Utilizing Low Cost Storage to Get the Most Out of Your Storage Infrastructure

Steve Arrington

Product Marketing Manager
Network Storage Solutions

HP WORLD 2003
Solutions and Technology Conference & Expo

Low Cost Storage Drivers

- Economic slowdown
  - Stagnate corporate IT budgets
- Demand for digital information continues to grow
  - More information being created, stored and shared among more people
  - Increasing amount of government regulation
- Technology advancements now make it much simpler and easier to implement low cost storage
  - Networked storage
  - Storage virtualization
  - Resource management

If all information is not valued equally, should all information be stored equally?

Information and Its Value

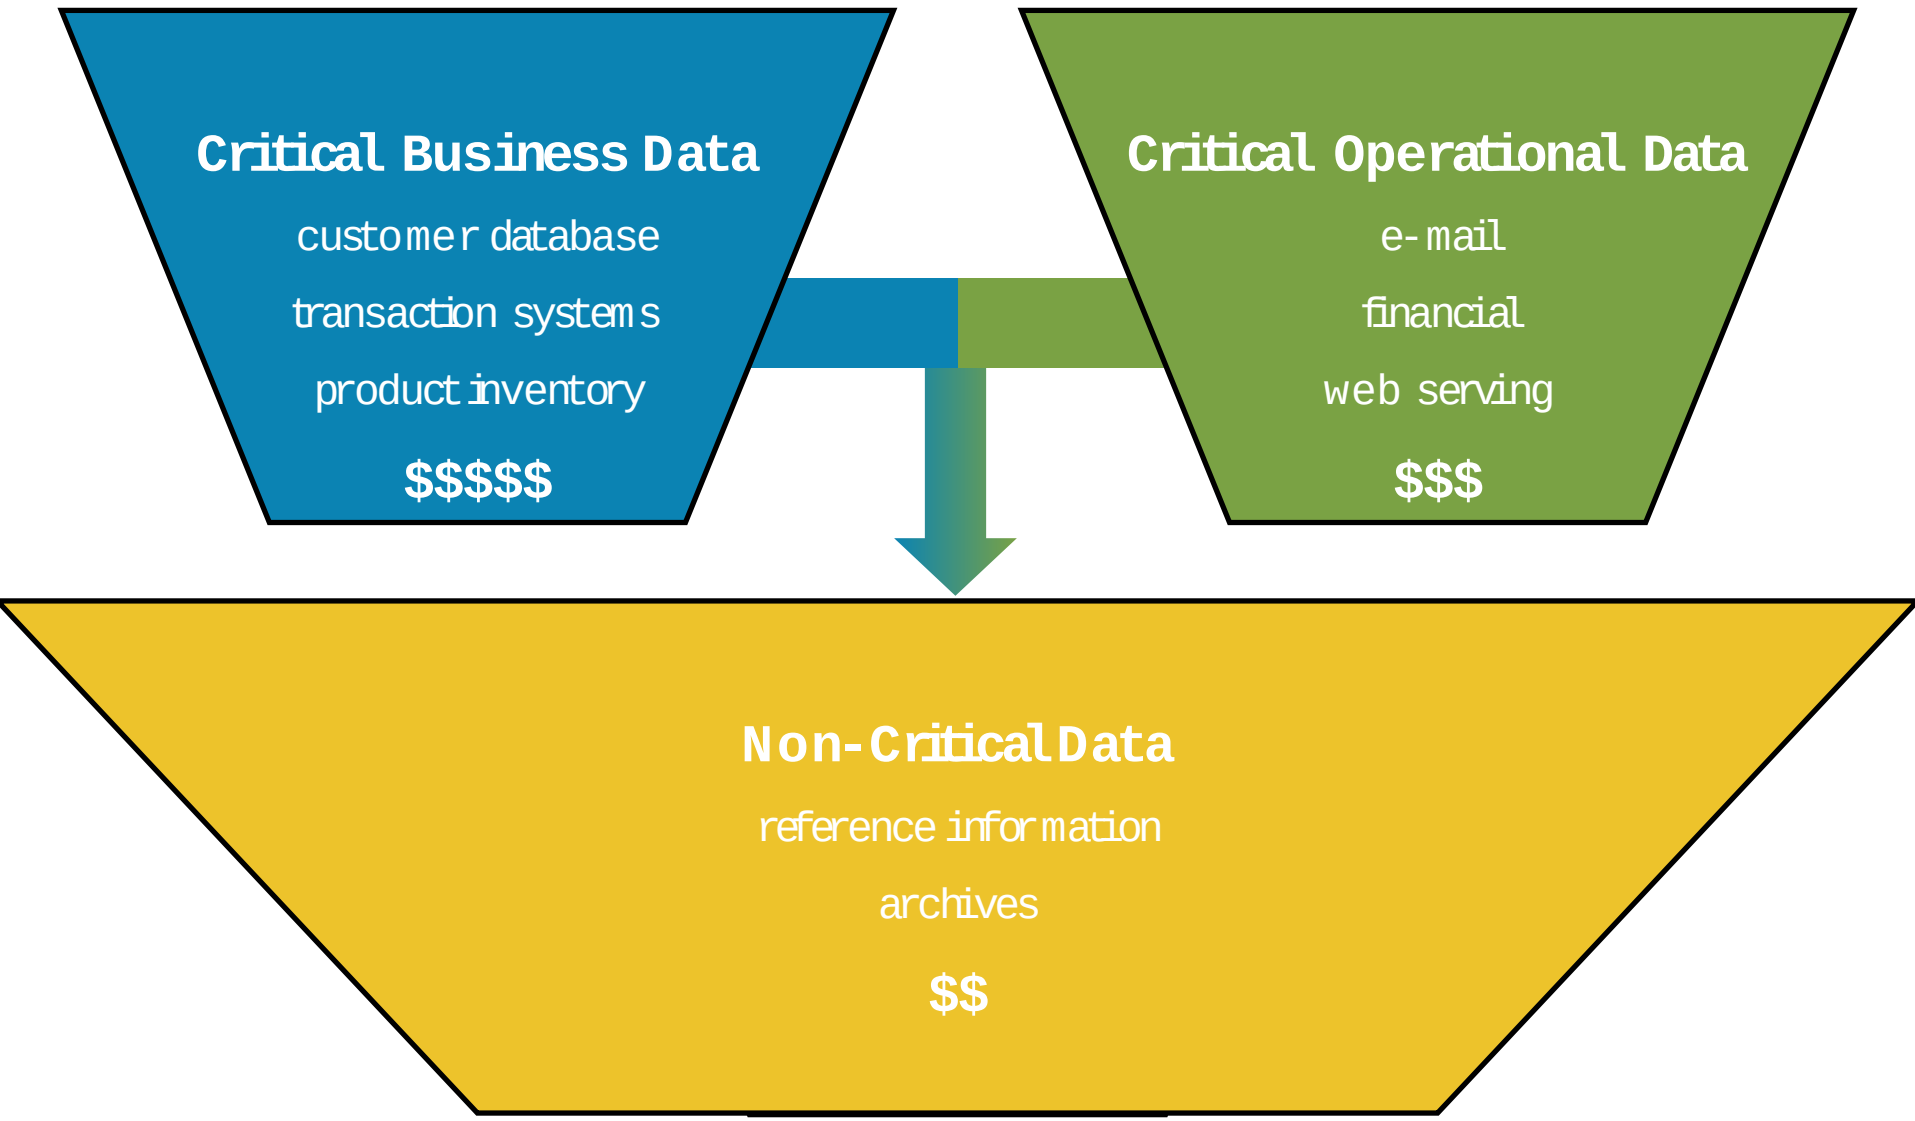

What's Driving Non-Critical Data Growth

The Increasing Value of Reference Information

- Industries transitioning from legacy, manual or paper based process to digital and electronic system
- Increasing government regulation of digital information
  - More than 10,000 regulations in the US alone
  - Regulations currently target a narrow set of records
  - retention period for some records can be measured in decades

Industries Impacted Most by Compliance Regulation

- Financial Services – 17 CFR 240.17a-3&4
  - Defines the types of records they must create and retain and for how long
- Healthcare – HIPAA
  - Describes security policies and procedures to ensure secure access, transmission and retention of personal health information
- Life Sciences/Pharmaceutical – 21 CFR Part 11
  - Describes the types of records to be retained and the use of electronic systems and records in place of paper and manual systems
- Government – DoD 2015.2
  - Certifies which applications or technology solutions an agency may implement to manage records
- Publicly Traded Corporations – Sarbanes-Oxley Act (2002)
  - Addresses corporate and auditor policies and procedures and the retention of records related to financial reports

Impacts Vertical Markets & Horizontal Applications

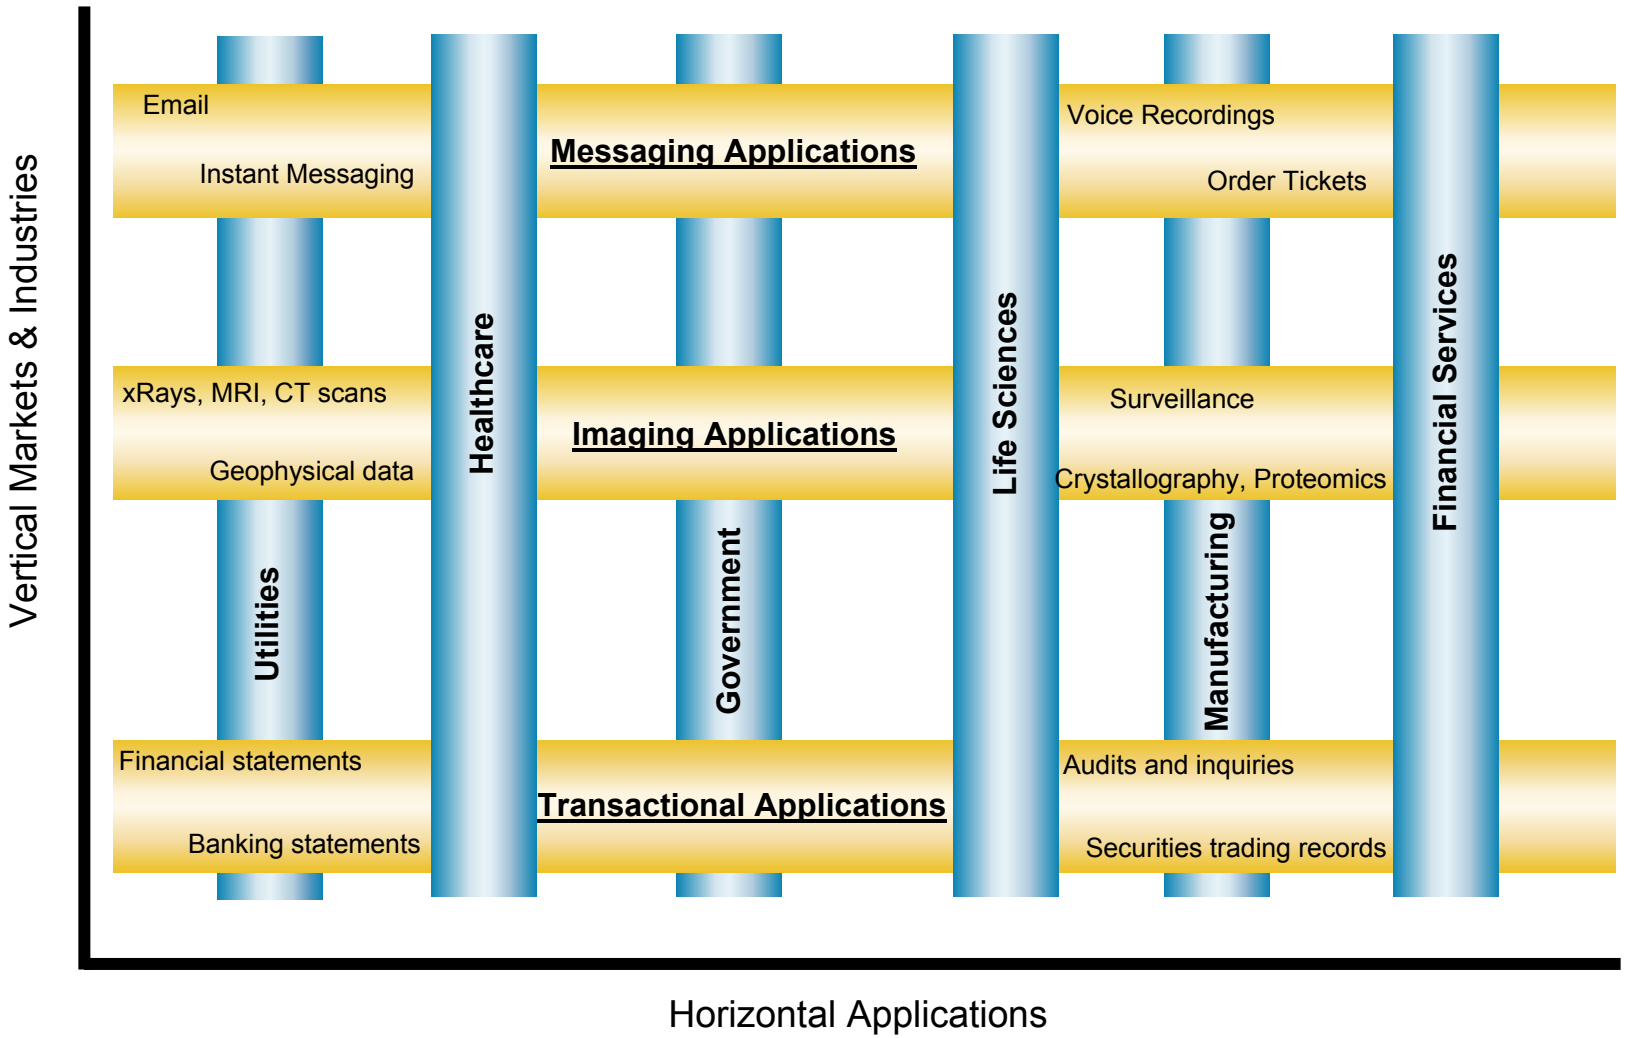

Retention Periods of Compliant Records

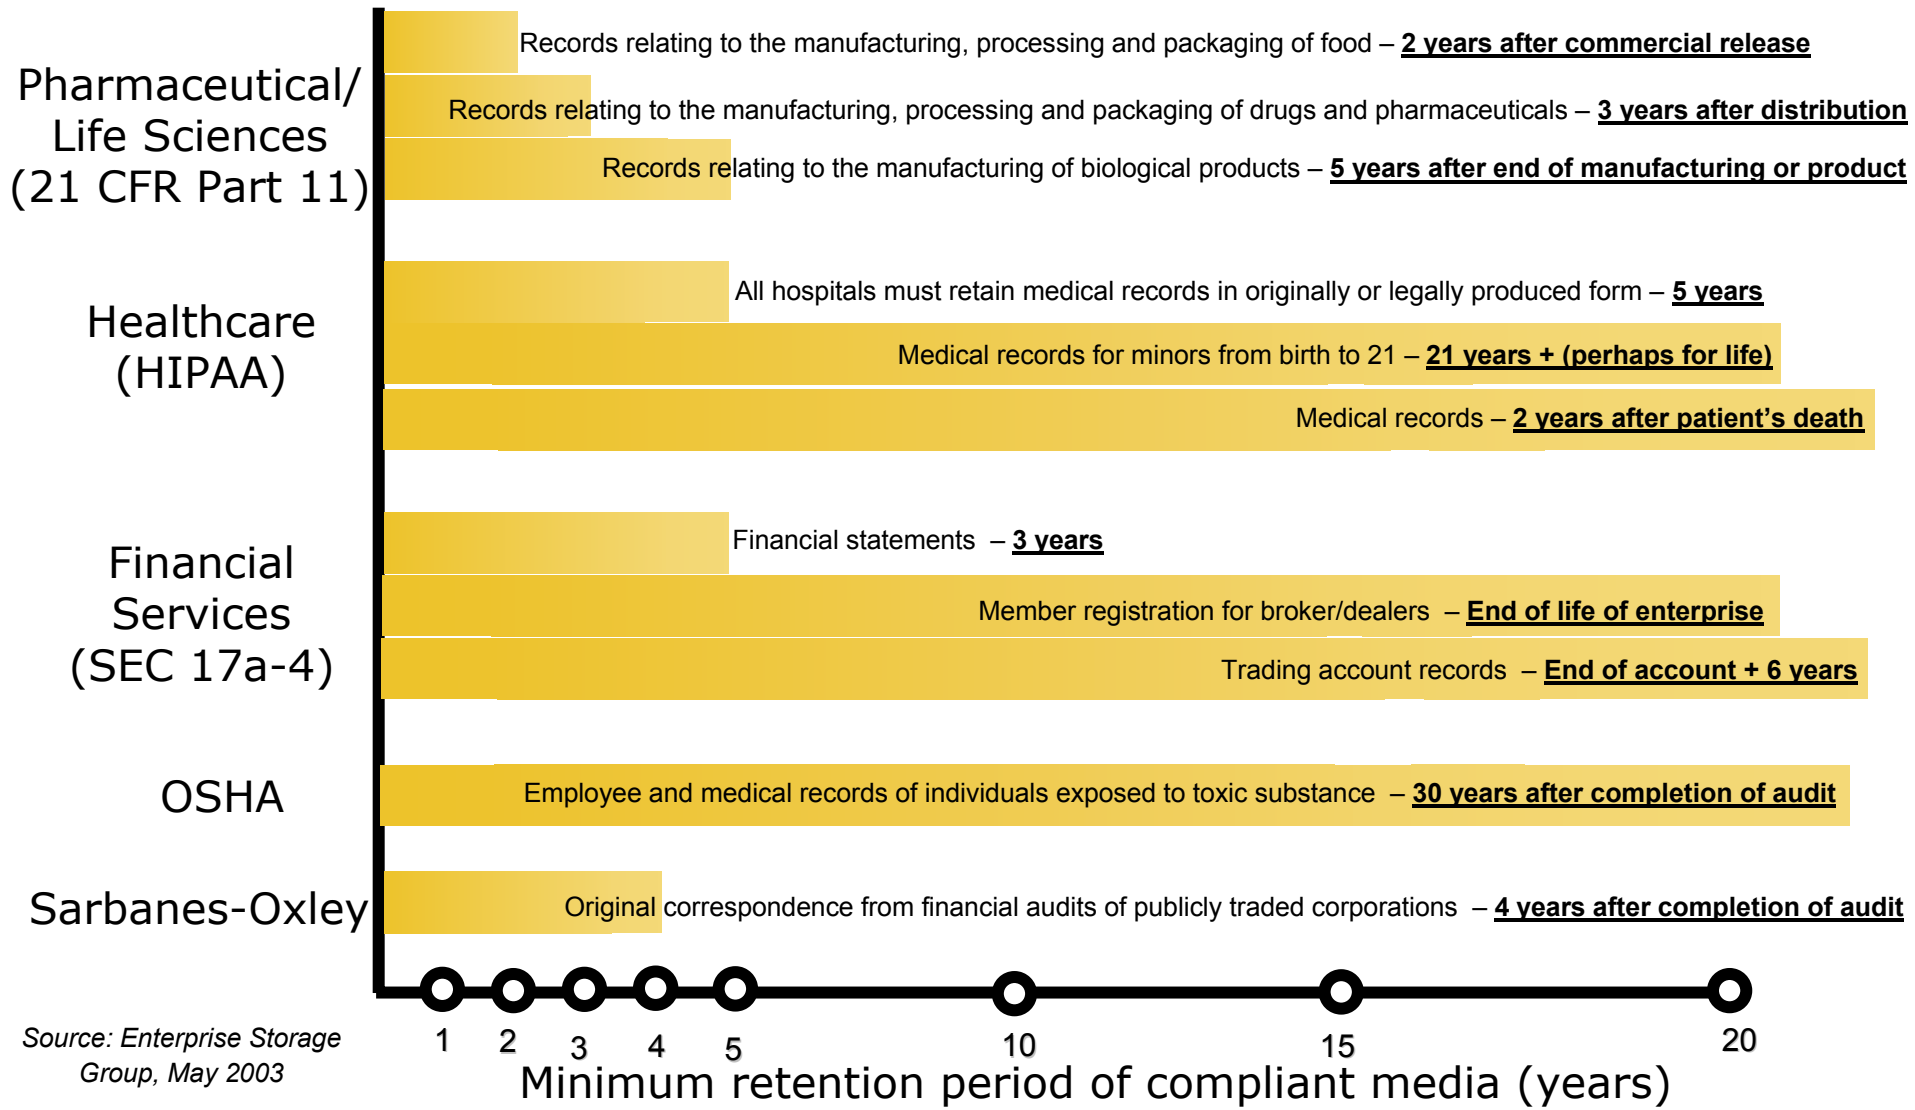

Source: Enterprise Storage Group, May 2003

Total Aggregate Capacity of Compliant Records

The capacity of compliant records will increase from 376PB in 2003 to 1,644PB in 2006, representing a CAGR of 64%

Worldwide scope; all vertical markets; Corporate, Commercial, and Government sectors

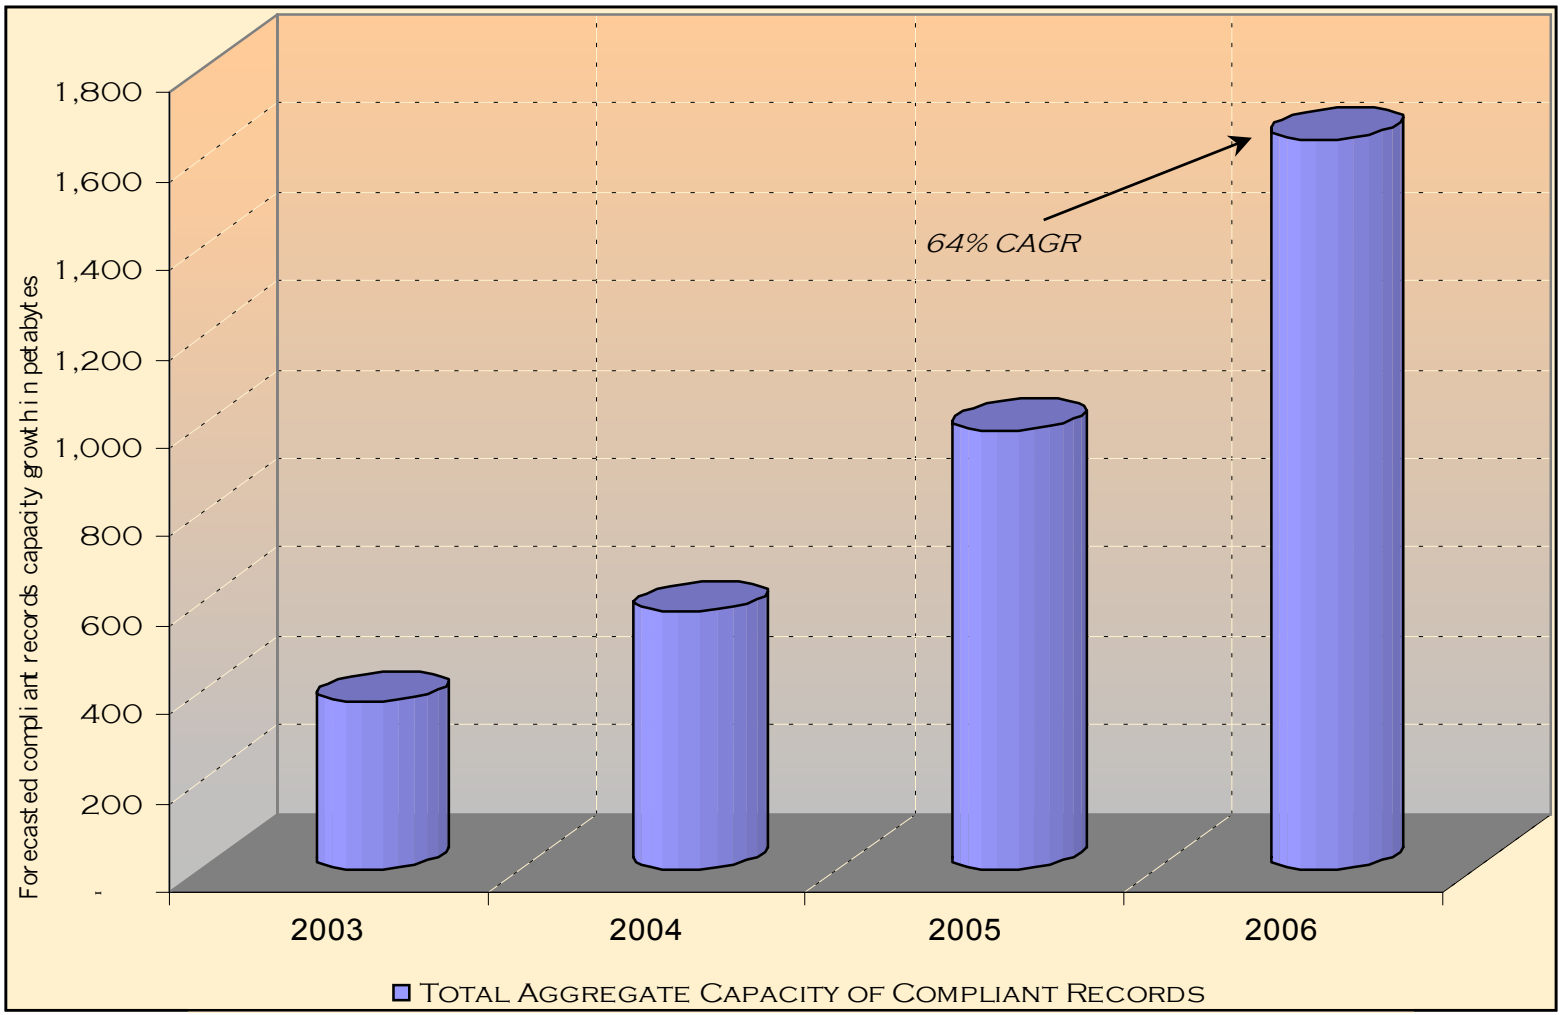

Source: Enterprise Storage Group, May 2003

Compliant Records Capacity by Media Type

While tape remains the most common media type for storing compliant records, the capacity of records stored on disk-based systems will increase 172% between 2003 - 2006

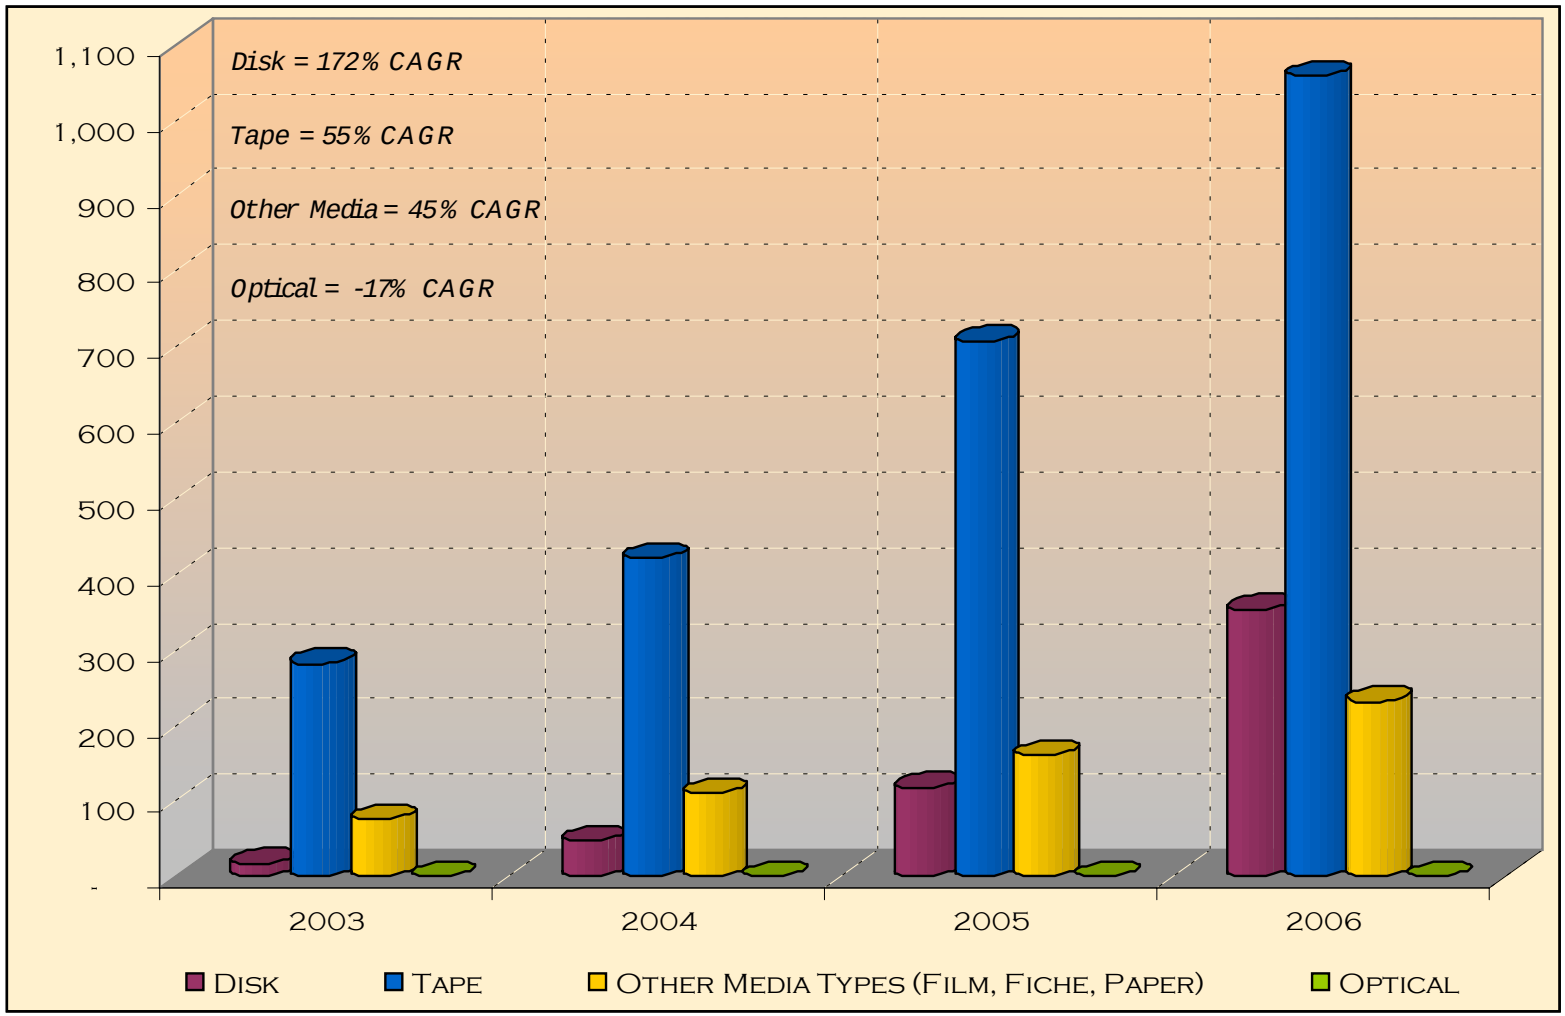

Worldwide scope; all vertical markets; Corporate, Commercial, and Government sectors

Source: Enterprise Storage Group, May 2003

Bottom Line of Compliance Regulation

- Regulations will expand in scope
  - Current regulations have set the stage
  - Rapid growth in the number of compliant records
  - Increased cost and complexity for businesses
- Increasing the need for lifecycle data management
  - Specified retention periods and disposition schedules
  - Costs are driving the need for new solutions
- Requirement for long-term discovery, legibility, authenticity and auditability
  - Failure to comply will result in penalities

Matching the Value of Data with the Value of Storage

Storage Quality of Service Characteristics

- Performance
  - Bandwidth – volume of data moved per unit of time
  - Throughput – speed of processing many small I/Os
  - Time to retrieve archived data
- Availability
  - Redundancy – multiple system components & copies of data
  - Remoteness – physical separation of systems & data
  - Recoverability – time to recover from an interruption
  - Durability of the storage media
  - Time associated with scheduled downtimes
- Cost
  - Acquisition cost
  - Operating cost

Matching Data Types to Quality of Service

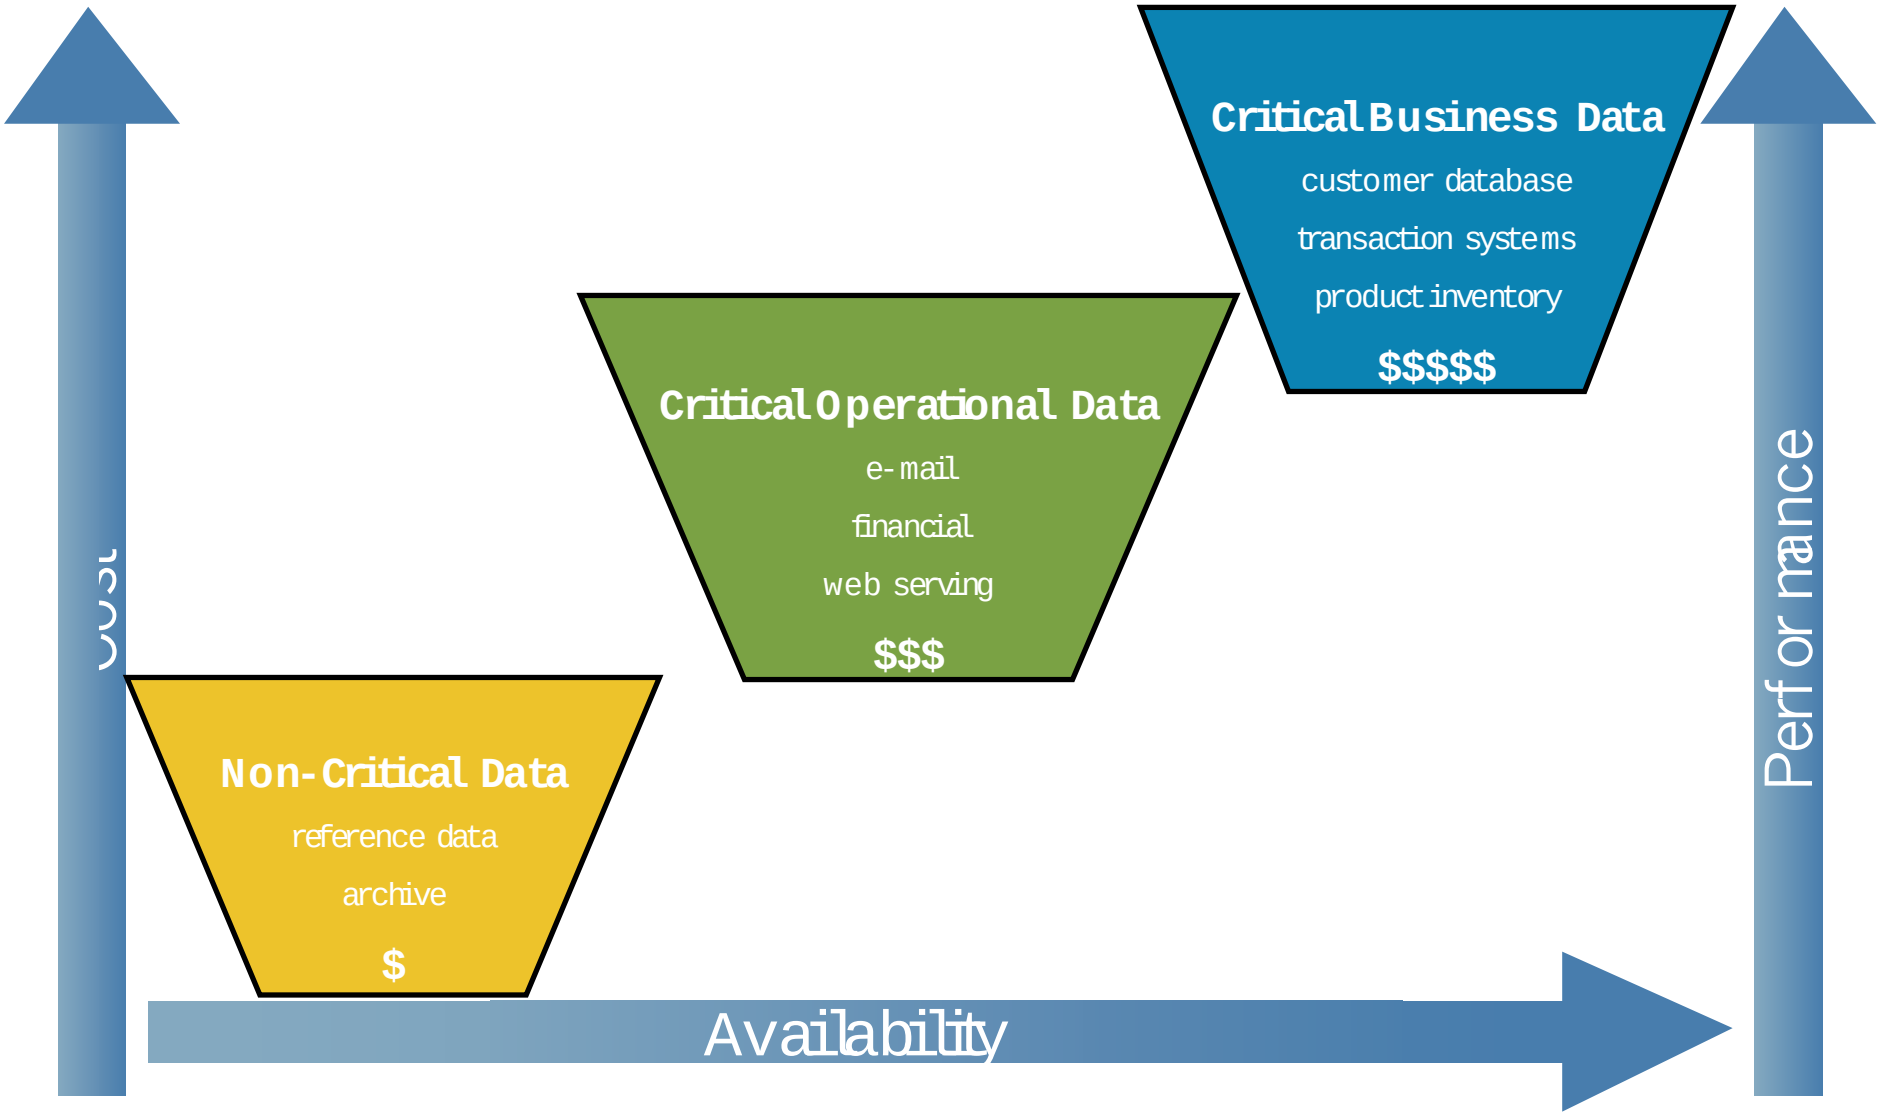

Traditional Architectural View of Storage

- Primary Storage
  - Monolithic
    - Huge capacity combined with very high-end features
  - Modular
    - Varying capacities and features “pay as you grow”
      high-end
      mid-range
      entry-level
  - “General purpose, good for everything” storage arrays
- Secondary Storage
  - Tape libraries and optical jukeboxes

Changing View of Storage

Categories change to align with how the storage is used rather than with the size of the organization or the amount of storage

- Primary storage
  - Fastest most reliable online storage for critical data
- Specialized storage
  - Cost effective online/near-online disk based storage for non-critical data
  - Combines hardware and software to perform specific tasks
- Secondary storage
  - Continues to provide long term retention and disaster recovery via tape libraries and optical jukeboxes

Desktop Class Drives in the Enterprise

It's a Quality of Service Trade-Off

- Cost
  - ATA drives are 1/3 the price of enterprise class drives
- Performance of ATA vs. enterprise class drives
  - Avg. read seek time: 12.5ms – 8.5ms vs. 4.7ms – 3.6ms
  - Rotational speed: 5,400 – 7,200 rpm vs. 10K – 15K rpm
  - ATA drives don't have command tag queuing
- Reliability
  - ATA drives: ~500K hours MTBF at 40% duty cycle
  - SCSI/FC drives: 1.2 million hours MTBF at 100% duty cycle
  - Drive warranty: 1year for ATA vs. 3 years for SCSI/FC
  - Increased time to recover from failure with ATA drives

Low Cost Storage Solutions

- Data Protection
  - Disk-to-Disk Backup
  - Asynchronous Remote Replication
- Content Management
  - Email Archiving
  - Hierarchical Storage Management (HSM)
  - Fixed data pools
    - Medical Imaging
    - Web Serving/Streaming Media
    - File & Print

Simple Tiered Storage Solution Example

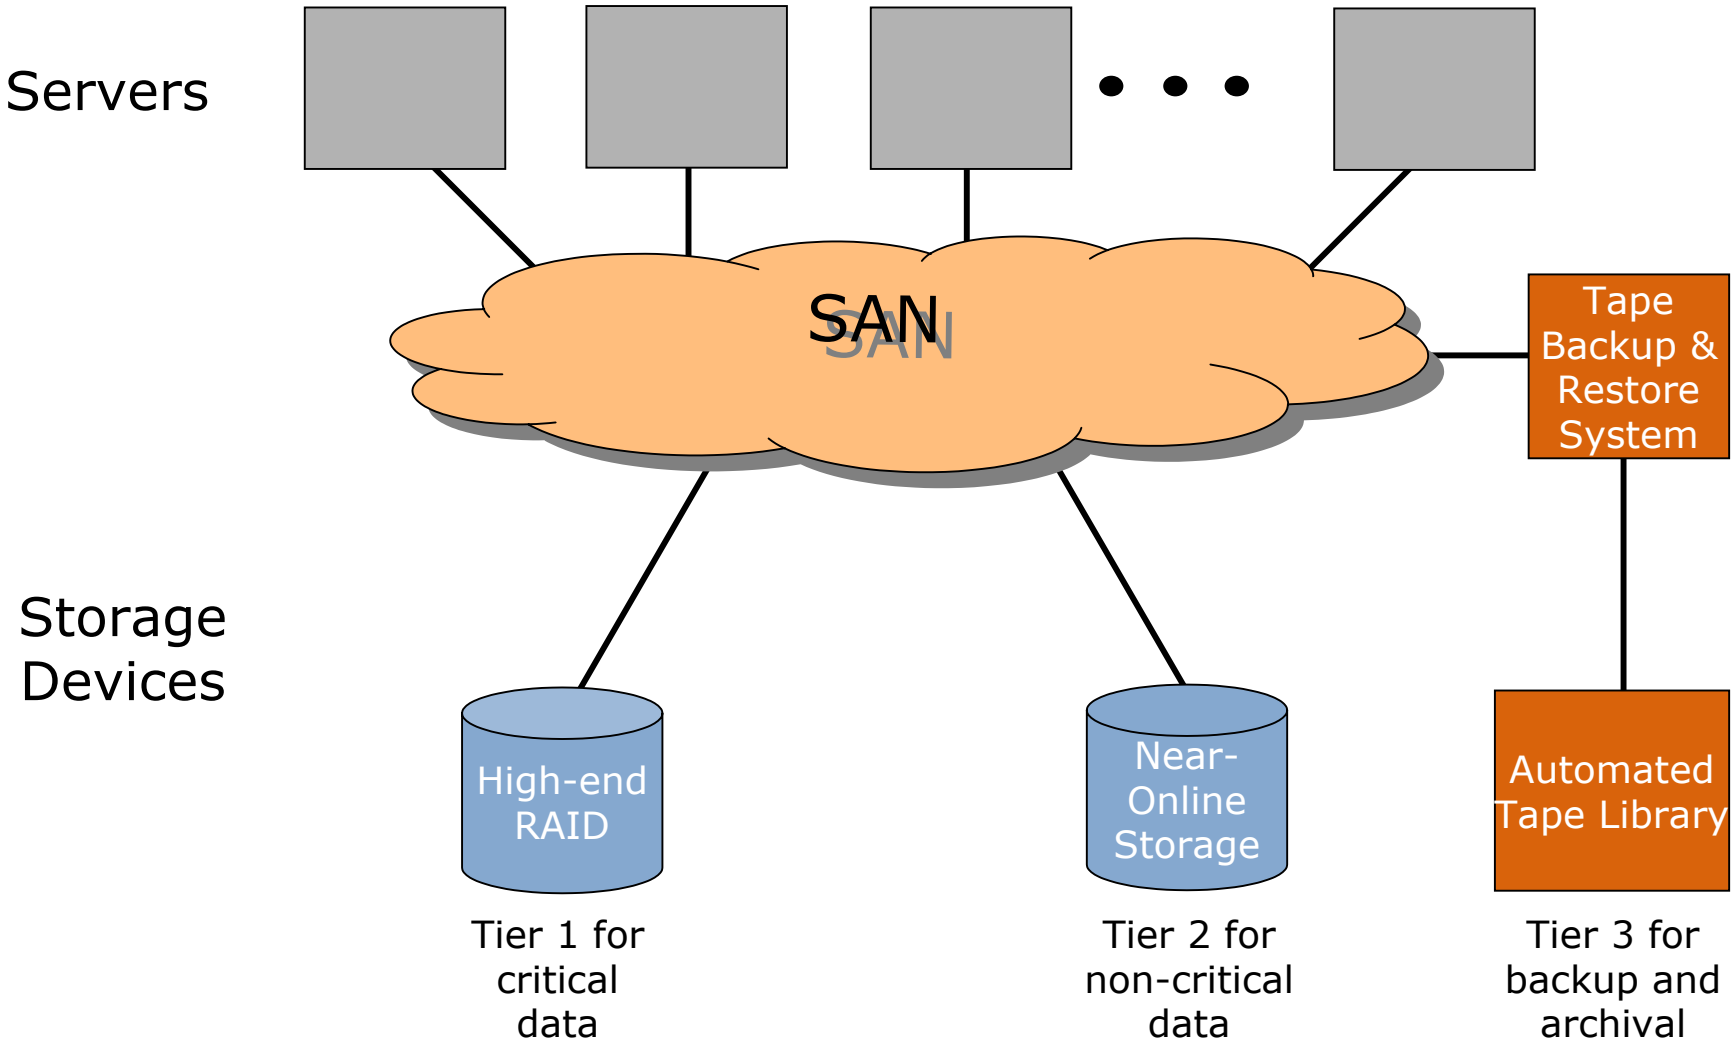

Return on Investment Example

10TB of High-end Array storage costing $600,000

50% of data requires high performance and availability

Add 5TB of Near-Online Array storage costing $50,000

5TB of non-critical data is migrated over

Recovered capacity	$300,000
Less Near-Online investment	$50,000
Immediate payback	$250,000

ROI = 500%

* does not include cost associated with installation, migration and administration

Managing, Migrating and Measuring

- Manual method
  - Individual device management
  - Periodic manual migration of data
  - Practical if storage is dedicated to specific applications
- Automated tools
  - HSM for policy based data movement between tiers
  - Archival software for pruning specific application records
  - Storage Virtualization to create separate pools or classes of storage from multiple array platforms
  - Backup software to move data from disk to disk to tape
  - SRM software to monitor, control, report and automate diverse storage resources

Multi Tiered Storage Solution Example

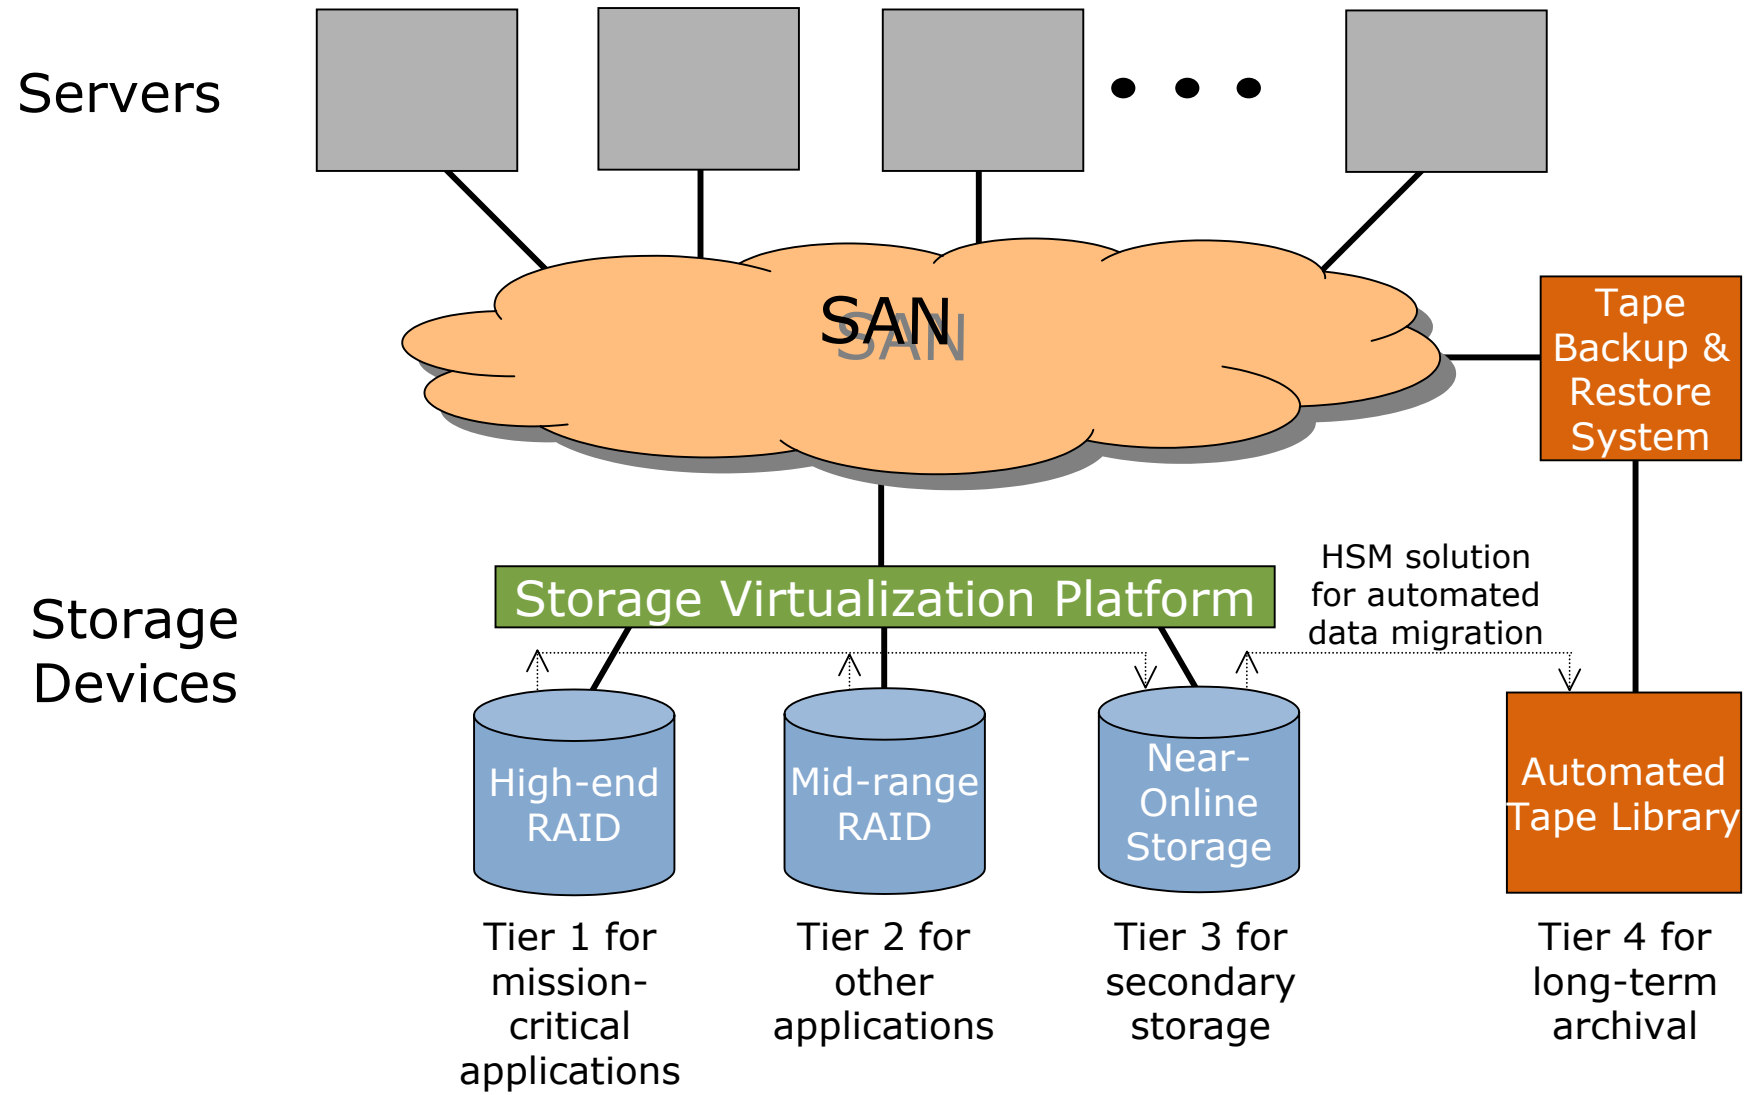

What hp has today

- SAN attach low cost storage using the MSA1000
- 6TB storage capacity per array at $.01 per/MB
- Qualified with HP Data Protector, Veritas NetBackup and Computer Associates BrightStor ARCserve
- Industry standard offer, non-proprietary APIs
- Solution blueprints
- Solution integration services

Disk-2-Disk Backup & Restore

Three scenarios:

1. Backup and incremental backup's to MSA drive array then backup off-line to tape or optical, on/off-site
2. Full backup to tape then incremental backup to MSA drive array
3. Backup and restore to/from MSA drive array only (not recommended for full data protection solution)

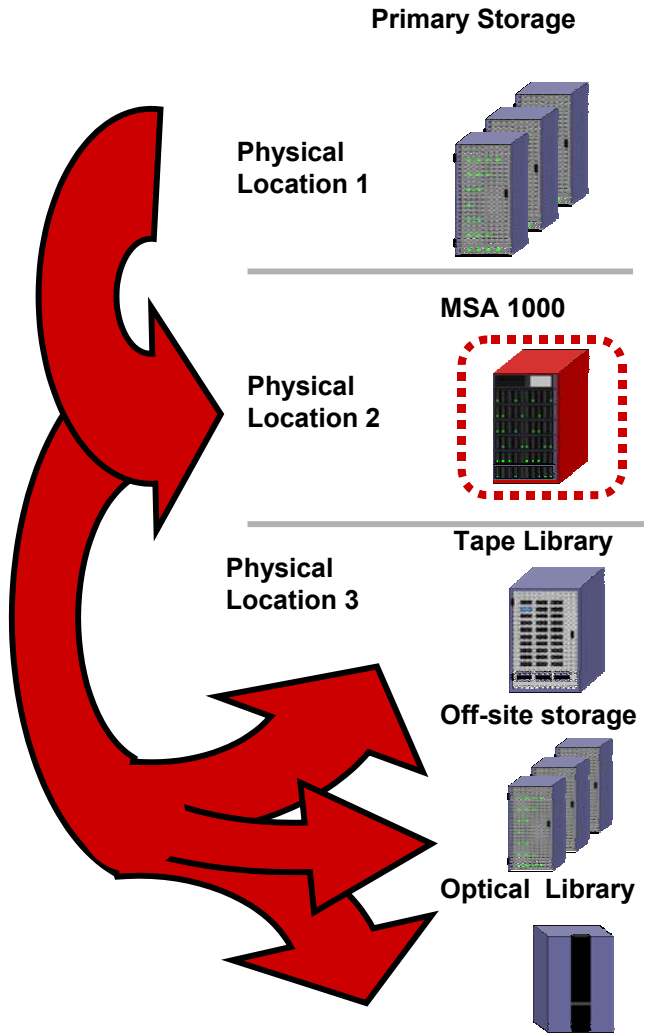

Disk-2-Disk-2-Tape Backup Solution Rack

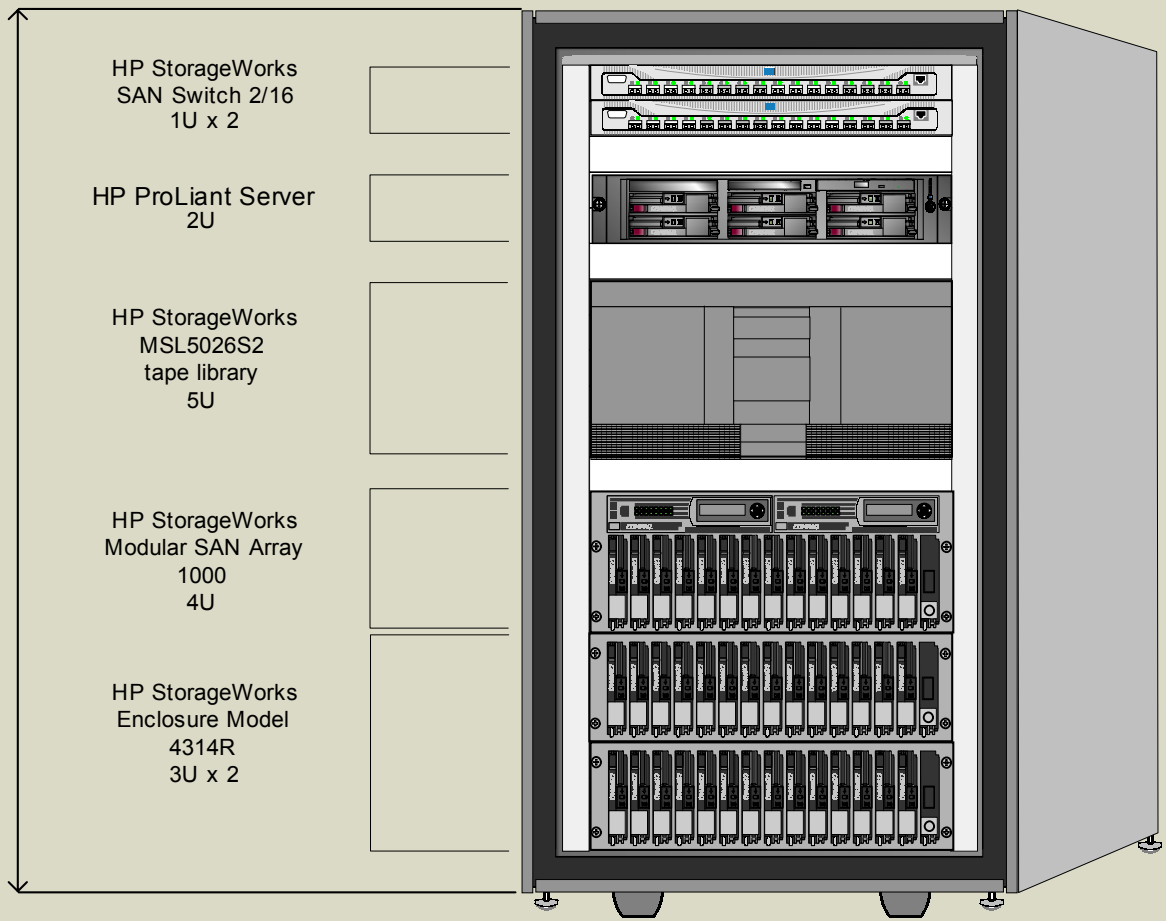

hp's OpenView Storage Management Software

- Storage Area Management
  - Storage Area Manager – centralized, enterprise-wide device, capacity and performance management
  - Storage Provisioner – establishes multiple storage offerings based on service level and storage attributes
- Information Lifecycle Management
  - Storage Data Protector – for enterprise data protection and disaster recovery
  - Storage Media Operations – automated tracking and management of offline storage media
- Virtualization Software
  - Continuous Access Storage Appliance – abstracts physical storage to create separate logical storage pools

Serial ATA Drives

hp's choice for the future

- Additive features to parallel ATA
  - Point-to-Point configuration eliminates bus sharing overhead
  - Roadmap starts at 150MB/sec
    - Future speeds of 300MB/sec
  - Smaller cable with 7 pin connector
  - One meter cable length
  - Hot-plug drive replacement
  - Command queuing (very basic compared to SCSI/FC)
  - Cyclical Redundancy Checking (CRC)

Data to Disk Segmentation

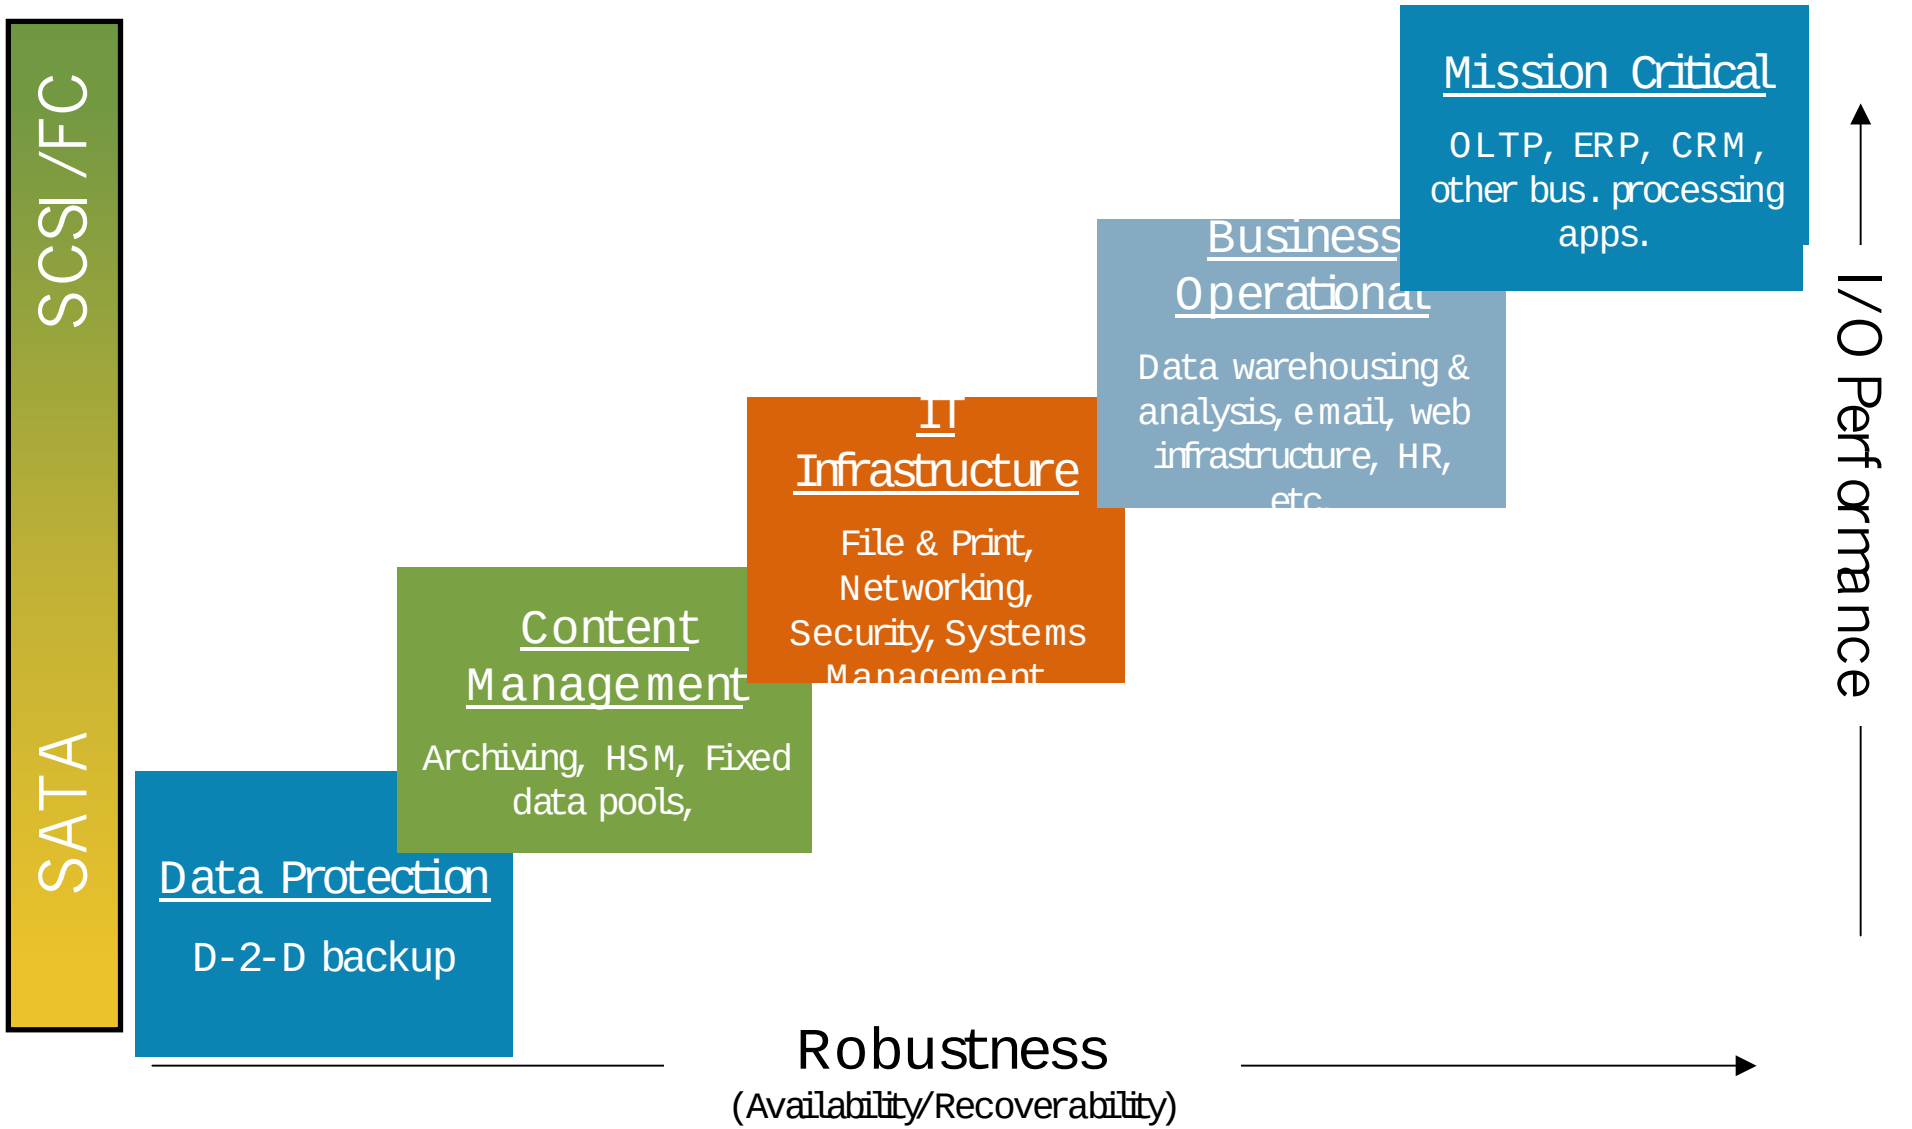

Future Storage Solutions

- MSA1000 will continue to add features and build upon a heterogeneous environment
- First half 04, hp will release a new network storage product utilizing SATA drive technology
- Software solutions, both in-house and through partnerships, are in development for email archiving and content management.
  - Legato partnership announced on June 2nd
  - More to come in the near future

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