New SCSI transports (iSCSI, iSER, SAS, SATA): A comparative study

Alok Gupta
R & D Manager
Hewlett-Packard Company

HP WORLD 2003
Solutions and Technology Conference & Expo

Agenda

- Typical storage architecture
- Brief overview of new transports
  - iSCSI
  - iSER
  - SATA
  - SAS
- Relative benefits

Newer SCSI Transports

- Numerous newer SCSI Transports are emerging
  - iSCSI (internet SCSI)
  - iSER (iSCSI Extensions for RDMA)
  - SATA (Serial ATA, not really SCSI)
  - SAS (Serial Attached SCSI)
- Who are the Incumbents
  - pSCSI (Small Computer Systems Interface)
  - FC (Fibre Channel)
  - ATA (not really SCSI) …

SCSI Architecture

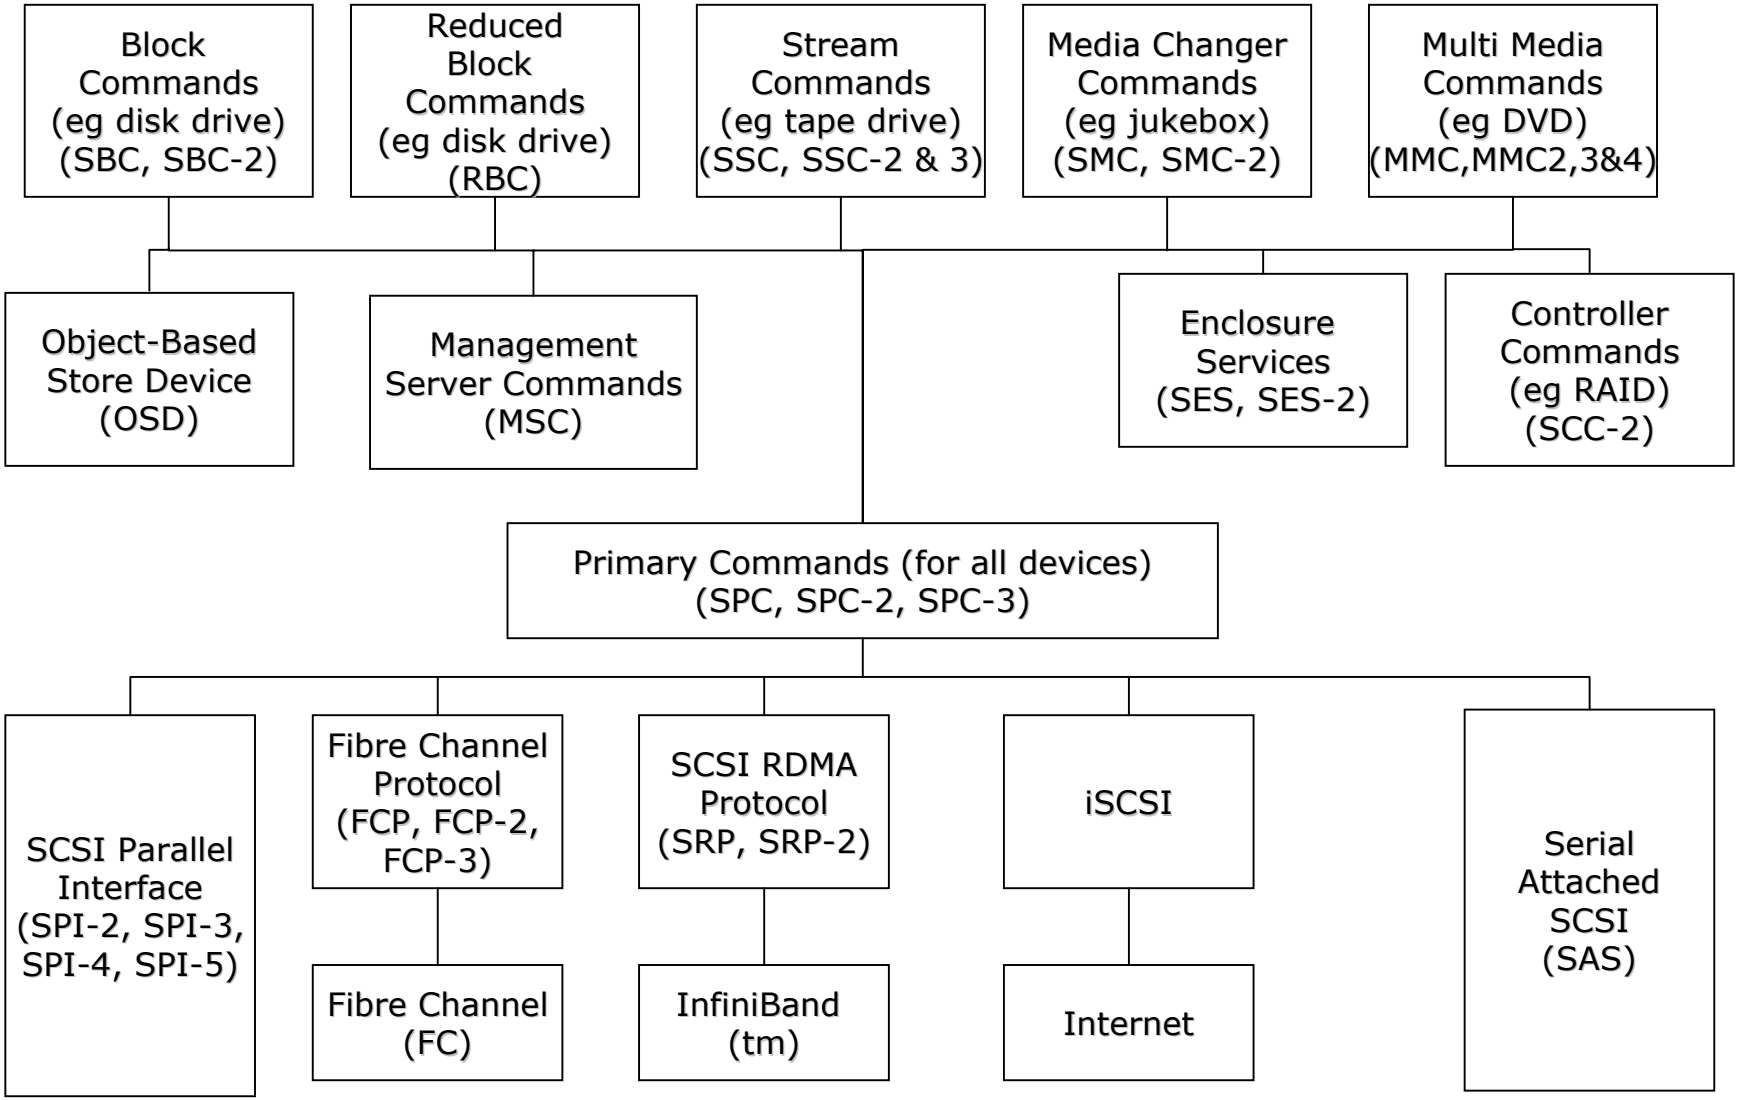

Typical Storage Architecture (Today)

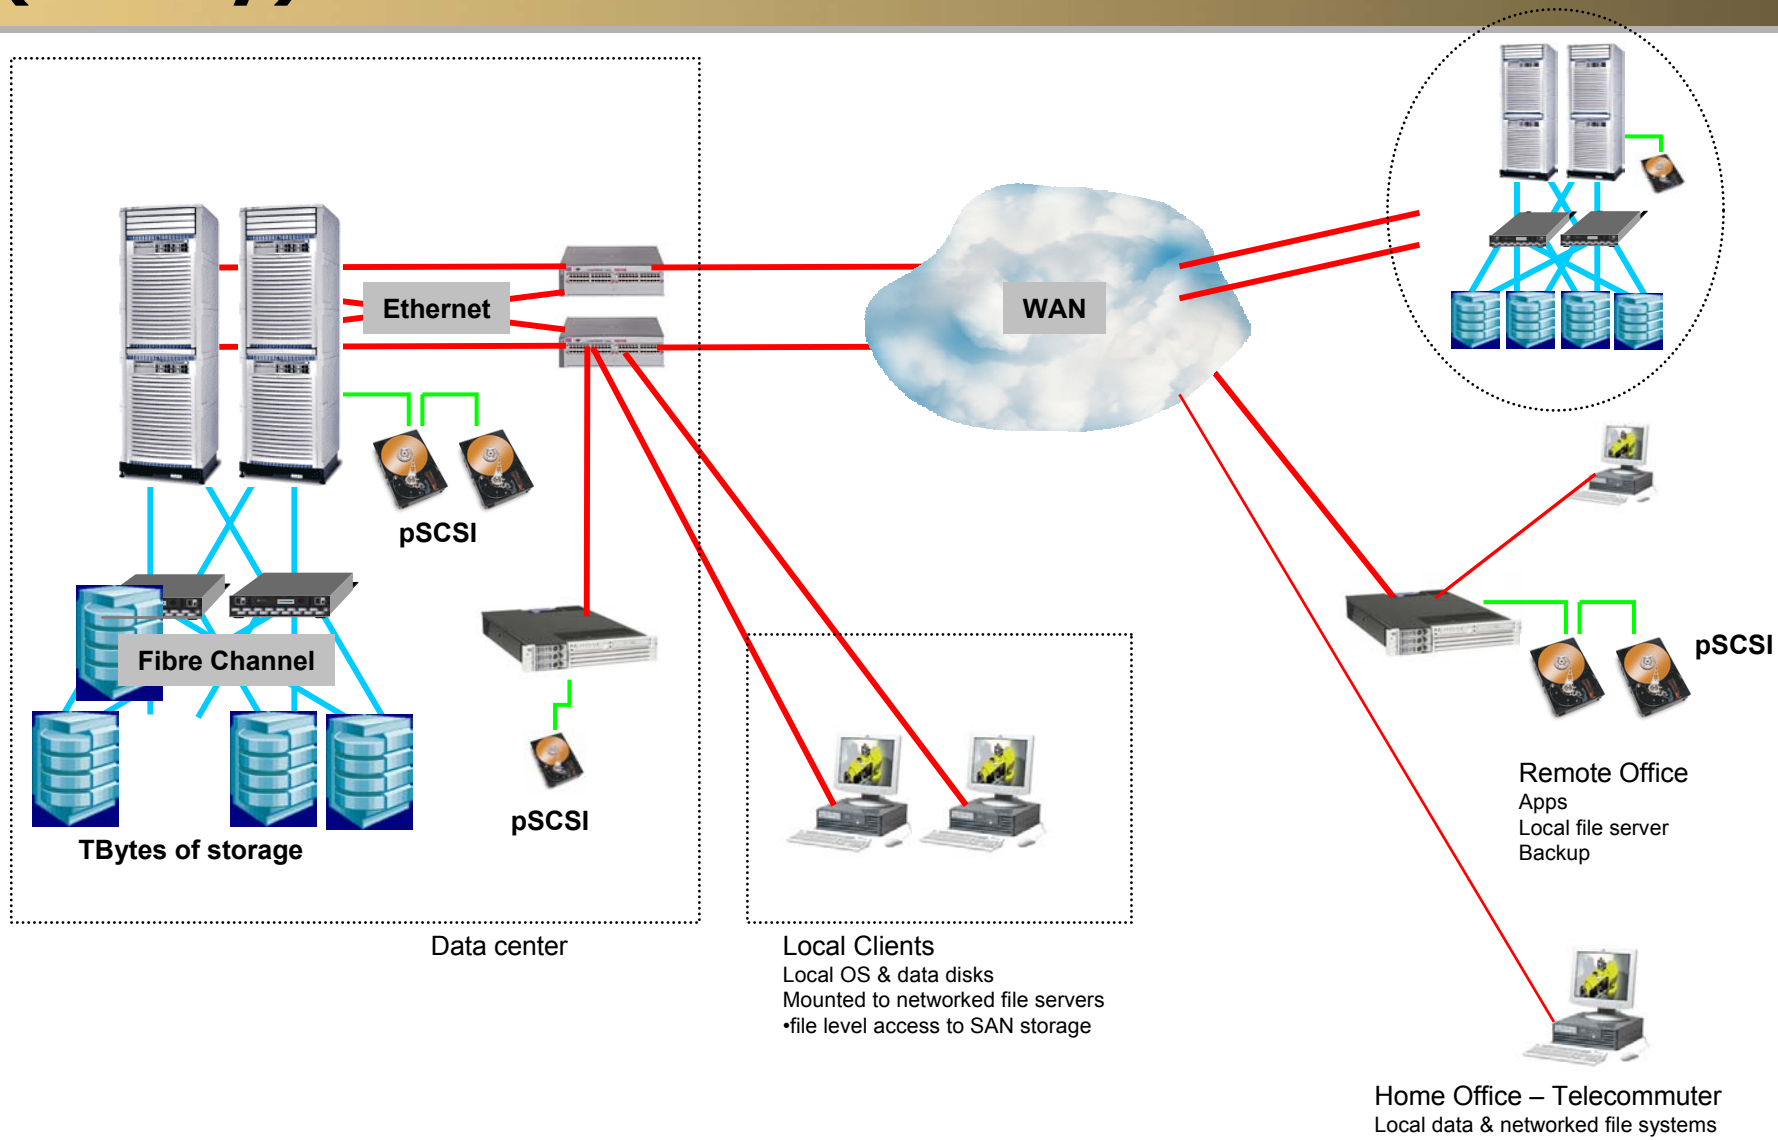

iSCSI Overview

- SCSI transport protocol that maps the SCSI family of protocols onto TCP/IP

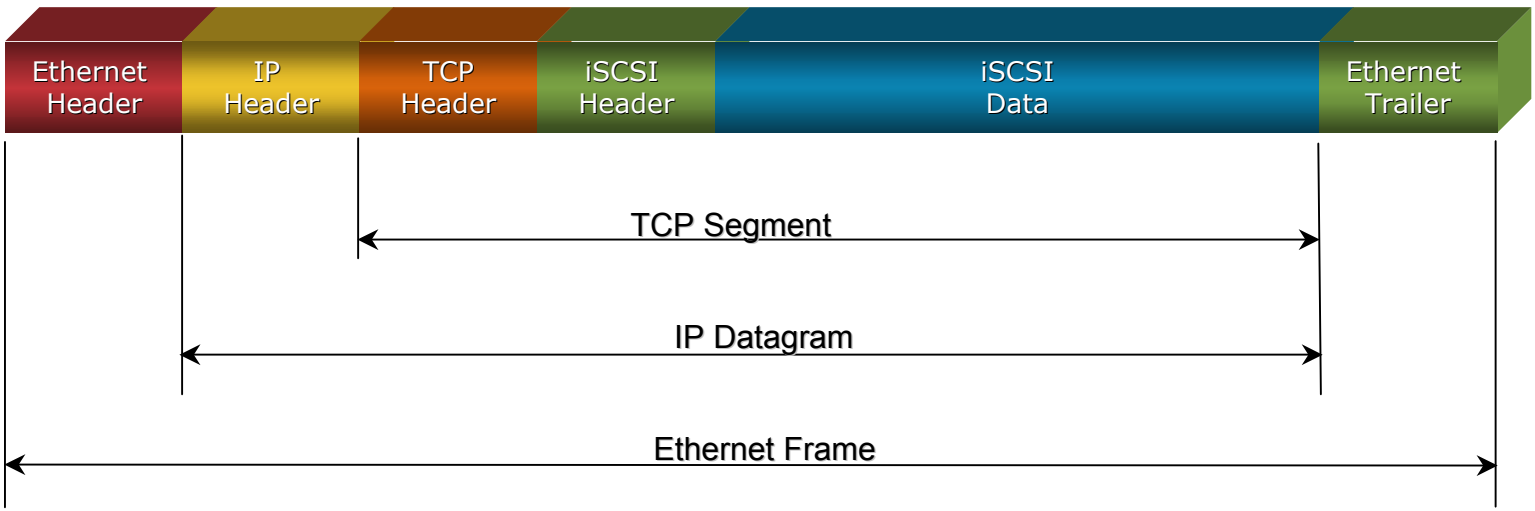

iSCSI Layer Overview

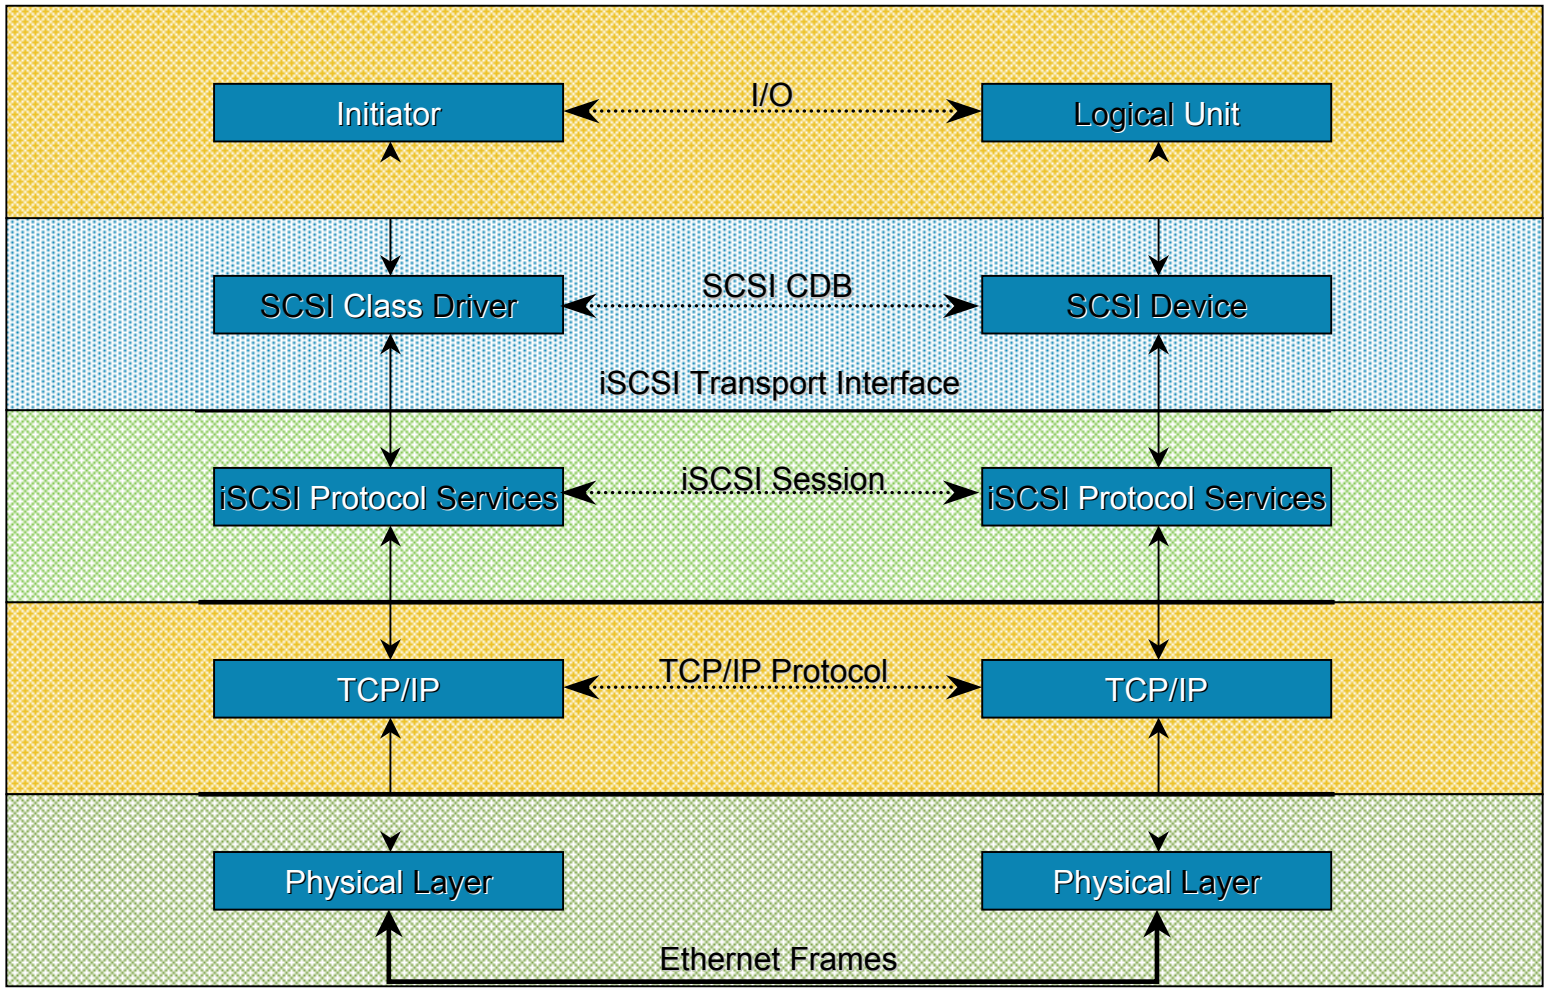

iSCSI and Fibre Channel

- iSCSI protocols can be exported fully into Host Bus Adapters
  - Can get added benefits when the same HBA is used for raw Ethernet and/or TCP/IP Offload Engine (TOE)
  - Can take advantage of networking debugging and administrative tools
- Discovery can be through static configuration, Service Location Protocol (SLP) or internet Storage Name Service (iSNS)
- iSCSI is a session based protocol

Type of iSCSI Targets

- iSCSI Target can be a gateway to Fibre Channel
- iSCSI Target can be a gateway to pSCSI
- iSCSI Target can be a storage array, tape drive or tape library with native interface to iSCSI
- iSCSI Target can be a virtualization device

iSCSI Value Proposition

- EtherEverywhere Vision (widely deployed and understood)
  - Convergence of mass storage and networking fabric
  - Brings in QoS, vLAN and Security
  - Networking specific load balancing and fail-over features
  - Provides seamless conversion from local to wide area networks
  - Ability to use mature networking tools for supportability
  - Piggyback advances in Ethernet link speed
  - Brings in economies of scale
- Provides block I/O unlike NAS

Market Segments

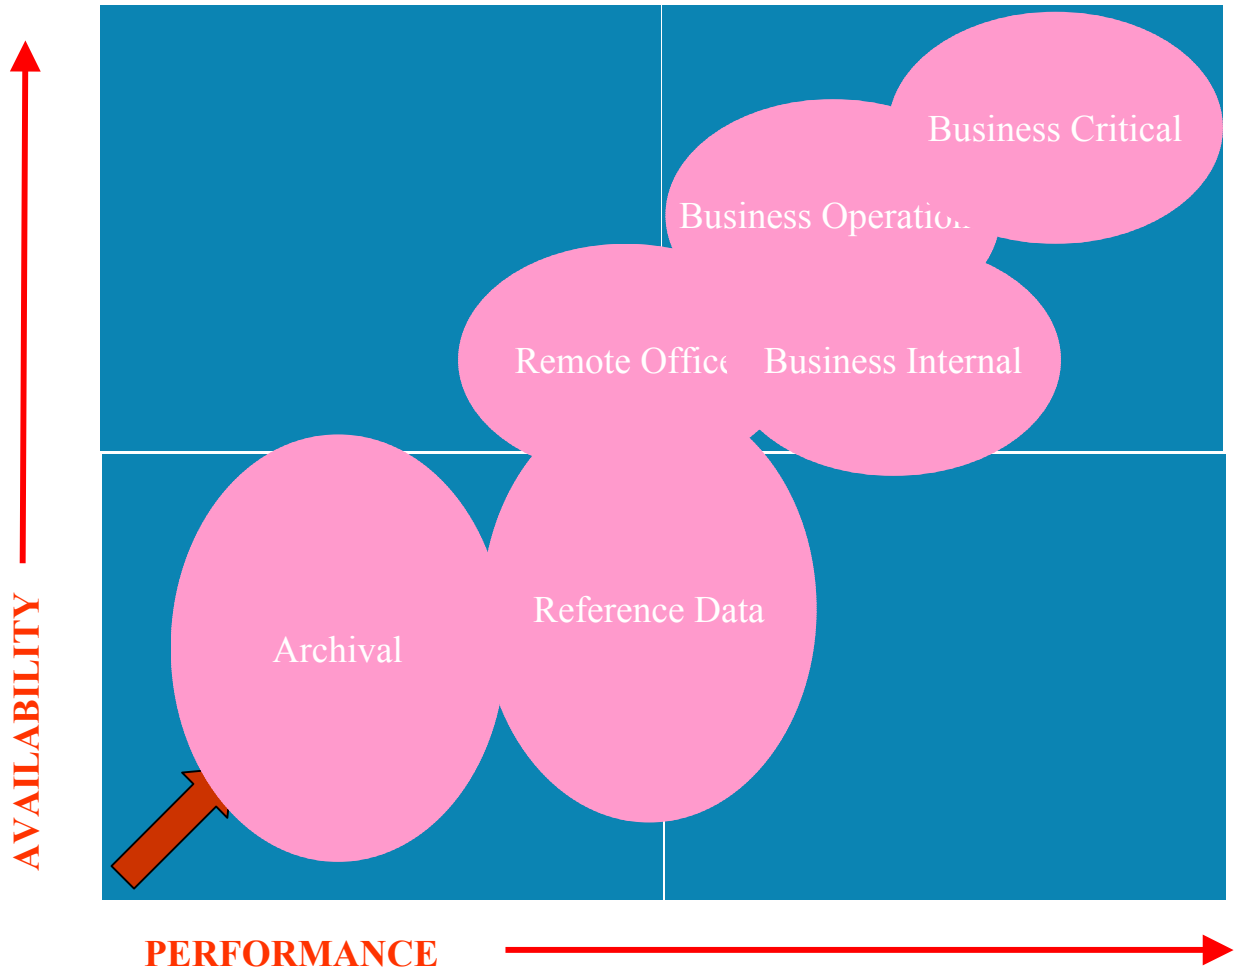

iSER Overview

- iSER (iSCSI Extensions for RDMA) provides RDMA capability to iSCSI by layering iSCSI on top of the RDMAP (Remote Direct Memory Access Protocol)
  - RDMAP provides read and write services directly to applications and enables data to be transferred directly into user buffers without intermediate data copies
- Designed to enable scaling at high speeds
- iSER mode is negotiated during iSCSI login
  - All iSCSI interactions thereafter use RDMAP messages
  - Target drives SCSI writes and reads (exception unsolicited data)
- No changes to iSCSI authentication, security and text mode negotiations

iSER Layering

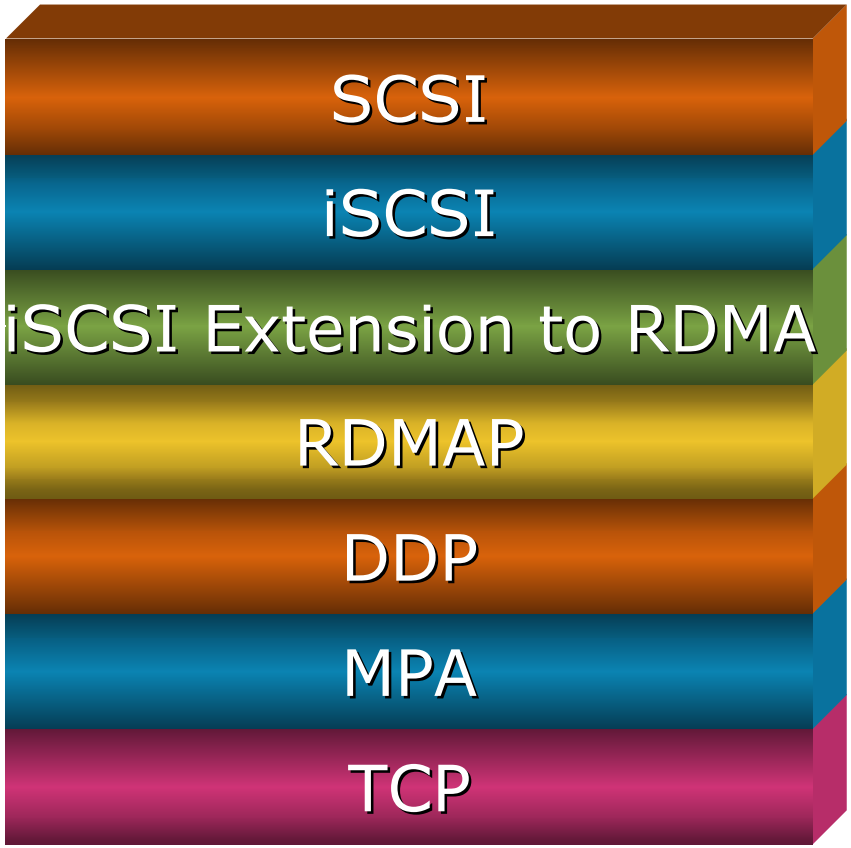

DDP: Direct Data Placement
MPA: Marker PDU Aligned for TCP
TCP: Transmission Control Protocol

iSCSI Potential

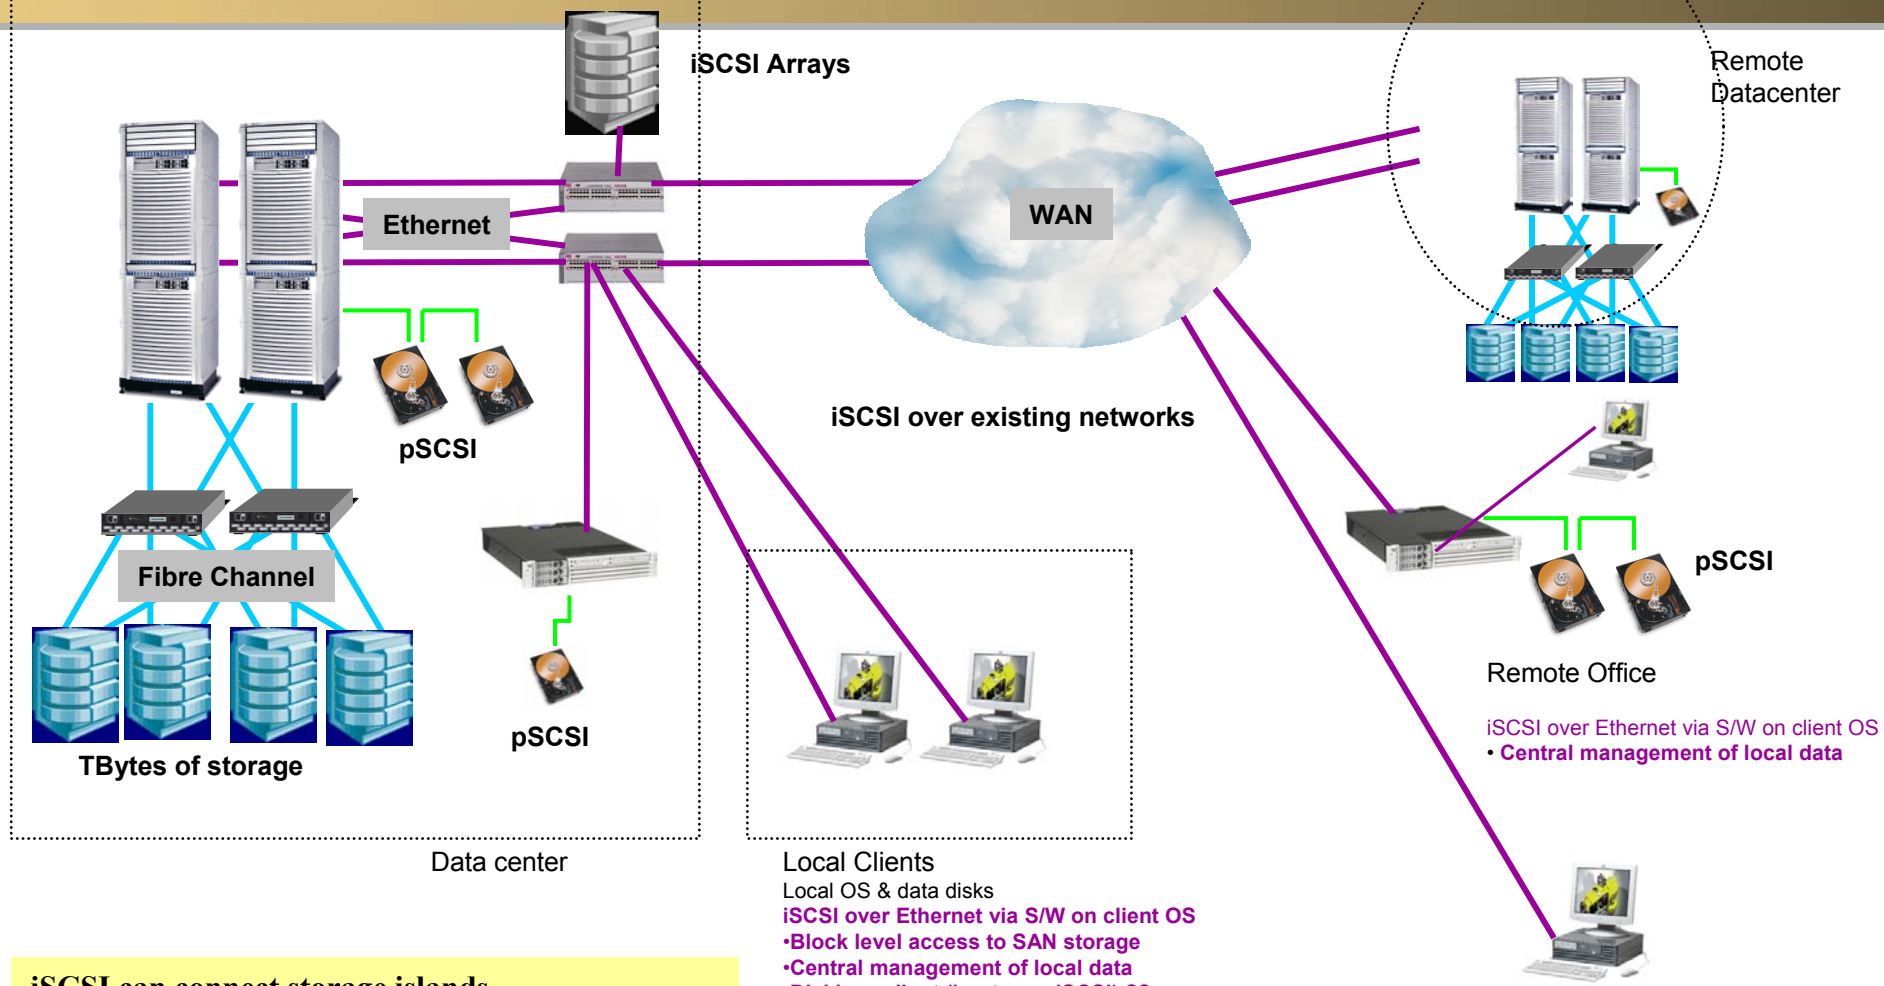

iSCSI can connect storage islands
legacy servers can connect to FC SAN
- clients can access SAN storage at block level
- allows central management of local client data
- effective protocol to connect remote data centers

Serial Interfaces (SAS & SATA)

- Issues with of parallel storage interface (eg SCSI, ATA)
  - Clock frequencies have increased to keep pace with bandwidth requirements
  - Increased “cross-talk” and signal distortion
  - Large number of cable conductors involved
- Serial interfaces aim to mitigate these issues
  - Transition from parallel to serial interface needed to meet future increase in performance
  - Enables simpler cabling and improved system reliability
  - Serial ATA (SATA), a replacement for parallel ATA
  - Serial Attached SCSI (SAS), a replacement for parallel SCSI

Serial ATA (SATA)

- Point-to-point interface that preserves ATA's low cost and software compatibility
- Desktop inheritance
  - Driven by low cost
  - Not optimized for fast access to data
- Can be used as "drive sets" and connected to server through any other SCSI interface such as FC or iSCSI
  - Lower hardware cost
  - RAID redundancy will take care of drive failures
  - Limited data integrity, but can be used for document imaging, backup and data streaming applications
- Transparent with parallel ATA

ATA Standards Architecture

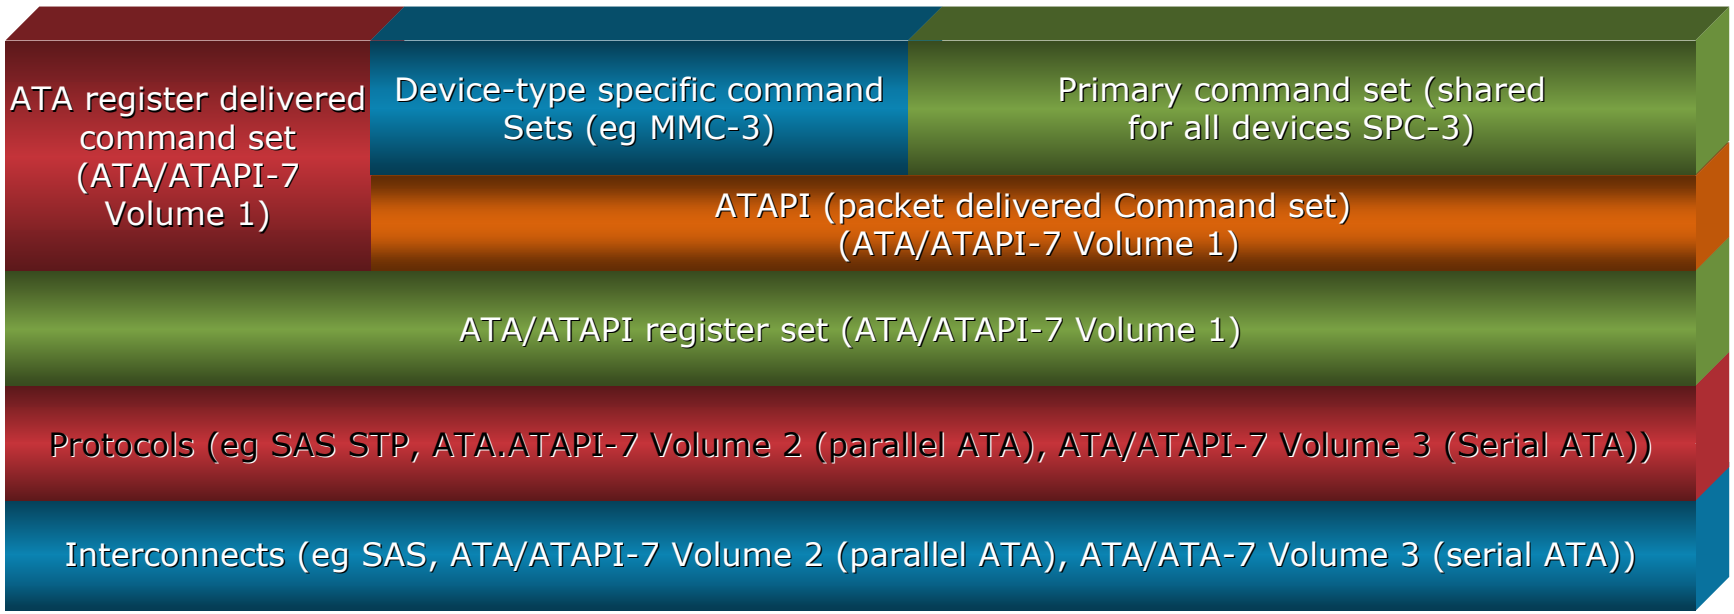

Serial Attached SCSI (SAS)

- Original promoters of SAS interface-Compaq, IBM, LSI Logic, Maxtor and Seagate
- SCSI protocol transported over serial interface
- Point-to-point full duplex, dual-port interface
- Provides SATA-to-SAS compatibility
- Three transport protocols
  - Serial SCSI Protocol (SSP) to support SAS disk & tape drives
  - Serial ATA Tunneling Protocol
  - Serial Management Protocol (SMP) to support SAS expanders
- Expander allows connectivity to 64 ports

SAS Layering

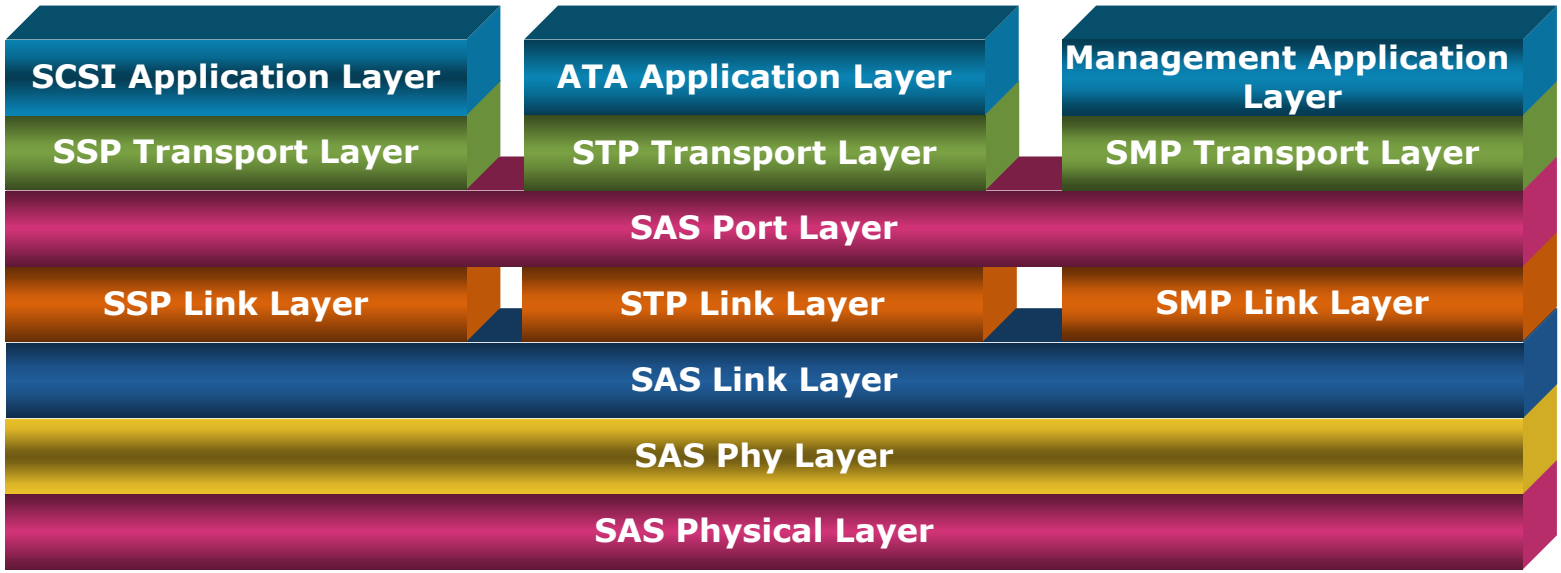

SAS/SATA Topology

Direct attach =
Number of drives limited to number of ports in the HBA

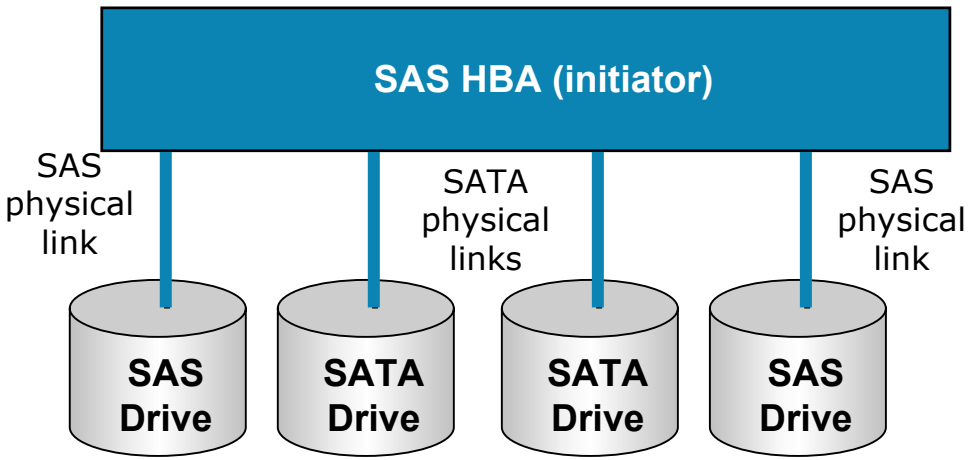

SSP (Serial SCSI Protocol) used to communicate with SAS drives

SATA (Serial ATA) used to communicate with SATA drives over SATA physical links

Expander attach =
More drives than HBA ports

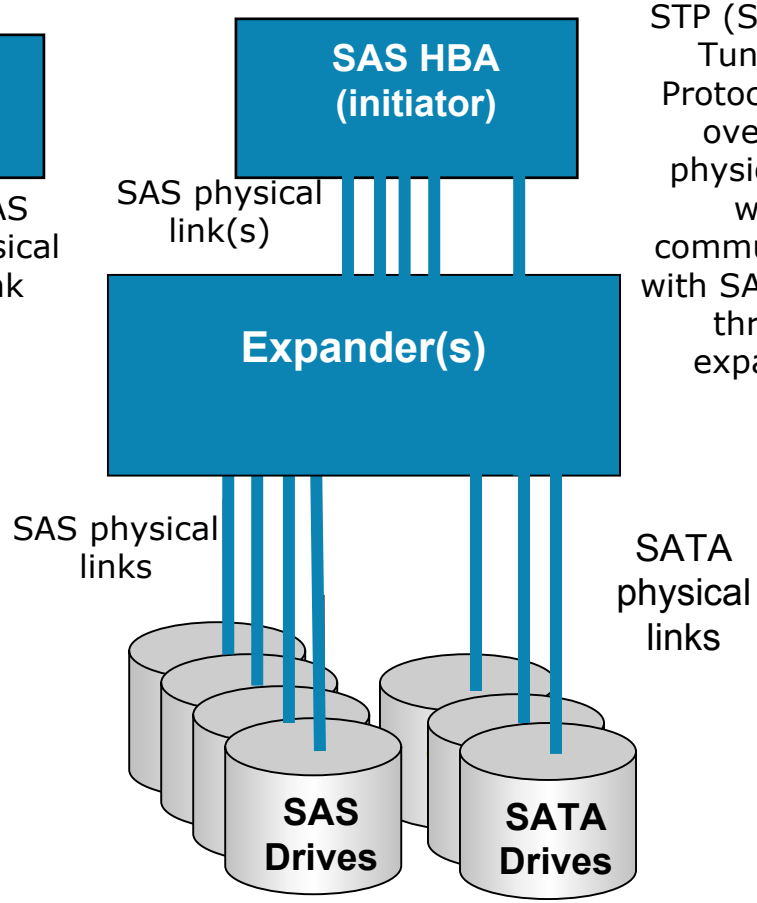

STP (Serial ATA Tunneling Protocol) used over SAS physical links when communicating with SATA drives through expanders

Technology Comparison

- Performance capability
  - Higher Bandwidth
  - Rich command feature set
- Scalability
  - Long cable length
  - Physical device addressing range
- Cost per gigabyte
- Cost effectiveness
  - Low frequency access
  - Streaming and sequential data
  - Archival data

Technology Comparison

- Reliability and availability
  - Extensive error recovery techniques
  - Multi-initiator support
  - Higher MTBF

Technology Comparison

	Fibre Channel	iSCSI/iSER	SATA	SAS
Performance Capability (Link Speed)	High (1,2,4,10 Gb/sec)	High (1,10 Gb/Sec)	Low (1.5, 3 Gb/Sec)	Medium (3, 6 Gb/Sec)
Connectivity	100 km with optical	IP connectivity	1m internal	10 m external
Cost per gigabyte	High	High/Low	Low	High/Low
Cost effectiveness	Business Critical	All applications	Archival, reference data	Business internal & operations
Reliability & Availability	High	High	Low	High

Market Segments

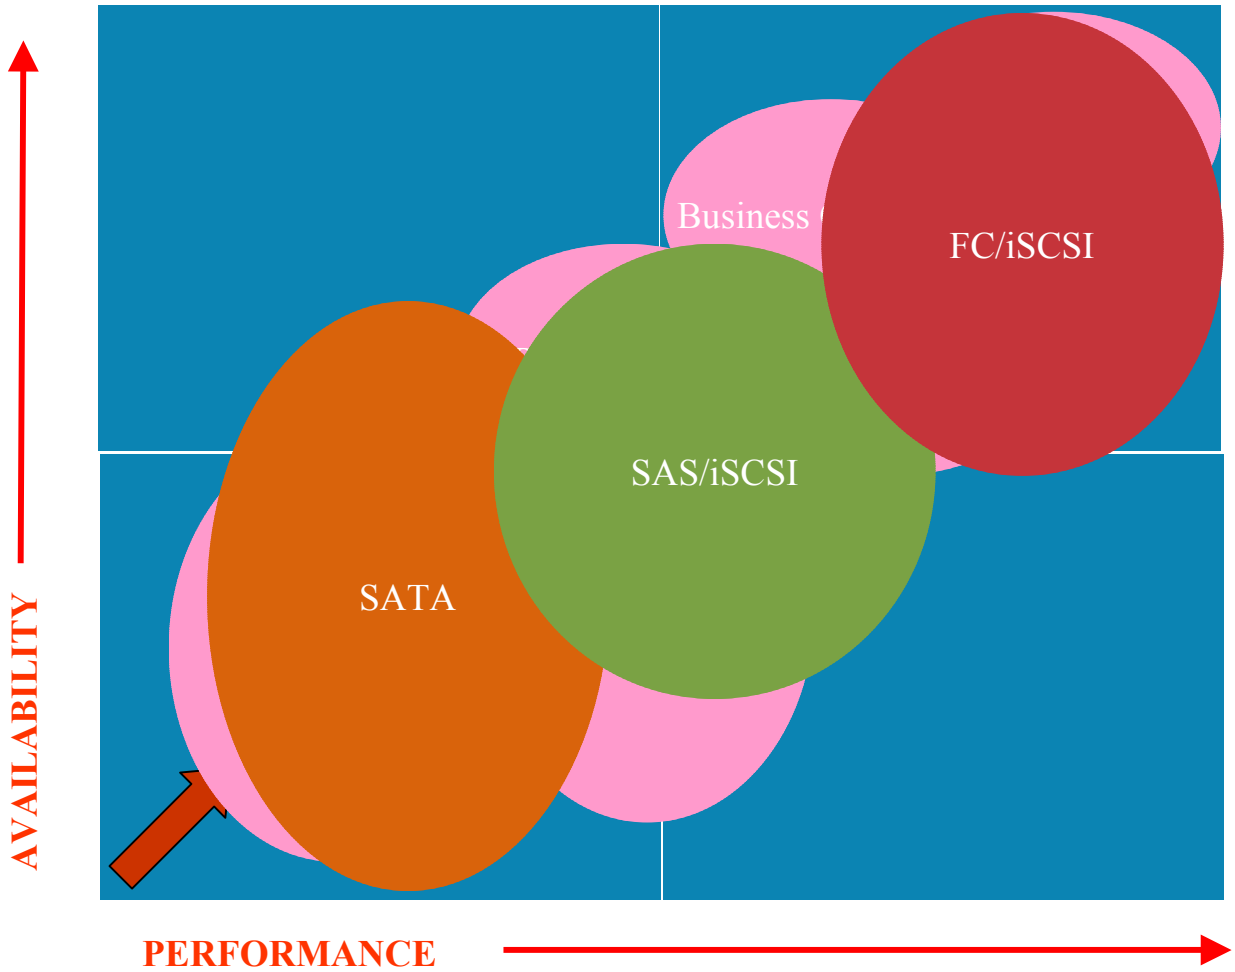

Interex, Encompass and HP bring you a powerful new HP World.