STORAGE PLUS 2006
storage networking and beyond

Tiering Up – Tiered Storage & Tiered Access

SATA, SAS, FC, FATA & Their Impact on Application Management

Tuesday, September 12, 2006, San Francisco, CA

Presented by:

Greg Schulz

Founder StorageIO Group

Author “Resilient Storage Networks” (Elsevier)

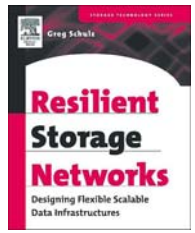

STORAGE PLUS 2006
storage networking and beyond

ByteandSwitch

What is tiered storage?

- Nothing new, we have had tiered storage for years
  - Disk and tape, fast disk and slow disk, on-line and off-line
- More than Disk vs. Tape, Fibre Channel vs. SATA. Etc…
  - Tiered protection, access, QoS, performance, media
- Align technology to functionality and economic value
  - Transactional vs. storage intensive, on-line vs. off-line
- Non-stop to HA, low RTO/RPO to high RTO/PRO
  - Level of accessibility, survivability, retention and preservation

Storage trends – setting the stage

- More data being generated and stored longer
- Increasing reliance up data being available
- Awareness of internal and external threats to data
- Larger amount of storage now accessed and stored remote
- Copies of data continue to increase as do logs and journals
- Storage I/O performance gap continues to widen
- Continued focus on dollar per GByte or TByte

Server and I/O performance gap

See "Data Center I/O Performance Issues and Impacts" white paper at www.storageio.com

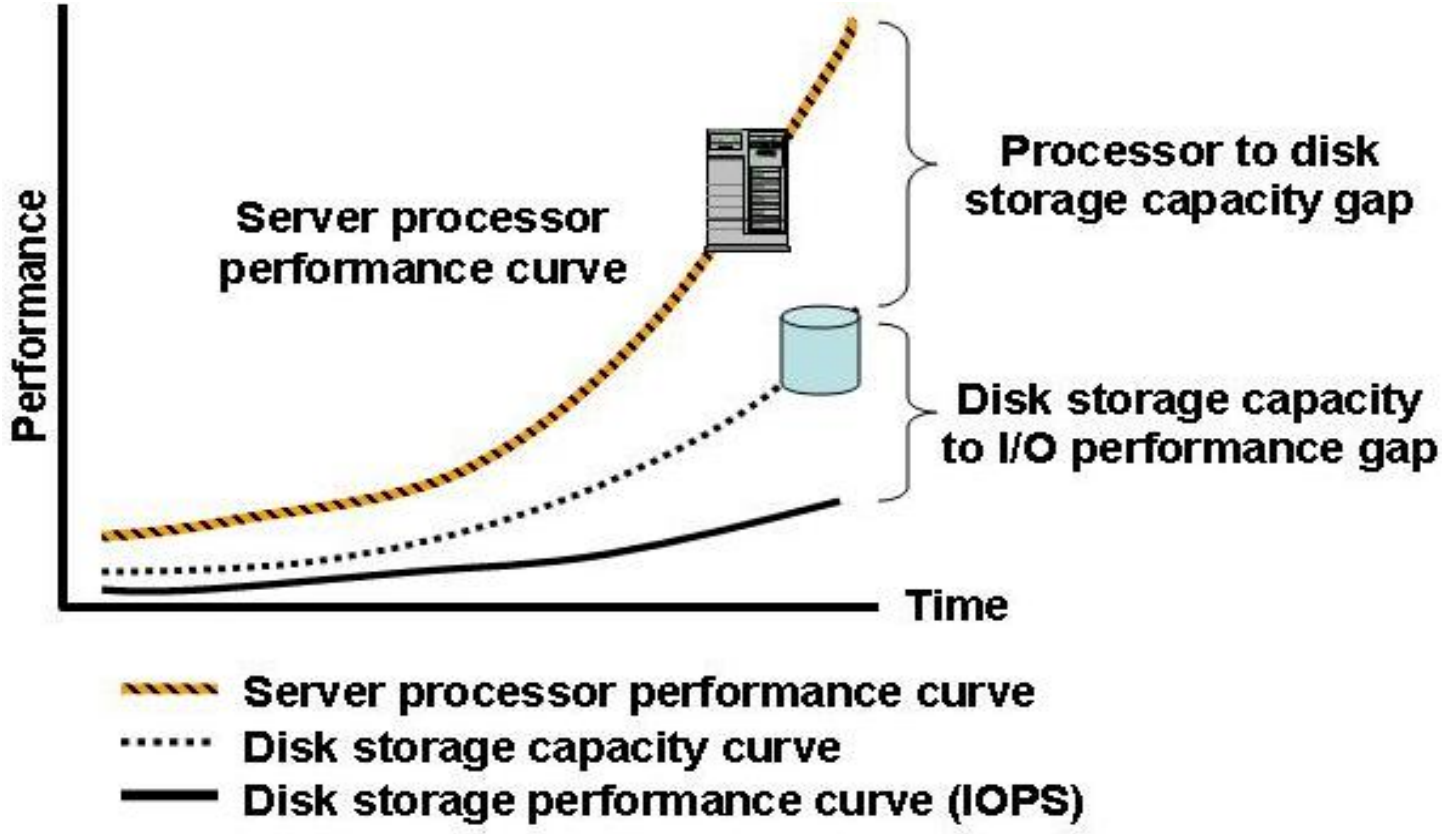

Why do you need tiered storage?

- Reduce storage management costs and complexity
- Accelerate performance while improving utilization
- Stretch IT dollars while enhancing service levels
- Enhance data protection and security of your data
- Regulatory compliance and data preservation
- Stability and scale… Consolidation… Virtualization
- Ask yourself if you will really benefit from tiered storage?

Tiered Storage – The Big Picture

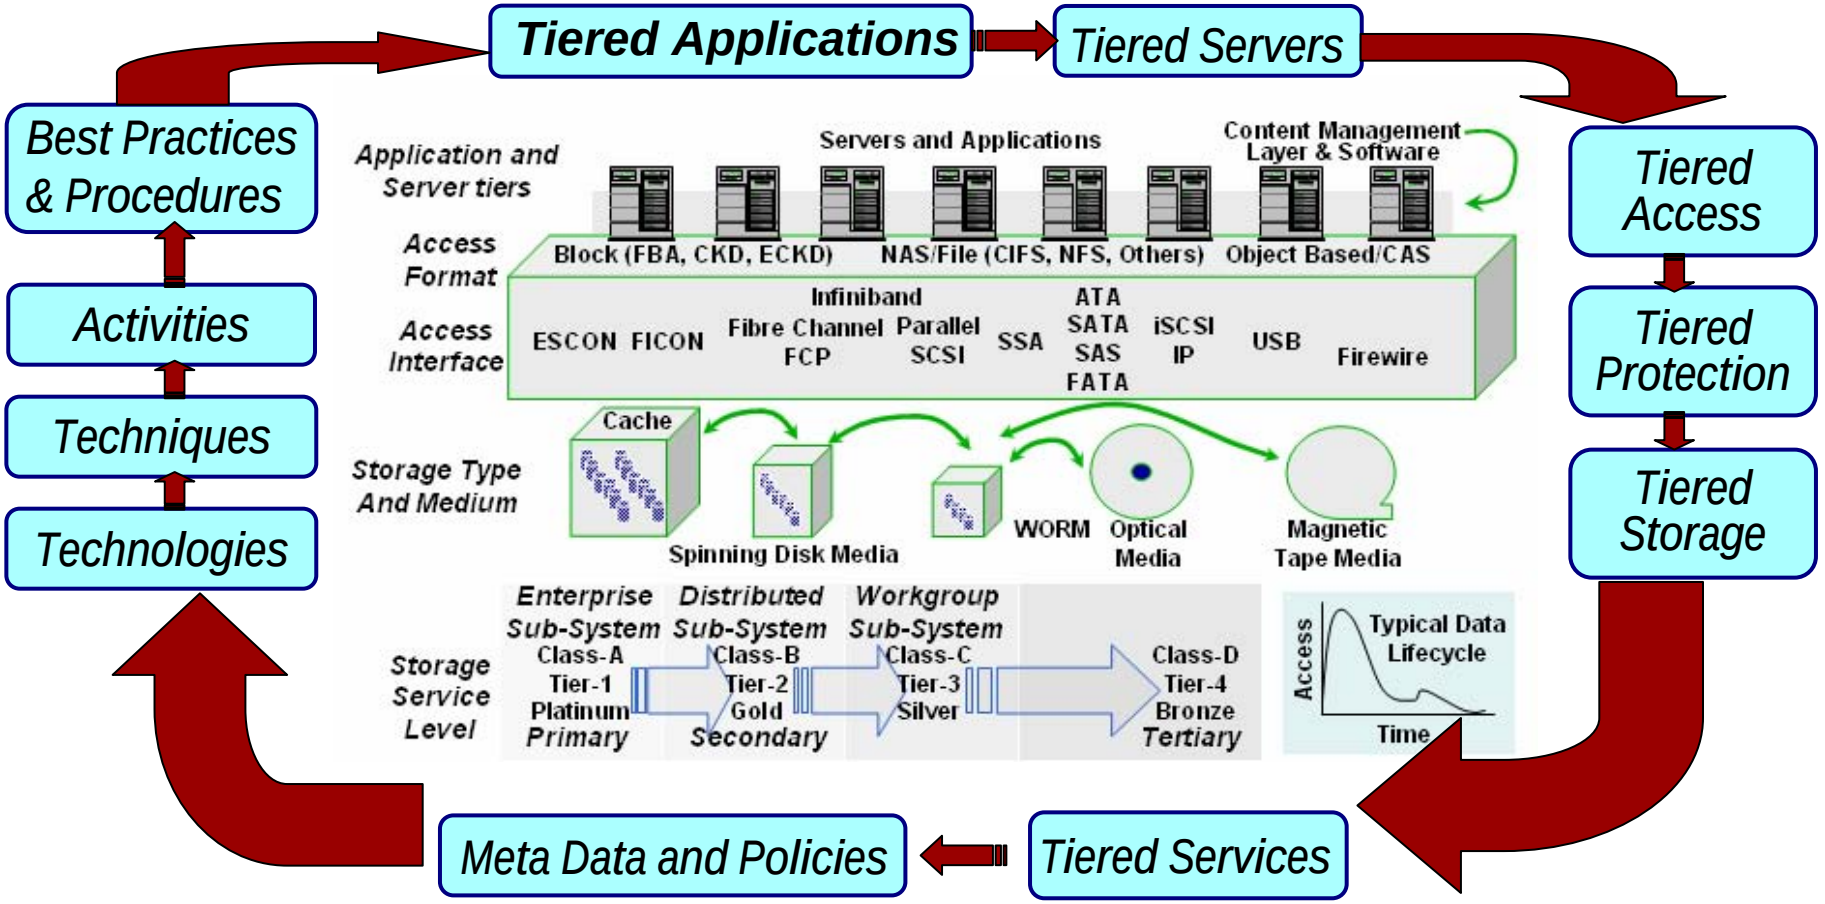

STORAGE PLUS 2006
storage networking and beyond

ByteandSwitch

What are your tiering options

- Tiered access (how you get at and access data)
  - Direct attached storage (DAS), network attached storage (NAS), storage area network (SAN), serial attached SCSI (SAS), serial attached ATA (SATA), iSCSI, Fibre Channel (FC), InfiniBand, universal serial bus (USB), block (SAN, iSCSI, SAS, SATA, USB), file (NAS – NFS & CIFS), object/content addressed (CAS)
- Tiered protection (how the data is protected)
  - Clustered, RAID, security, CDP, backup, PIT/snapshot, mirror/replication
- Tiered media (how the data is stored)
  - Disk (JBOD, SBOD, RAID, Cluster), tape, optical, virtual tape library (VTL)
  - FC, FC near line (aka FATA), SATA, ATA, SAS, SSD, LTO, DVD
  - High performance, high capacity, desktop vs. enterprise

Tiered storage access options

- Interfaces and protocols, direct attached, networked
- Block access (DAS, SAN)
  - Fibre Channel – point to point &switched; 1/2/4/10Gb; longer distance
  - iSCSI – point to point or switched; 1Gb & 10Gb; lower cost
  - InfiniBand – point to point, switched; 10Gb+; short distance
  - SAS & SATA – point to point and local switch over short distance
  - USB/Firewire – entry level and low end
- File (NAS) and object (CAS)
  - NFS & CIFS via IP/Ethernet – Data and storage sharing
  - CAS via IP/Ethernet and APIs or gateways – Data archiving

Tiered storage protection options

- Tiered protection (how the data is protected)
  - Redundant controllers, active/active, active/passive, clustered
  - RAID – Mirroring, parity and dual parity with proactive rebuild
  - Continuous data protection (CDP), PIT copy, snapshots
  - Backup and restore, archiving and data preservation, WORM
  - Local and remote mirroring and replication with consistency groups
  - Encryption and security to protect and secure stored data

Tiered storage media options

- Tiered media (how and where the data is stored)
  - RAID arrays (cache/centric, distributed/modular, entry level)
  - Virtual tape library (VTL) and disk (or disc) libraries (VDL & VOL)
  - High performance – 15K FC and SAS disk drives or SSD
  - High capacity – Near line FC/FATA, SATA, Tape, Optical(DVD/Blue ray)
  - Enterprise class vs. desktop disk drives
  - Near-line vs. off-line storage

Storage trends and directions

What to watch and look for

- Faster interfaces and networks
  - PCI-Express (PCIe) & new chipsets eliminate host bottlenecks
  - Fibre Channel going from 4Gb to 8Gb or other speeds
  - 10Gb Ethernet improving, iSER/RDMA needs software to work
  - InfiniBand doing well in high performance environments
  - PCIe I/O virtualization (IOV) fabrics and networks
- RAID and controller optimization for performance
  - Rapid drive rebuild and copy along with dual parity (RAID6)
  - Some controllers optimized for small random, others large I/O
  - Power management for controllers and disk drives
  - Transparent data migration between storage tiers/pools

Techniques and technologies

Application and data characteristics

Tier-1 Business Essential	Customer facing systems, billing, email, OLTP and others essential to keep business running High performance, time sensitive, large workloads
Tier-2 Business Critical	Applications that are critical to keep business running however they can be un-available for limited periods of time. Moderate performance
Tier-3 Business Important Nice to have	Important to support the business perhaps adding some competitive value however not required to keep business running, low performance need
Backup/Nearline	Copies of data to restore deleted or corrupted files and to help support BC and DR functions
Long term archive	Data stored for long periods of time for compliance and preservation for possible future access if needed

Techniques and technologies

Where tiered storage technologies fit

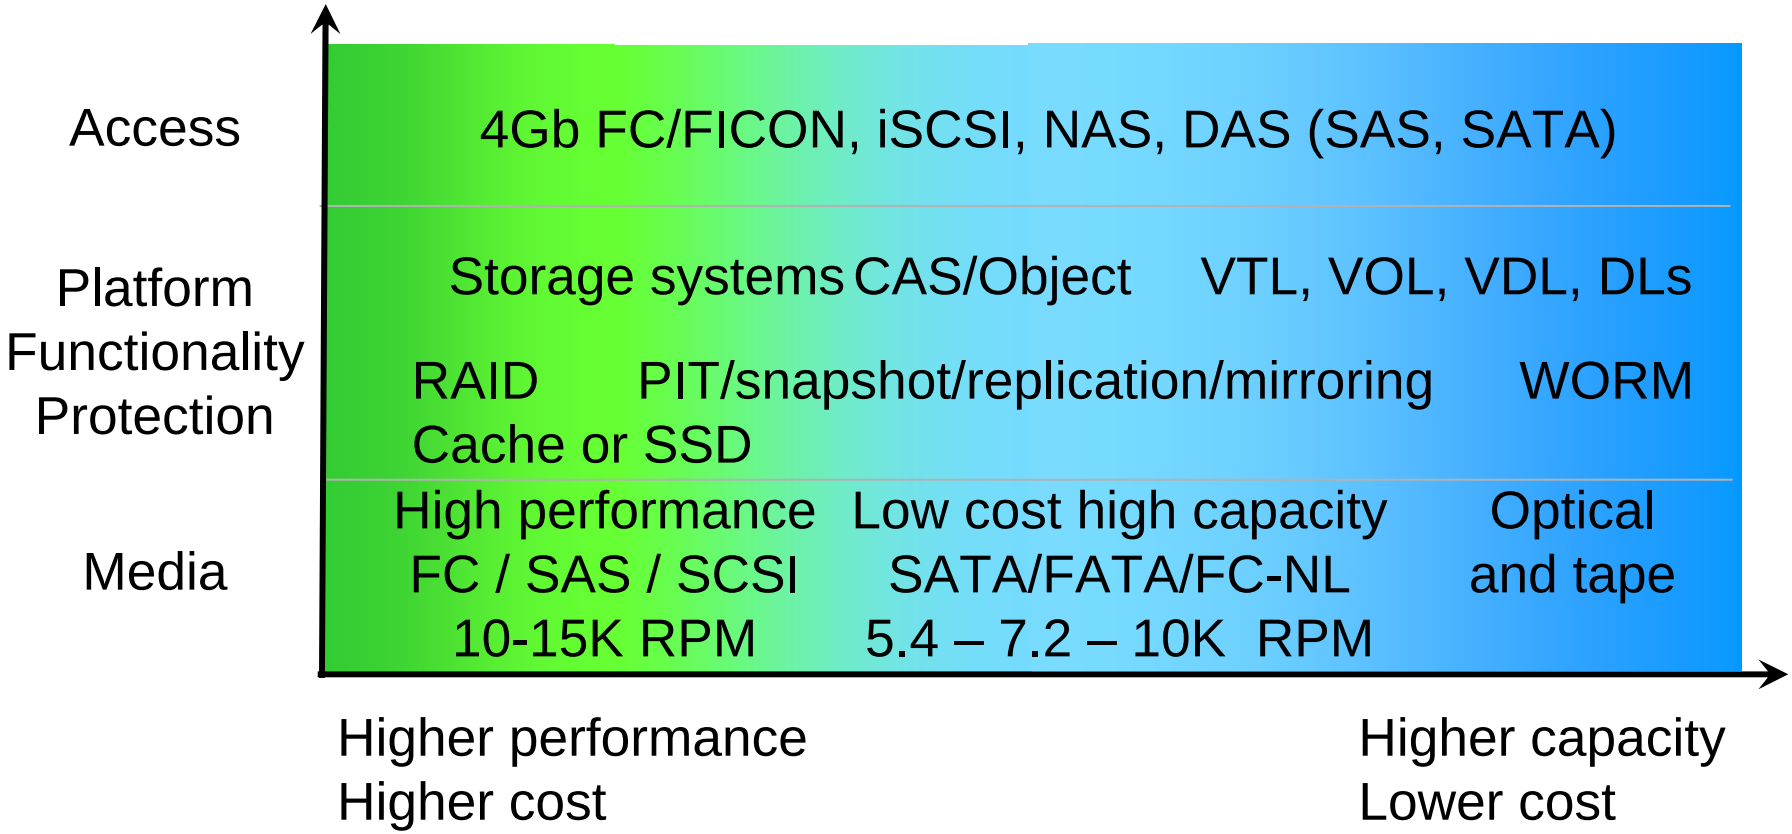

Techniques and technologies

Some typical application storage characteristics

Backup	Sequential large I/O (bandwidth), reads and write, scheduled
Archiving	Sequential I/O, writes, few reads, capacity centric
OLTP	Random small I/O (latency), mixed reads and writes, HA
Database	I/Os vary with type of DB as well as DB components
Email	Sequential, random on indices, mixed workloads
File server	Sequential, size varies by application, reads and writes
Video	Large sequential reads once stored, capacity centric

Tiered storage example

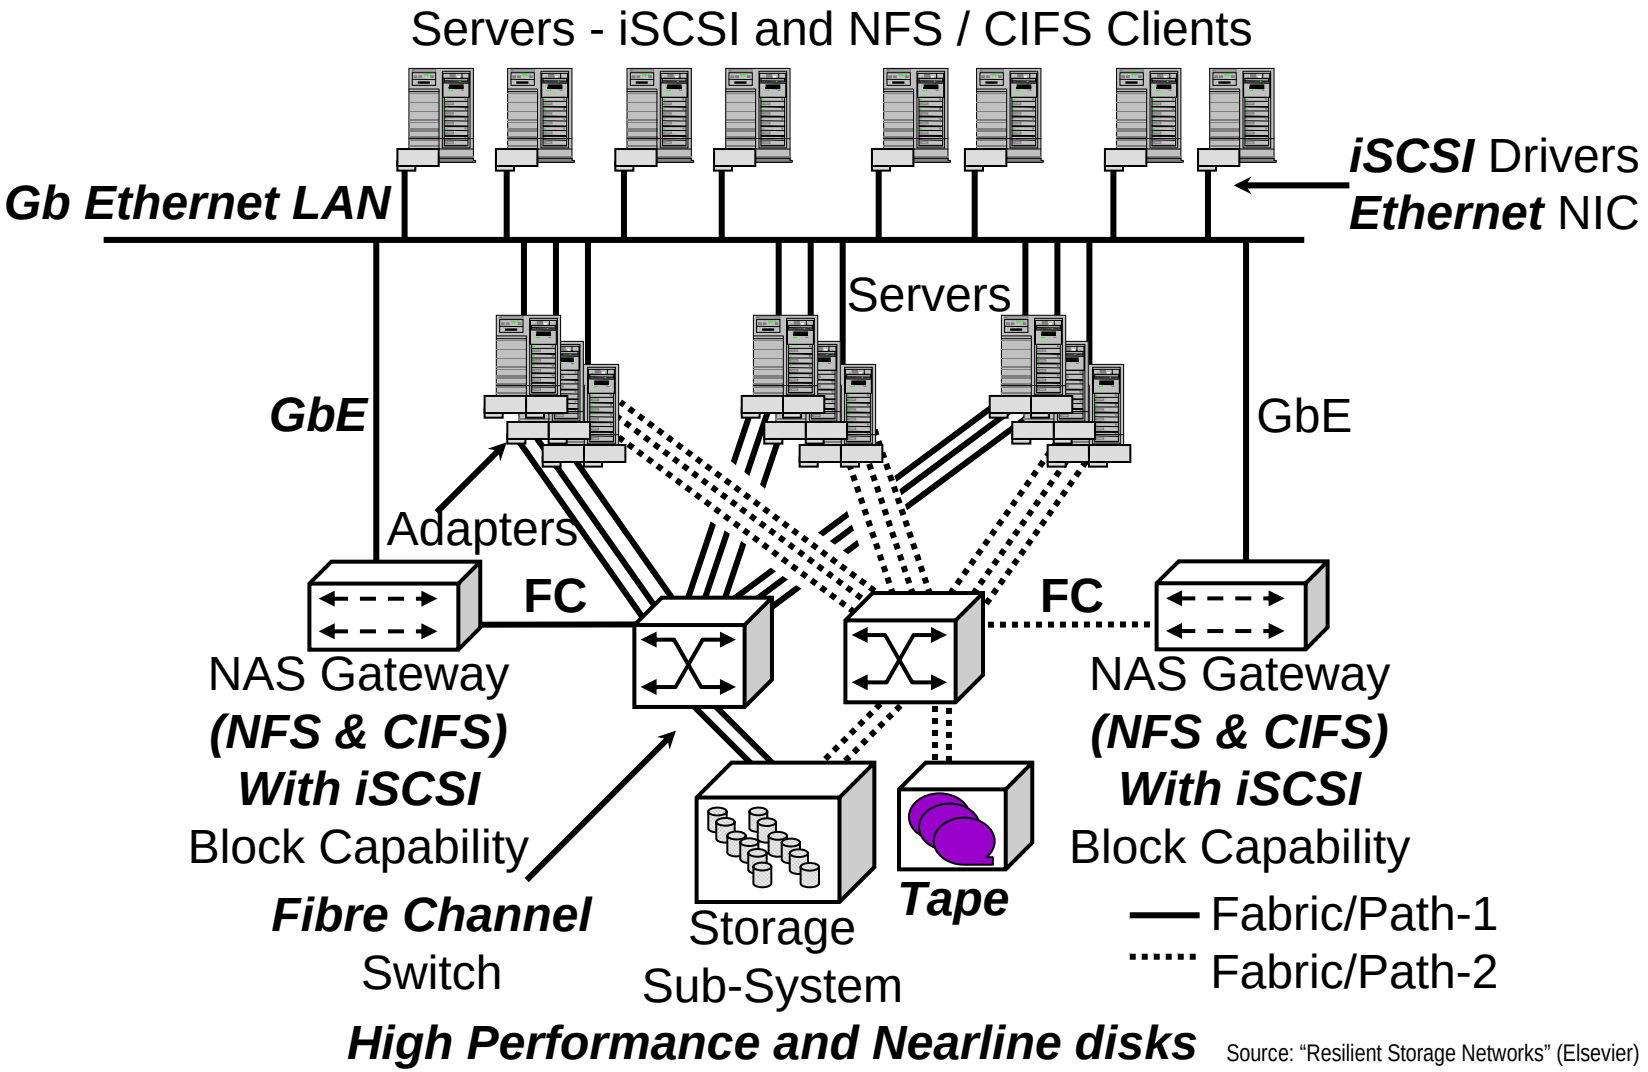

Source: "Resilient Storage Networks" (Elsevier)

Tiered storage example

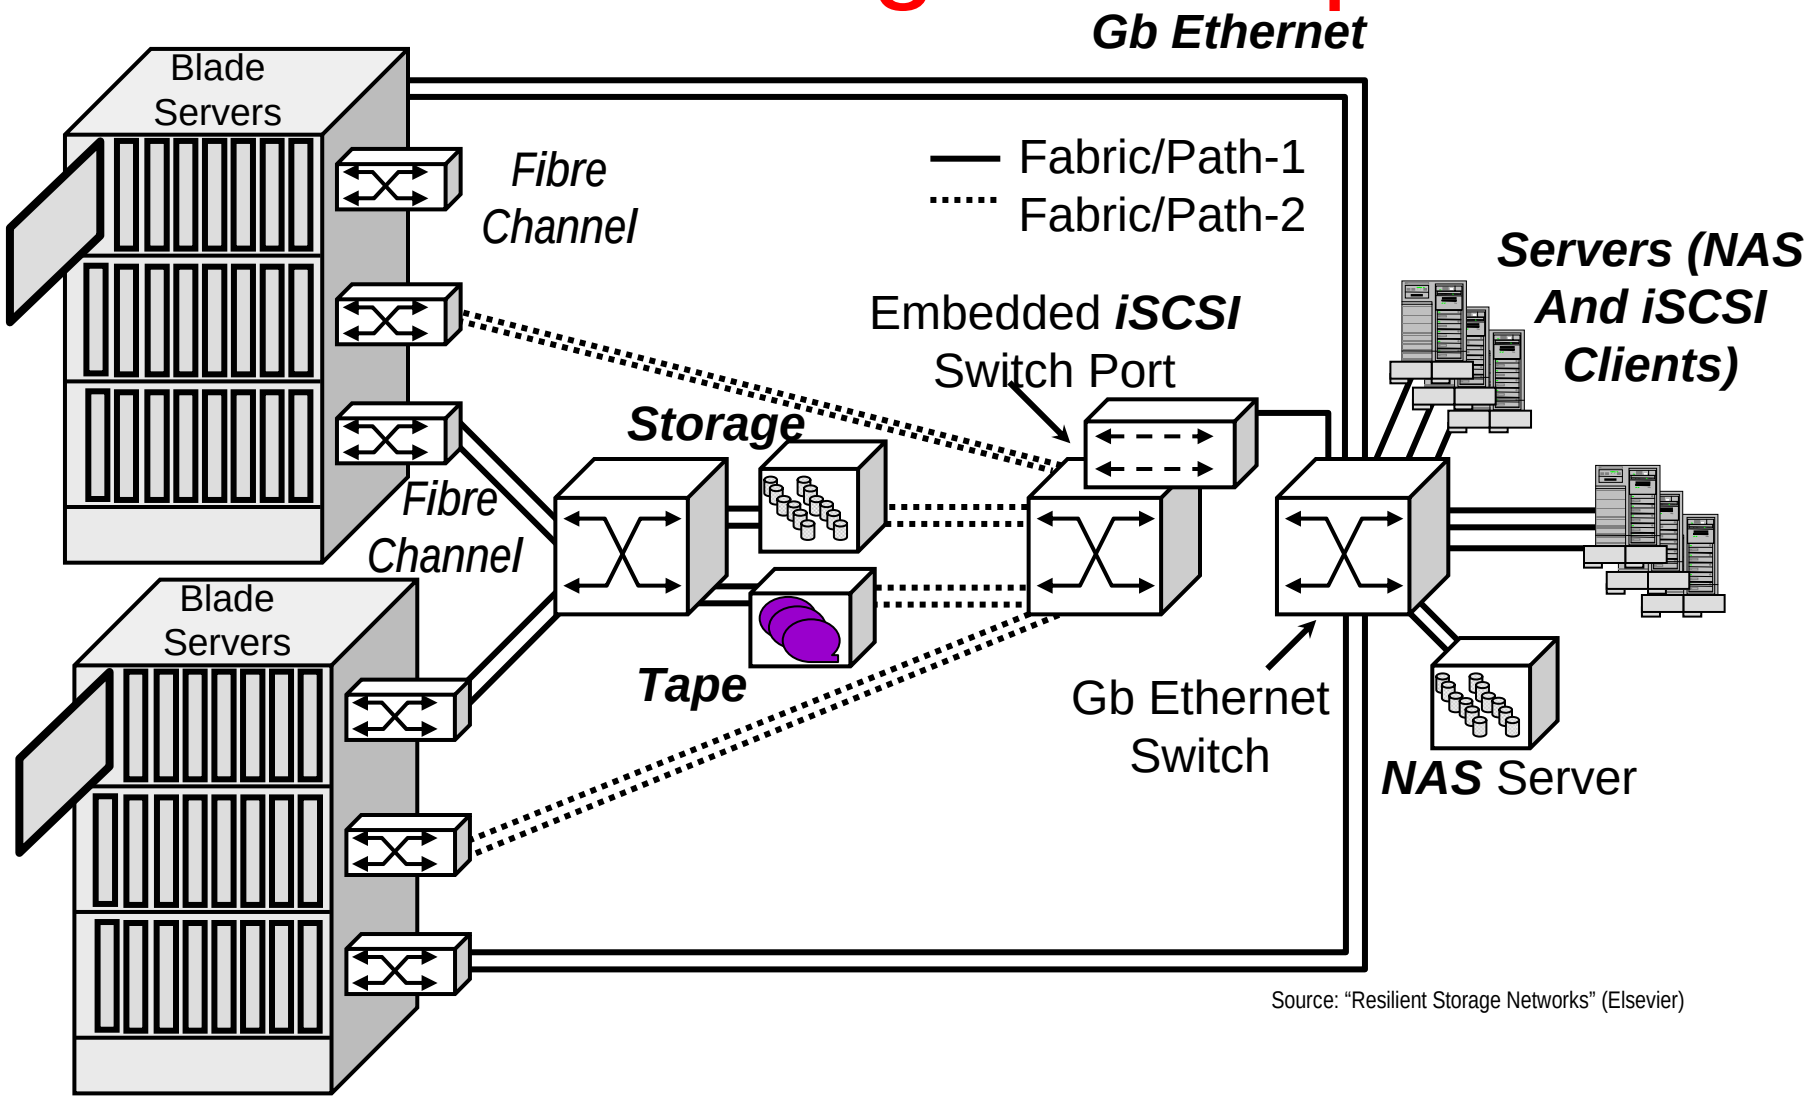

Source: "Resilient Storage Networks" (Elsevier)

STORAGE PLUS 2006
storage networking and beyond

ByteandSwitch

Comments and tips

- Data growth will seek out any and all available capacity
- Focus on data protection, preservation and accessibility
- Align defenses and protection to match applicable threats
- Leverage compaction and archiving technologies
- Storage capacity utilization increases result in I/O issues
- Look beyond cost per GByte/TByte at effective value
- Negotiate to reduce software costs, improve productivity
- Understand how archived data will be retrieved in the future

Closing comments

- Caveats of consolidation to watch out for
  - Reduced performance with improved capacity utilization
  - Single points of failure with consolidation if HA ignored
- Design for performance and growth
  - Archiving and compaction can improve overall performance
  - Eliminate complexity to reduce overhead and cost
- I can be reached at greg@storageio.com
- Of course the obligatory plugs:
  - Resilient Storage Networks – Elsevier
  - www.storageio.com (see portfolio for articles etc)

SNIA Recommended and endorsed Reading

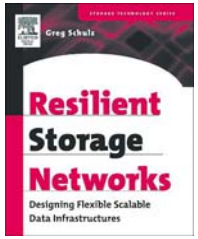

STORAGE PLUS 2006
storage networking and beyond

ByteandSwitch

STORAGE PLUS 2006
storage networking and beyond

Thank You

StorageIO

greg@storageio.com
www.storageio.com

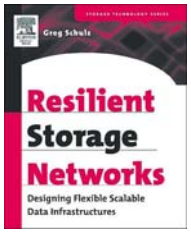

STORAGE PLUS 2006
storage networking and beyond

ByteandSwitch