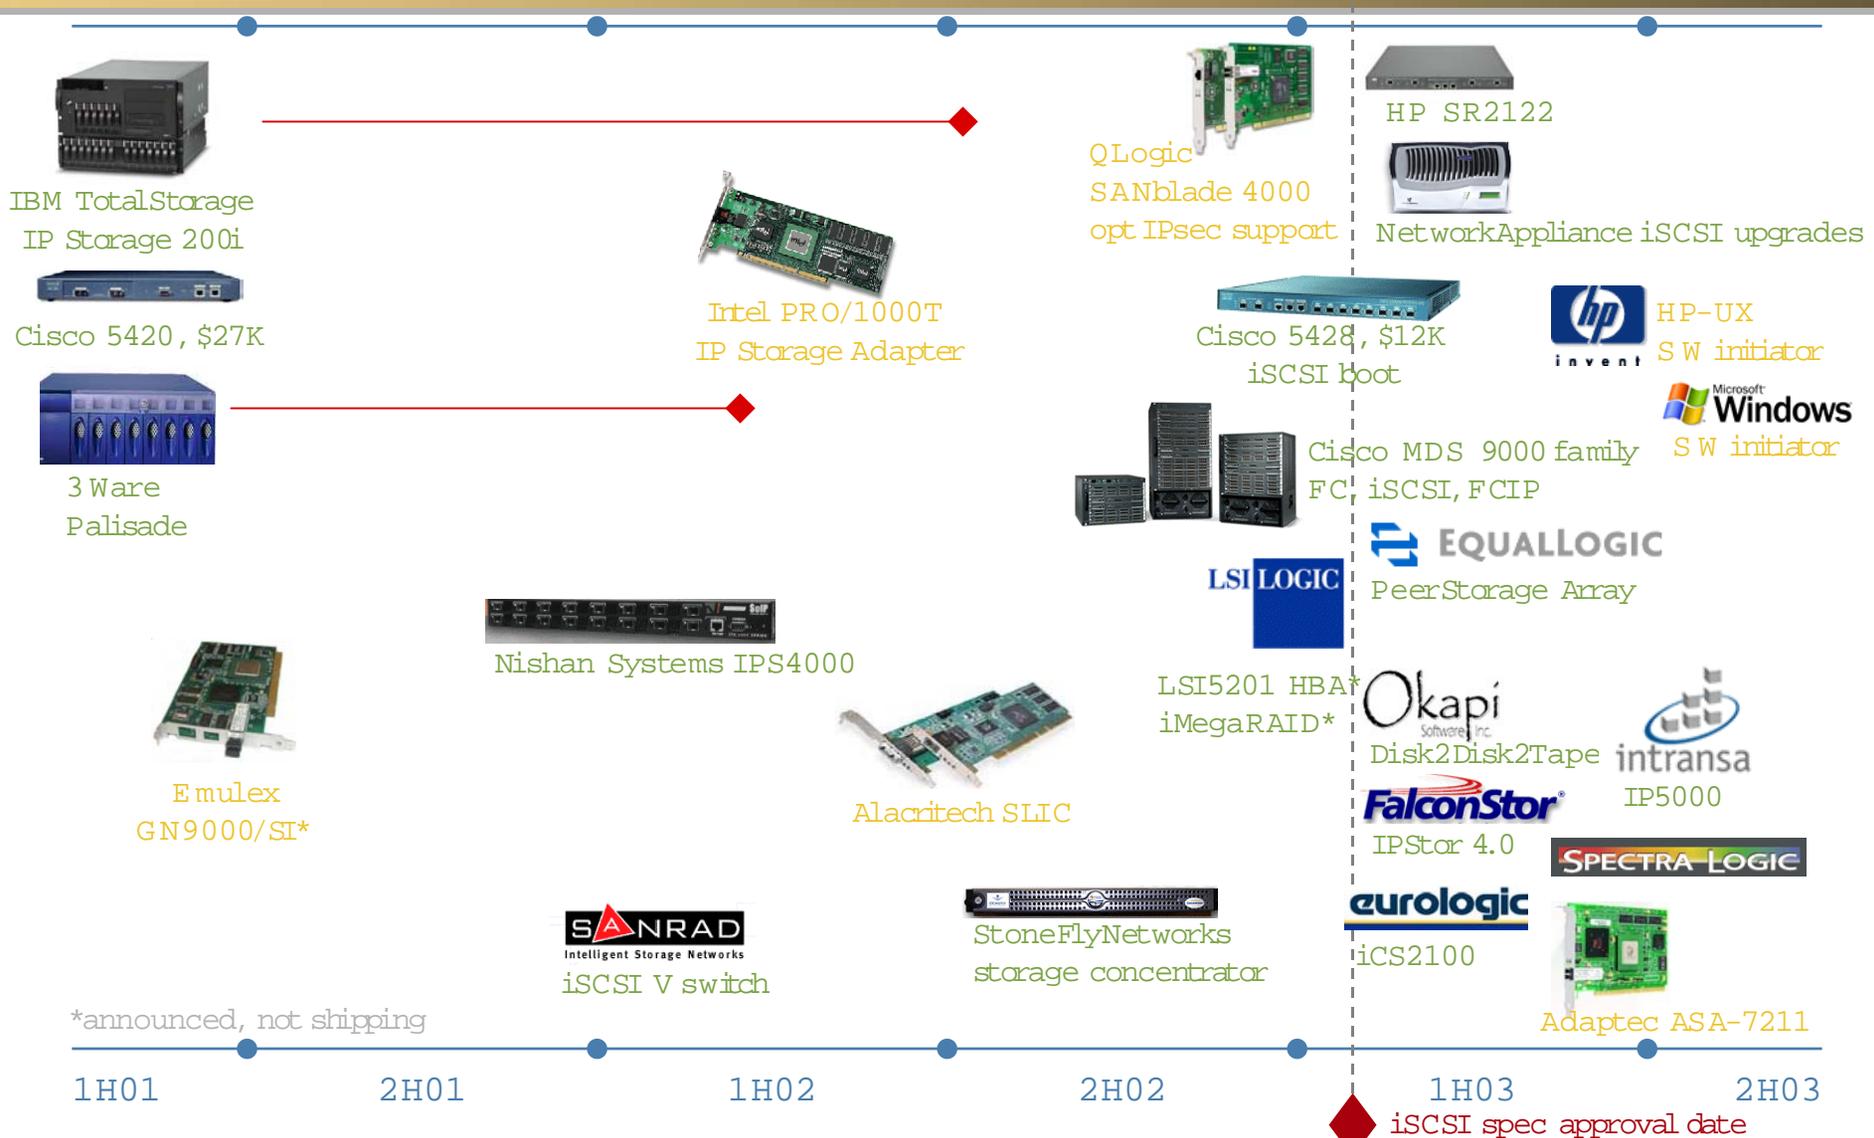
■ <http://www.ece.cmu.edu/~ips>

- RDMA Consortium is developing iSER (iSCSI Extensions for RDMA)

■ <http://www.rdmaconsortium.org>



iSCSI initiators, targets, routers, storage switches (publicly announced)



The local bus



Local Bus Options

- PCI-X 1.0 (today)
 - PCI-X 133mhz will double the performance of the host bus from 533 MB/sec to 1066 MB/sec PCI-X provides the bandwidth necessary for today's applications
- PCI-X 2.0 (tomorrow)
 - Doubles the performance again to 2GB / sec, using the same 64 bit, parallel, bus
- PCI Express (after PCI-X 2.0)
 - Transition from a parallel bus to a serial interconnect
 - For HP, this transition starts in 2006

Local bus summary

- hp server slot focus for next 2 – 3 years is on PCI-X 2.0, not PCI Express
 - A competitor may introduce Express 1.0 next year with the EVO form factor
 - HP position is that Express 1.0 will not have a sustainable ecosystem next year, nor will it have compelling customer value
 - Express 1.0 IS the right answer for desktops, however
- hp will transition to Express when it makes sense for our customers (target 2006)
 - This will be Express 2.0, likely with the REVO form factor

Q & A

