



# Session 3419: ATA and Serial Attached SCSI



Levi Norman  
Server Storage Marketing  
Hewlett-Packard

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







# Overview

- Enterprise and desktop drive design and features
- Parallel to Serial Transition
- Segmentation

# Anatomy of a disk drive




# Enterprise versus desktop devices



## Performance Mechanics

- Larger magnets improve seek times
- Performs well in multi-drive environment
- Stiff top cover minimizes susceptibility to Rotational Vibration




## Electronics

- Components handle higher data rates
- Maximize performance via SCSI firmware
- Dual processors

## Robust Head Stack Assembly

- Rigidity to stay on track
- More heads
- Performs well in multi-drive environment



## Smaller Diameter Disk


- Improves Seek
- Stable surface assuring reliability
- More platters

## Motor


- Higher RPM
- Tighter Runout (variation)

# Disk Drive Design Trade-offs


7200 RPM ATA/SATA HDD



10K RPM SCSI/SAS/FC HDD



15K RPM SCSI/SAS/FC HDD



## RPM

- The higher the RPM, the better the access times due to lower rotational latency
- The higher the RPM, the smaller the media size to minimize power consumption

## Media Size

- Larger media provides greater capacity
- Larger media requires a lower maximum RPM to maintain rotational stability
- Larger media means longer seek times due to greater seek distances and smaller magnets

## Actuator assembly

- The larger the actuator, the longer the seek time due to increased inertia

## Actuator magnet

- The larger the magnet, the better the seek performance

# Disk Parameter Comparisons

	Enterprise	Enterprise	Desktop
Form Factor	3.5"	2.5" (SFF)	3.5"
Interface	SCSI/FC/SAS	SCSI/FC/SAS	PATA/SATA
Spin RPM	15K / 10K	10K	7.2K
Capacity	36 – 300GB	36 – 73GB	40 – 250GB
Seek time	3.5ms / 4.7ms	<4.5ms	9.5ms
Random I/O	~300 / ~220	~240	~70
RVI rads^2	21	21	5
MTBF	1.5M/hr 24x7 100% duty	1.5M/hr 24x7 100% duty	600K/hr 8hr/d 20% duty
Error rate	1 in 10 <sup>16</sup>	1 in 10 <sup>16</sup>	1 in 10 <sup>15</sup>
Warranty	3 years	3 years	1 year

Have richard verify if desktop capacity needs to increase to 500GB

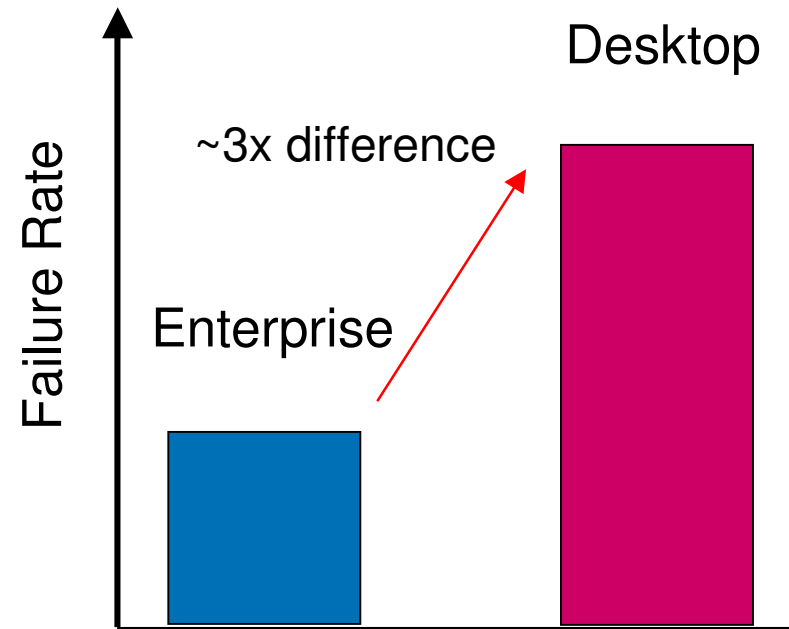
# Hard Drives Features

Enterprise	Desktop
High IOP's	Low cost
High Availability	Capacity
Multi-user	Single User
Performance	Low Acoustics
Scalability	
High Temperature	
Manageability	
24x7 reliability	8-10hr/day reliability
100% duty cycle	20% duty cycle

# Disk Drive Reliability

There is a significant difference in how MTBF is specified with respect to temperature, usage, and work load between **desktop** and **enterprise** disk drives

- Enterprise drives MTBF is based on 24 x 7 and continuous operation
- Desktop and Mobile are based on 8 hour per day use and low duty cycles



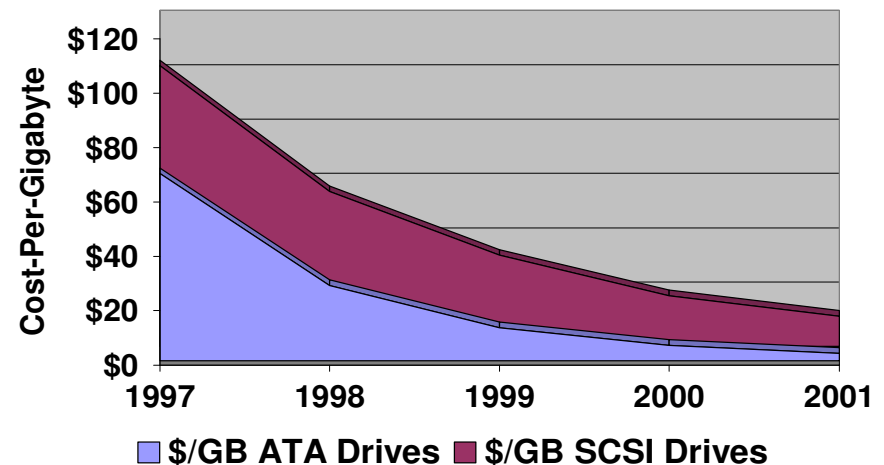
If used 8760 hours  
100% duty cycle




# Summary of disk differences

- ATA drives are designed to meet the needs of PC and consumer applications
  - These apps demand the lowest cost per gigabyte
  - Best \$/GB trading off performance and reliability
- SCSI drives are designed to meet the needs of mainstream servers, workstations, and RAID storage applications
  - These apps require the lowest \$ per IOPS and high availability
  - Best \$/IOP, best multi user performance, with highest reliability

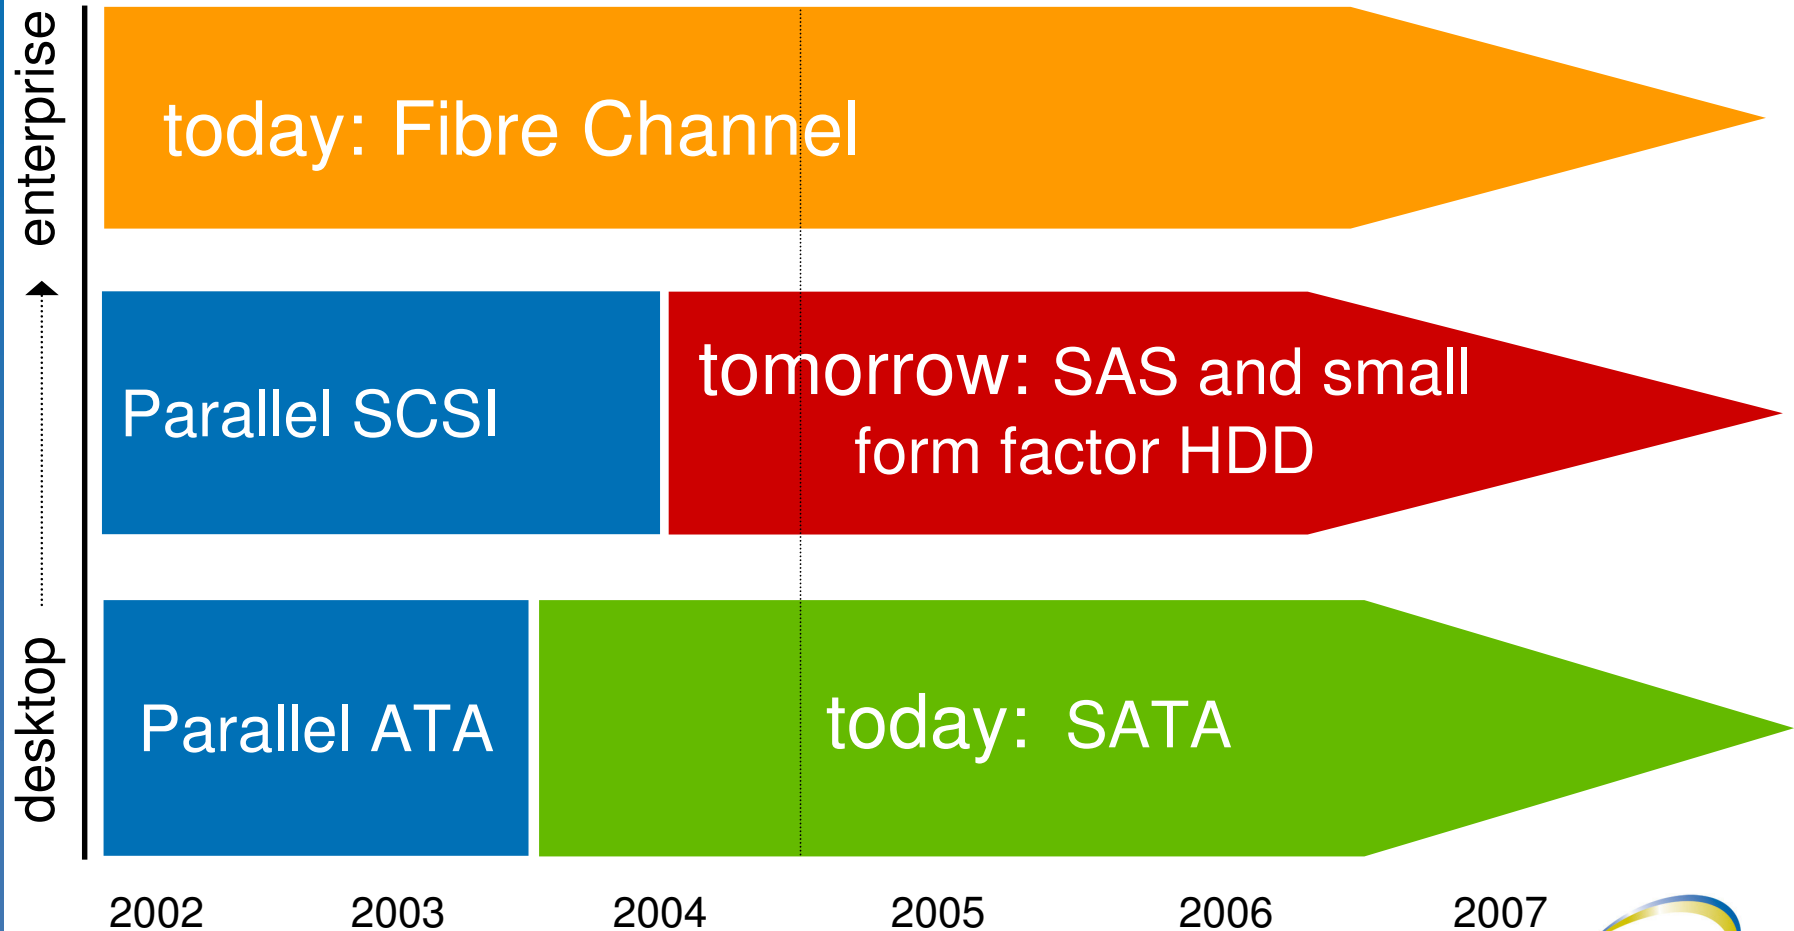
**ATA Disks: Best Cost-Per-Gigabyte**



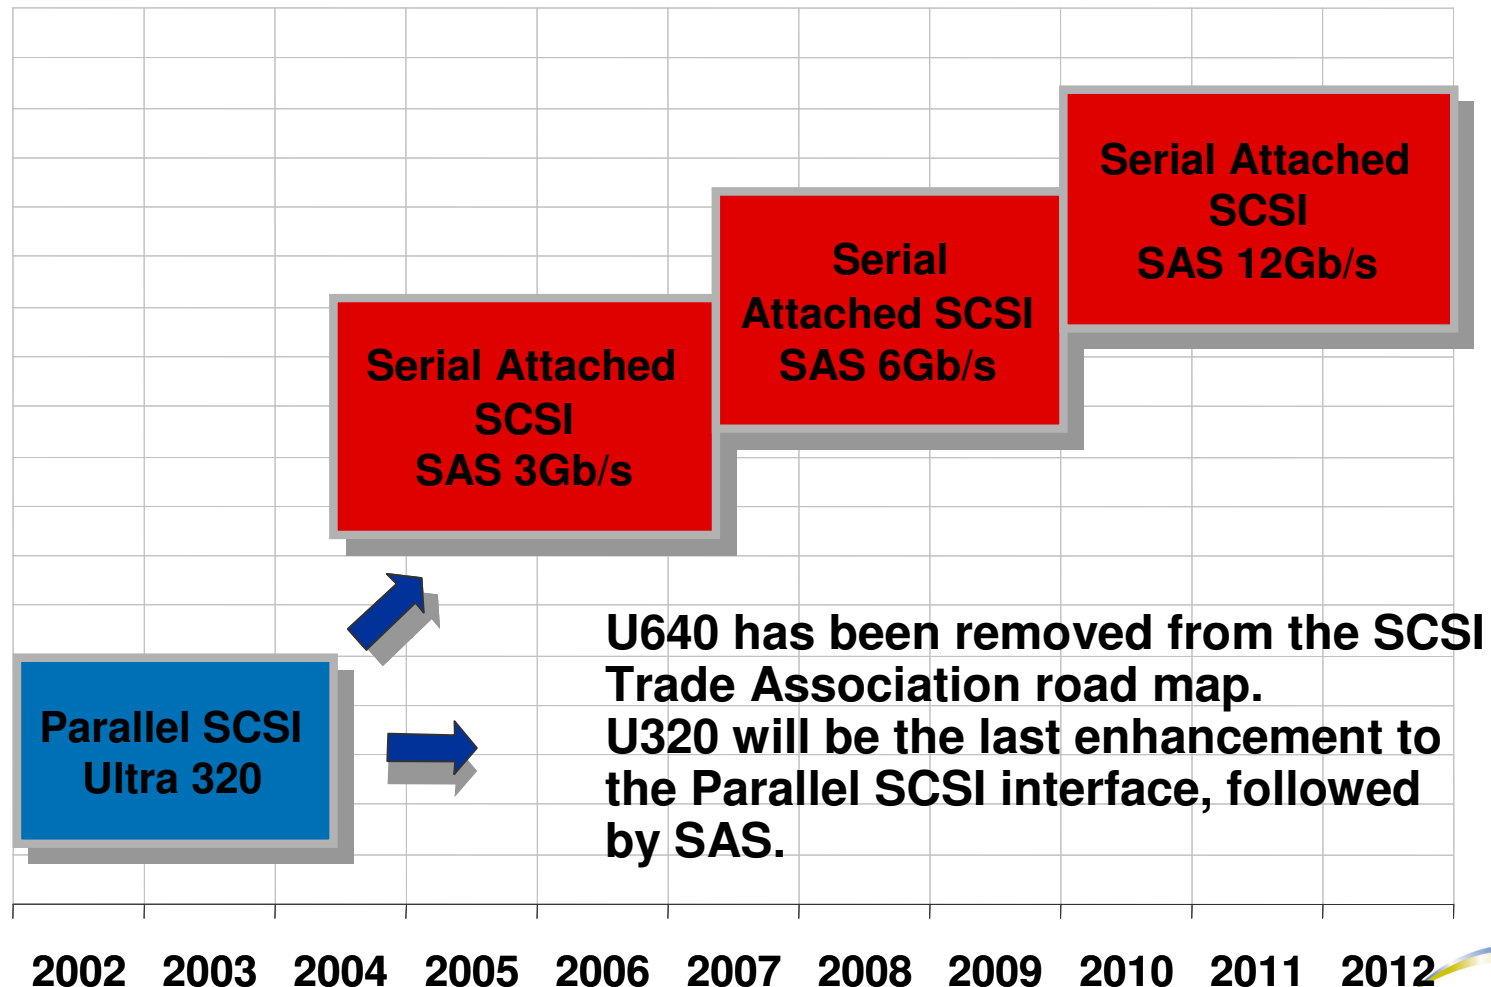
**SCSI Disks: Best Price-For-Performance**



# Parallel to Serial Transition




# SCSI Trade Association Roadmap



# What is Serial?

- Replacement parallel ATA and parallel SCSI
  - 100% SW compatible “drop-in”
- “Star” topology (point-to-point, no hubs)
  - Each device gets full bandwidth
  - No bus arbitration/collision overhead
- Serial ATA (SATA) and Serial Attached SCSI (SAS) have attributes and features that extend its capabilities
  - Provides additional capabilities such as hot plug and 1st party DMA
  - SAS provides for using SAS and SATA in the same box



# Serial Attached SCSI (SAS)


## features and benefits

serial attached SCSI feature	serial attached SCSI benefit
thinner cable & fewer signals	improved chassis airflow & cooling resulting in easier cable routing for density faster HDD access times and improved IOPS
disk/backplane interoperability	customer flexibility and scalability simplifies IT support higher bandwidth per pin
leverages industry standards & current software stack	faster time to market and ease of transition improved interoperability performance
3 Gb/s (300MB/sec) performance	improved overall bandwidth growth roadmap to 6 Gb/s resulting in investment protection
smaller connectors	small form factor drives dual-ported hard disk drives
point to point topology	dedicated, scalable throughput resulting in improved IOPS support for 1096 devices
HDD failure indication	system uptime and reliability 50% improvement in MTBF to 1.5 million hours
cyclical redundancy checking	improved data integrity
LVD signaling	less cross talk and power

# Serial Technology Enables Choice


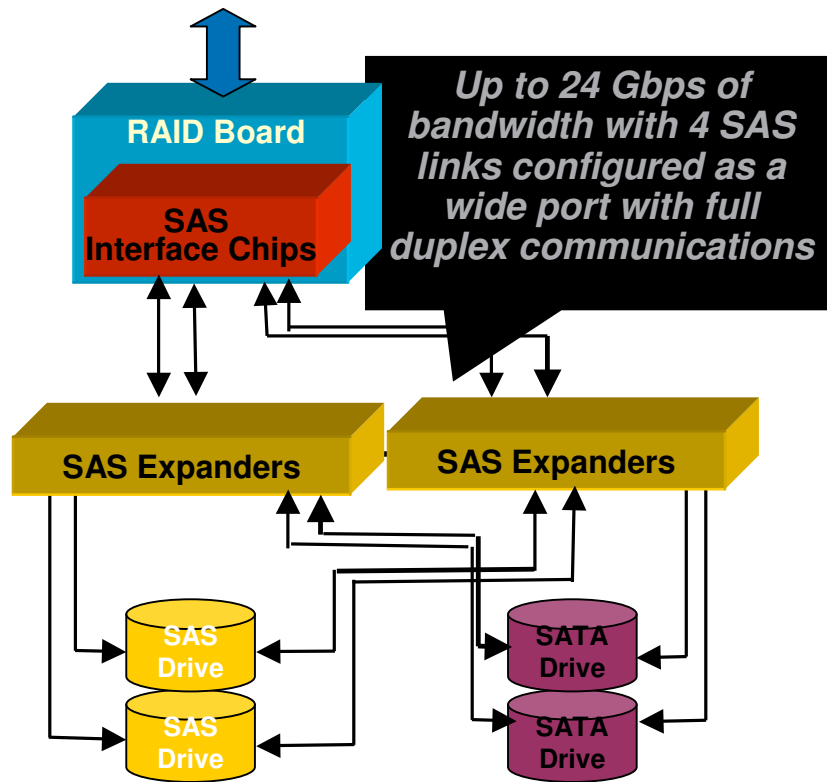


One backplane accommodates either SAS or SATA HDDs



- High-performance & highly-reliable serial attached SCSI disk drives can be used for mission critical and performance-oriented applications
- High-capacity SATA drives can be used for disk enhanced backup or non-mission critical reference data

# SAS Delivers Scalable Performance



- 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

# Summary Parallel to Serial Transition



## What's the Same

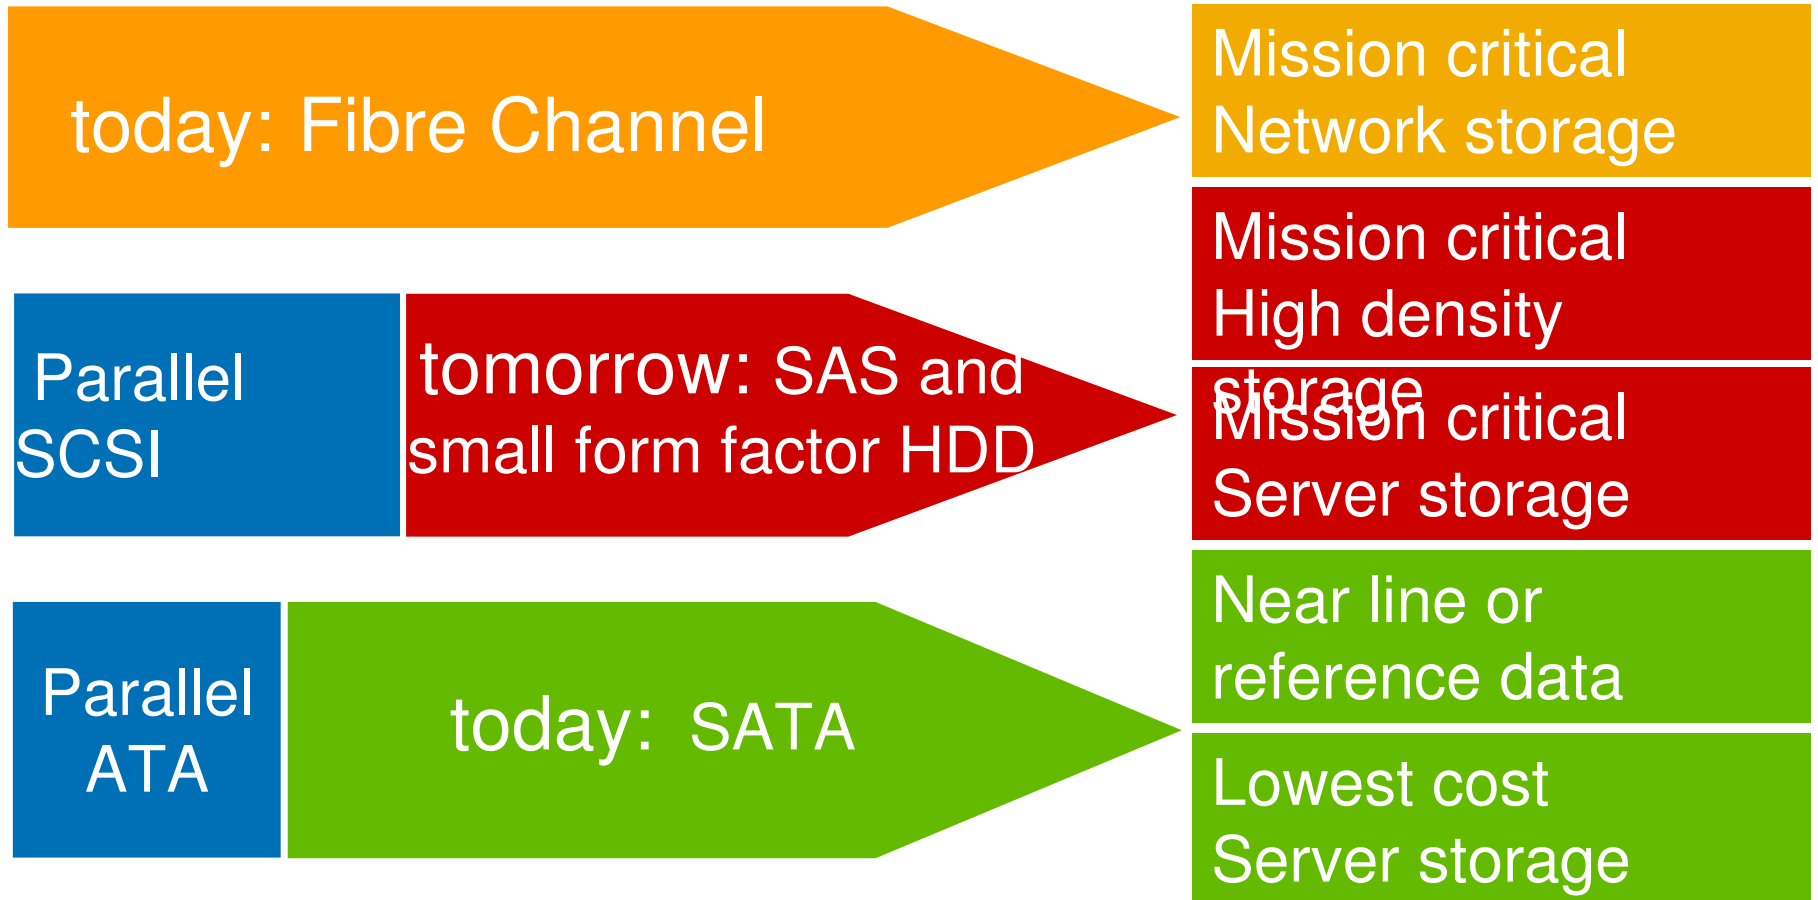
Builds off the SCSI protocol  
no new drivers  
same software  
SATA is entry level  
SCSI is enterprise level  
SATA is 8 to 10 hours, 5 days  
a week  
SCSI is 24x7x365

## What's New


Point-to-point  
dedicated bandwidth  
between controller  
and HDD  
Smaller, thinner cables for better  
airflow and cooling  
Greater flexibility  
mix SATA and SAS  
on same backplane  
Ability to connect more devices



# Datacenter storage segmenting




# HDD Market Segmentation

Networked Server High Density Storage	Mission Critical Storage
Critical To Customers	<ul style="list-style-type: none"> <li>• Transactional Performance</li> <li>• Guaranteed Availability</li> <li>• Mission Critical Reliability</li> </ul>
Driving Metrics	<ul style="list-style-type: none"> <li>• IOPS per GB (Performance/Cap.)</li> <li>• IOPS per U (Performance Density)</li> </ul>
Platforms For Products	<div> <div>           FC, SAS, SCSI            10K - 15K            RPM         </div> <div>  </div> <div>           Access time:            &lt; 5.7ms             Duty Cycle:            24 x 7x 365         </div> </div>

Access time – is it average seek 3.5 to 4.5 ms; access time 5.5 to 5.7

# HDD Market Segmentation

Near-line or Reference data Low Cost Server	Non-mission Critical Enterprise
Critical To Customers	<ul style="list-style-type: none"> <li>• Sequential Performance</li> <li>• High Capacity</li> <li>• 9-5, 5 days a week reliability</li> </ul>
Driving Metrics	<ul style="list-style-type: none"> <li>• \$ per GB (TB solutions)</li> <li>• \$ per GB (Small Capacities; Cost)</li> </ul>
Platforms For Products	<p data-bbox="961 1036 1713 1117">Near-line or Reference data &amp; Low Cost Server</p> <div data-bbox="961 1117 1184 1205"> <p>SATA, ATA</p> <p>7.2K RPM</p> </div> <div data-bbox="1314 1104 1432 1185">  </div> <div data-bbox="1453 1104 1696 1185"> <p>Access time: &lt; 13ms</p> </div> <div data-bbox="1453 1221 1747 1351"> <p>Duty Cycle: 8 to 10 hours, 5 days a week</p> </div>

# Technology Comparison

	SATA	SAS	Fibre Channel Arbitrated Loop
<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 2005
<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

# Drive Usage Summary

	SATA	SAS
Drive Strengths	Capacity	Reliability
	Lowest Cost Per GB	Best Price for Performance
Applications	Reference Data	Transactional Data
	Non-Mission Critical Applications (near-line storage, data tubs)	Mission Critical Applications (on-line transactions, banking)
	Entry level servers and bulk storage	Main stream servers and storage
Reliability	Desktop Class Architecture	Enterprise Class Architecture
	8 to 10 hours, 5 days a week 20% Duty Cycle	24x7x365 Duty Cycle 100%



# HP WORLD 2004

Solutions and Technology Conference & Expo

Co-produced by:

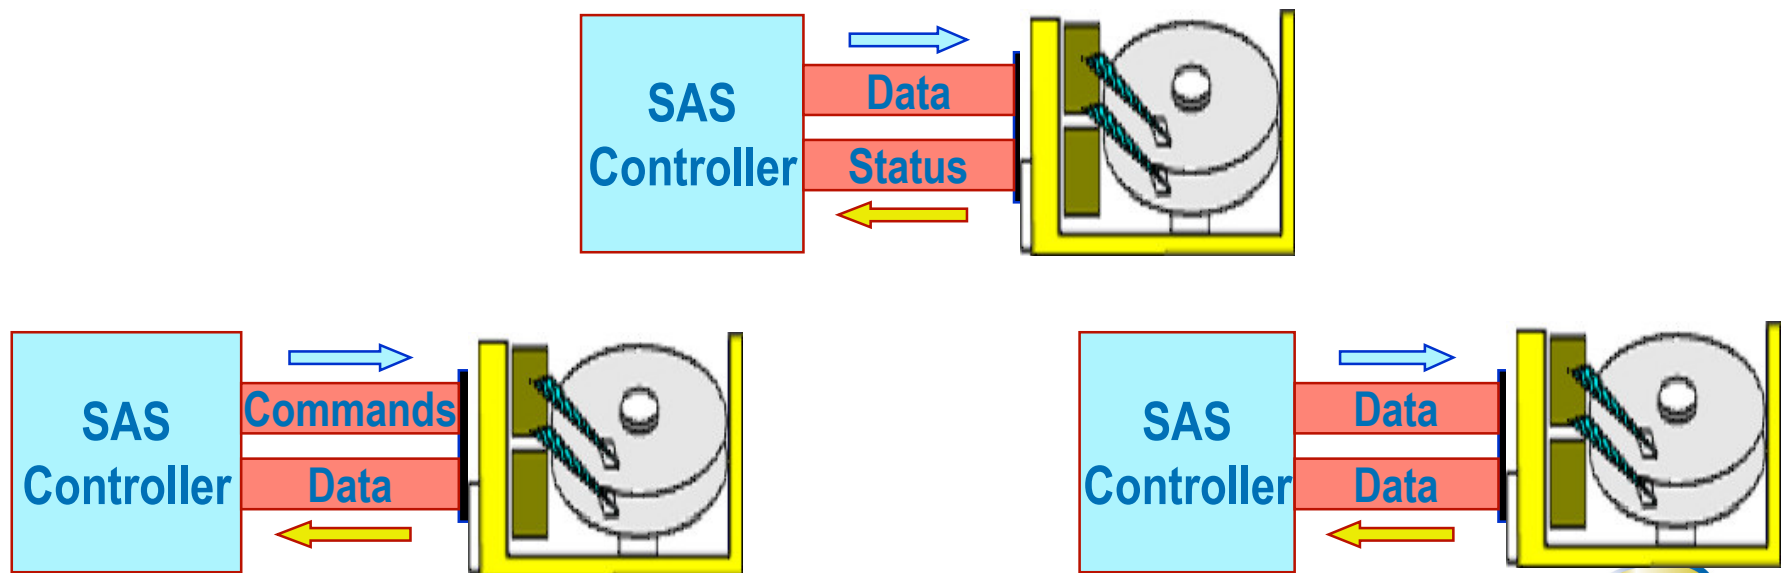


RECOMMENDED TRAINING VENUE FOR THE  
**HP Certified Professional**



# What is full duplex?


- The SAS interface allows for communication in both directions on a link simultaneously.
- SAS will ship full duplex first generation.





# What is dual port?

- The SAS interface allows for devices to have two limited full duplex links. (Limited to no more than two simultaneous data transfers.)
- Second generation SAS devices are targeted as dual ported.
- Controllers can have (n) number duplex links.



# What are multiple links?

- The SAS interface allows for ganging of multiple ports together to create “wide” links. Wide links provide increased bandwidth 1x, 2x, 3x, 4x, of 300MB/s!

