Power-Managed Storage:
Gain Control of Your Growing Data
(Longer Data Life and Lower Energy Consumption)

Aloke Guha
CTO, COPAN Systems

LISA'06
December 8, 2006
Washington DC

Data Explosion

More new information over the next 2 years than in all previous history

93% of *NEW* content is born digital

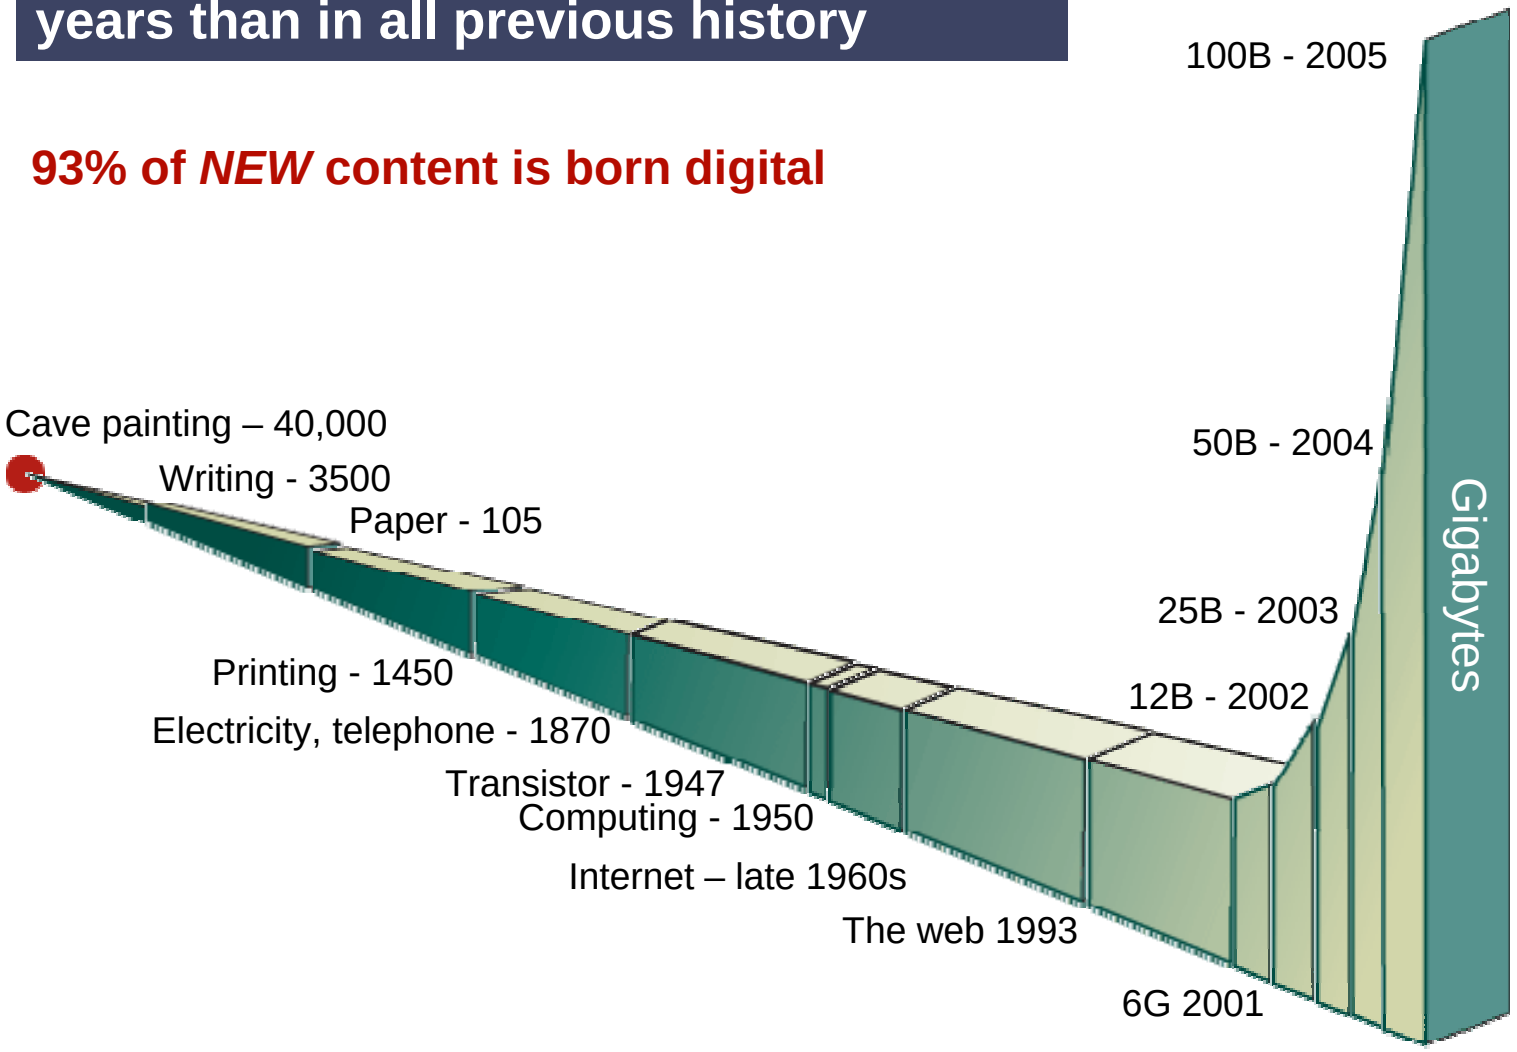

200 BILLION GIGABYTES OF INFO

Waves of Power

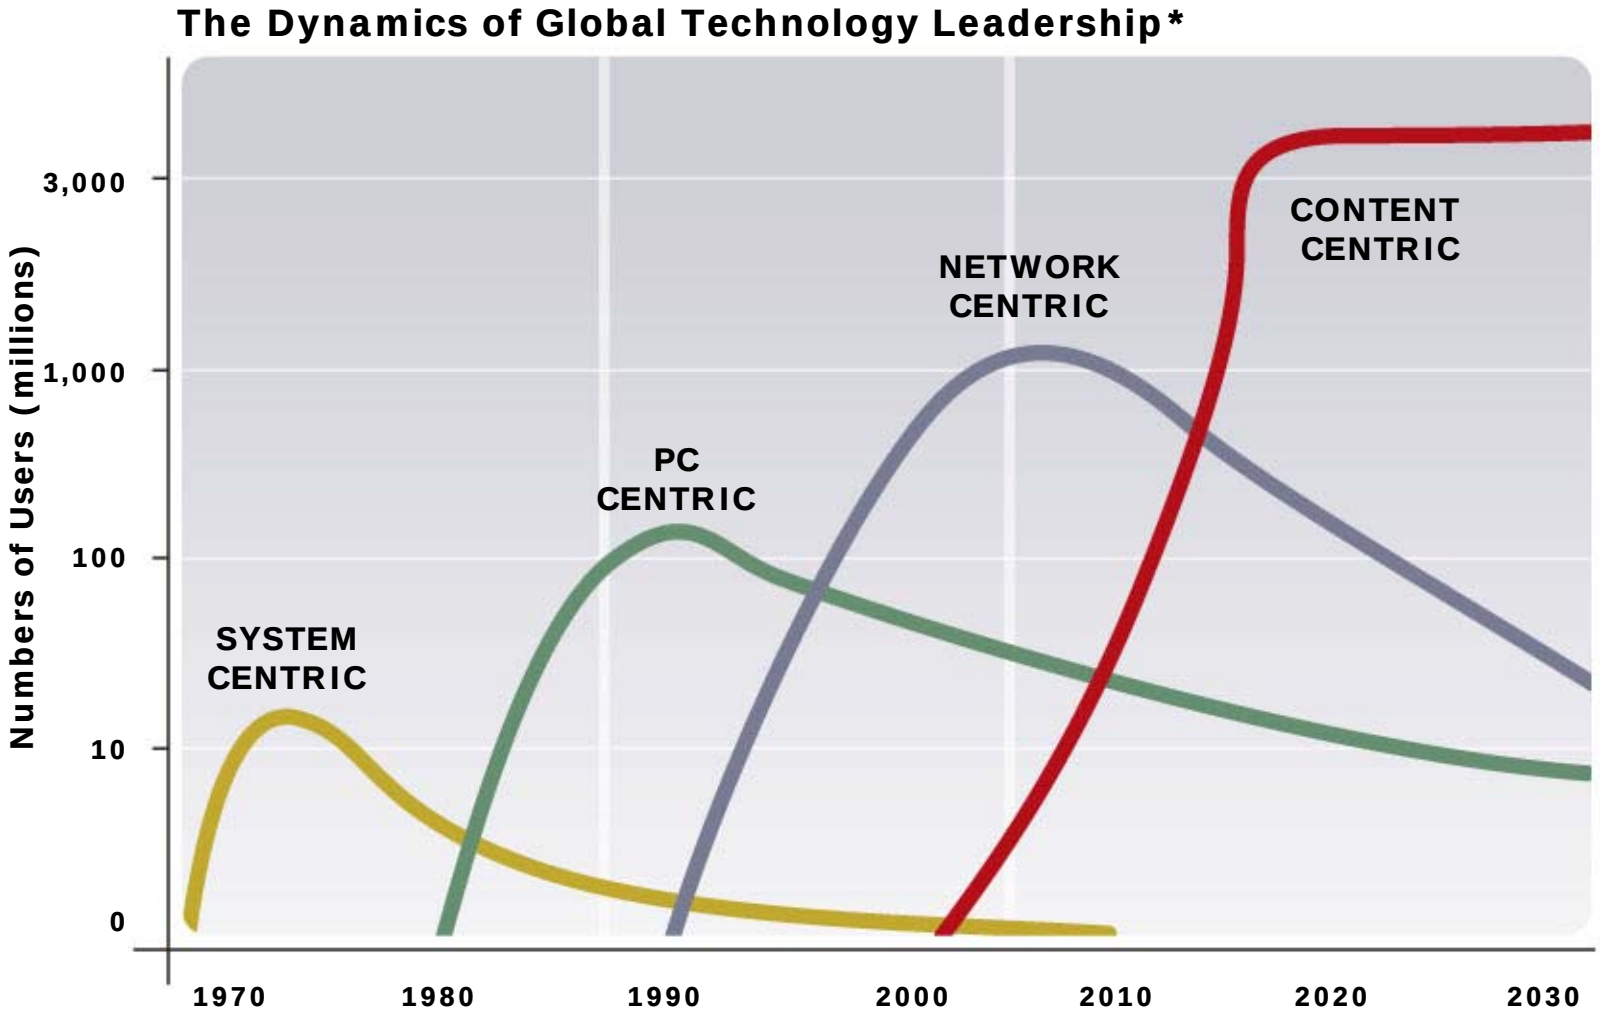

*David C. Moschella "Waves of Power: Dynamics of Global Technology Leadership 1964-2010," 1997

Waves of Storage Architecture

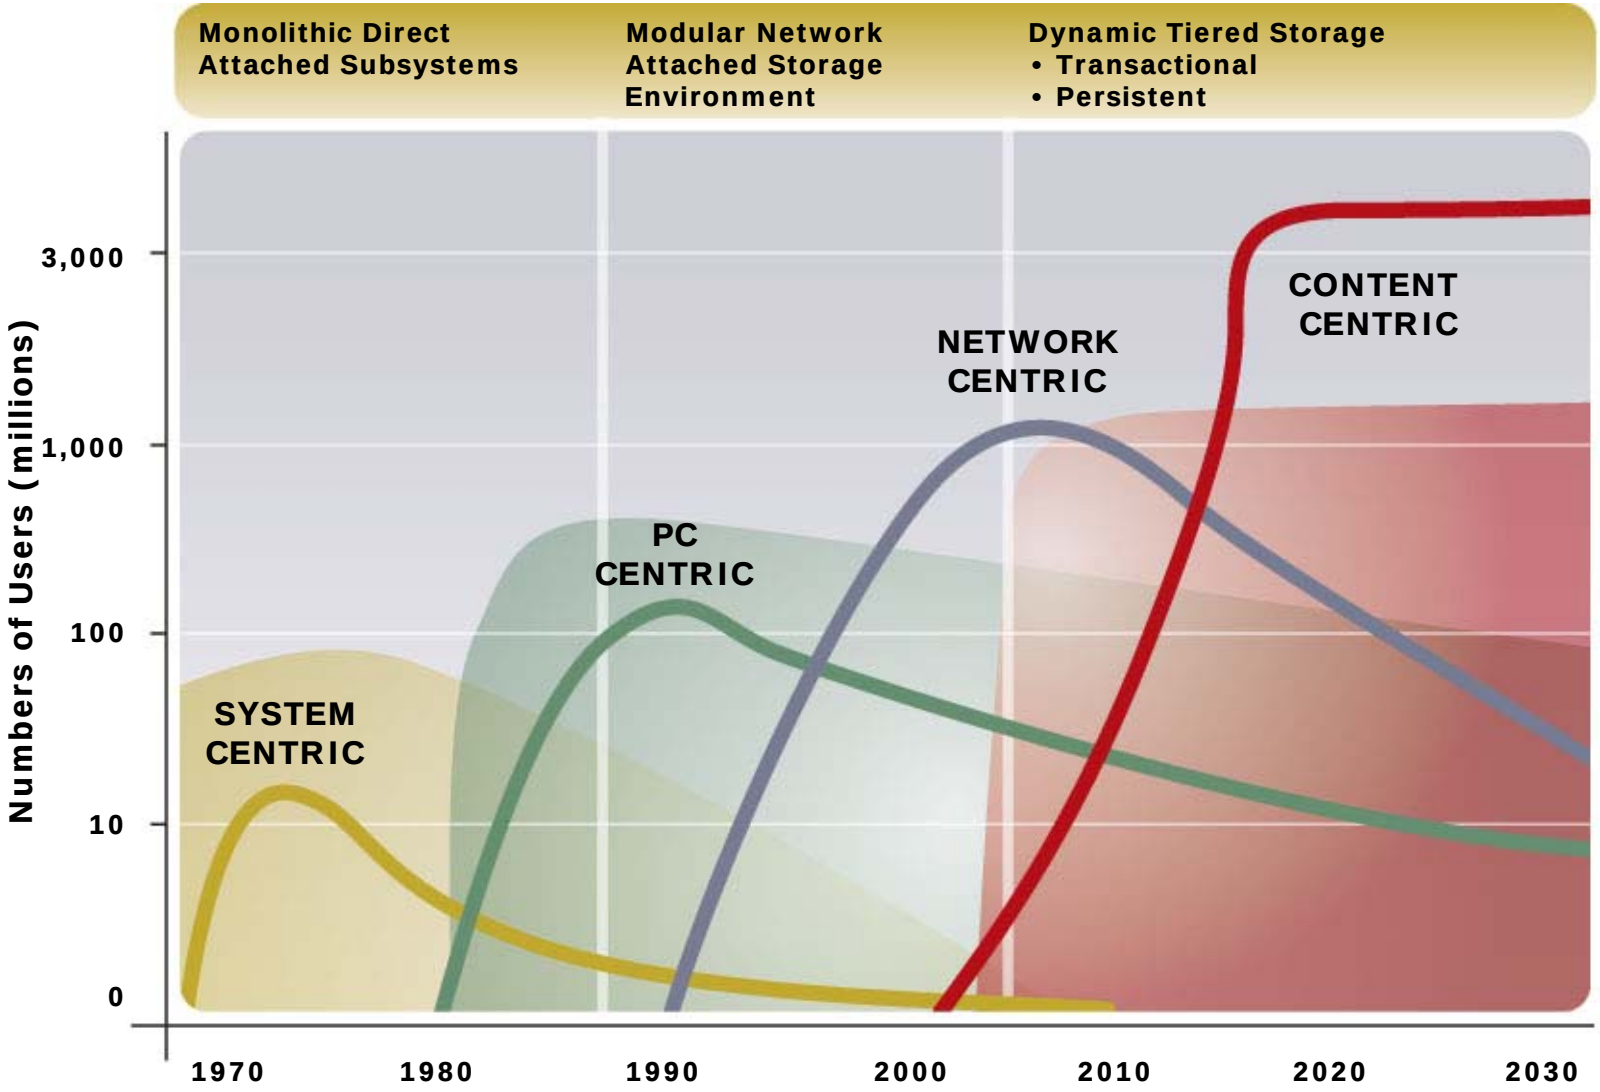

Data Center Storage Growth: Persistent Data

- **Persistent Data: *data that lives on long after it was first created***
  - Information captured or has aged to the point that it probably won't be changed or modified
  - Needs to be retained for future reference/access
- **Persistent Data (vs. Transaction Data)**
  - Long-term Retention (vs. Short Shelf Life)
  - Immutable (vs. Dynamic and Changing)
  - Usually High Data Rates (vs. IOPs)
  - Business Vital (vs. Business Critical)
  - Event-Driven (vs. Random Read/Write)
  - Data Integrity
  - Reference Content
  - Data Accumulation

Different Data–Different Needs

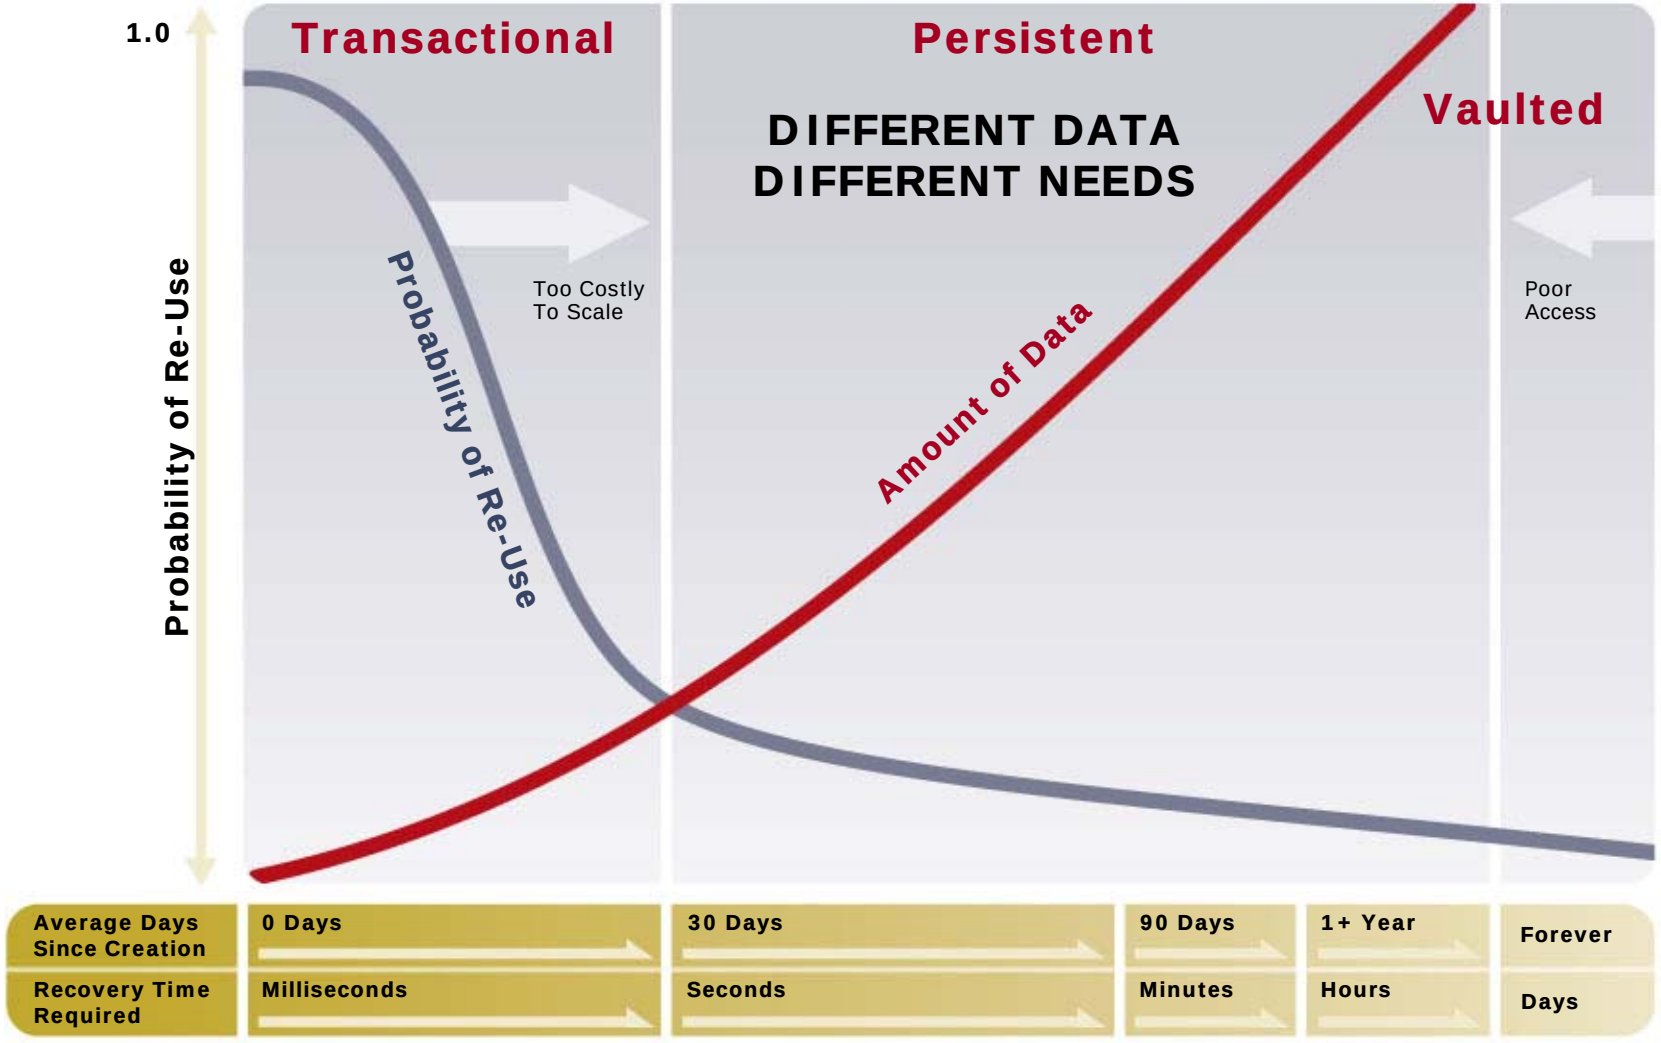

Transactional & Persistent Data

Transactional Data

Characteristics

- Highly Dynamic
- Short Shelf Life
- High IOPs
- Random Read/Write
- Information Capture & Creation
- Structured Data (mostly)
- Consistency Restrictions

Origin of Data

- OLTP
- Data Base Apps
- ERP
- Email (current)

Business Critical

Persistent data

Characteristics

- Immutable
- Long-term Retention
- Data Integrity
- Bandwidth Centric
- Event-Driven
- Reference Content
- Data Accumulation

Origin of Data

- Backup/Recovery
- Aged Transactional Data
- Unstructured Data Applications (mostly)
- Image Capture/Store
- Record/Document Management
- Archive

Data Protection

Data Migration

Archive Data

Business Vital

Challenges of Persistent Data

- **Large Volumes: typically consumes 80% or more of all storage in the data center**
- **Growth: more data is being retained on disk – tape is used for vaulting**
- **Cost: CAPEX**
- **Infrastructure Strain: OPEX**
  - Energy: power and cooling – beyond servers and growing
  - Reliability of long-term data – on traditional disk systems

  ⇒ Longevity – data migration/salvage/regeneration . . .

  ⇒ Footprint – data center floor space

Growing Energy Concerns

- **Data center energy consumption growing out of control**
  - ". . . in 5 years, energy costs will consume up to 40% of IT budgets" Gartner, Nov 2006
  - "33% of data centers expect to be out of power and cooling capacity by end of '07; 96% would be out of capacity by '11"
- **Data storage energy demands growing higher than that of servers**
  - Today costs are 20%-30% vs. over 30% for servers
  - Expect storage energy costs to be at parity with server energy costs in few years
  - Disk storage consumption growing exponentially
  - Power consumption growing linearly
- **Challenges in storage energy**
  - Data storage growth (esp. for persistent data) much higher than server growth
  - More data is persistent and is retained longer
    - 80% or more data retained for long-term: existing disk still consumes power
    - *More data needs to accessed online: regulatory, compliance, corporate governance*

Heat Density Trends

"*15% annual increase for servers,* DASD (disk systems) and workstations." (through 2010)

"7% annual increase for tape storage (through 2005)

The Uptime Institute Inc.: 2005-2010 Heat Density Trends in Data Processing, Computer Systems and Telecommunications Equipment (2006)

Long-Term Data Reliability Concerns

- **More persistent data kept online for accessibility, protection, integrity,**
  - Backup/DR, Archive, etc.
- **Most of the data are moved to lower-cost high capacity SATA drives**
  - SATA drives shipments to grow from 12.5% share in '04 to 52% of all TBs shipped in '09 (IDC June '05)
- **Retention period of data: 3 years to 7 years to 80 years to forever!**
  - Typ. warranty of disk systems < 3 yrs
  - Failure rates of SATA lower than Fibre Channel that has lower shelf life
- **Limited published or empirical data on SATA drive failure rates**

MTBF (hrs)	AFR (%)	Disk Specification
8,000,000	0.11%	
5,000,000	0.18%	
3,000,000	0.29%	
2,400,000	0.36%	
2,000,000	0.44%	
1,200,000	0.73%	Fibre Channel
1,000,000	0.87%	Fibre Channel
800,000	1.09%	
600,000	1.45%	SATA
400,000	2.17%	SATA
200,000	4.29%	
100,000	8.39%	

Current Tiered Storage Concept

Tier I	Tier II	Tier III	Tier IV	Tier V
Enterprise Disk	Modular Array	Object-Based File	Backup / Recovery	Off-Site DR Vaulting
Mission Critical High Transactional Support Storage	Business Critical Transactional Cost Storage	Low Cost SATA	Online Recovery	Disaster Recovery, Off-Site Data
				

Power-Managed Storage: MAID

- **What is MAID?**
  - Large number of power-managed disks
  - More than 50% of drives powered off
  - Power cycling by policy
  - Lower management and environmental costs, and longer drive life
- **Enhanced MAID for Long-Term Data**
  - Optimized: density and power – max 25% drives spinning
  - Three-Tier Architecture
    - Scales performance with capacity
  - POWER MANAGED RAID® software
    - RAID protection for power-managed disks
    - 1/3 cost of traditional disk
  - DISK AEROBICS® software
    - Disk reliability and data integrity

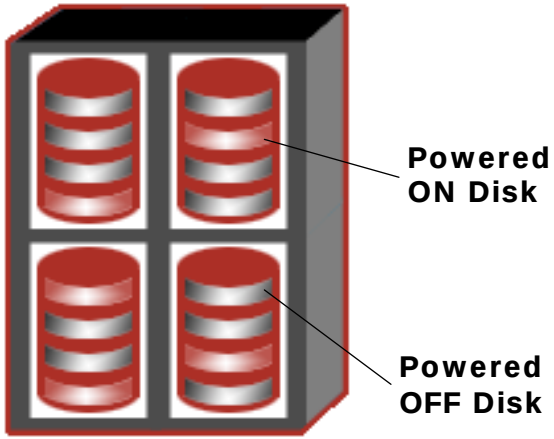

SNIA Definition

"A storage system comprising an array of disk drives that are powered down individually or in groups when not required. MAID storage systems reduce the power consumed by a storage array."

New Tiering Using Enhanced MAID

Tier I	Tier II	Tier III-IV	Tier V
Enterprise Disk	Modular Array	MAID (Massive Array of Idle Disks)	Off-Site DR Vaulting
Critical and Replicated Data	Lower Priority, Non-Replicated Data	File Reference Information Regulatory Data Online Recovery	Disaster Recovery, Off-Site Data

The New Data Center Tiering

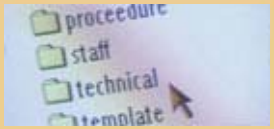

- Document Management
- E-mail

- Database
  - ERP
  - CRM
  - MRP

- PACS
- CAD/CAM

Transactional

Primary
- Tier 1
- Tier 2

Data Protection
- B/R: TSM
- B/R: NBU
- CDP

Data Replication
- IPStor

Email Archiving
- Symc EV

Database Migration
- Oracle ILM

File Migration
- VxFS, SAM-FS

Doc Mgmt/ECM

Persistent

Disaster Recovery

Vaulted

Example:
COPAN Enhanced MAID

Revolution 220A

Raw Storage Capacity	28 TB to 448 TB
Number of Shelves	1 to 8
Performance	Up to 5.2 TB per hour
Host Interface	Up to 4x 1-GigE
File Protocol supported	CIFS/NFS Access
Features	Automated Archiving, WORM, Retention, Versioning
Number of Files Stored	1 Billion
Power (448 TB)	4.9 KW

Revolution 220T/220TX

Raw Storage Capacity	28 TB - 448 TB
Number of Shelves	1 to 8
Performance	Up to 5.2 TB per hour
Host Interface	Up to 8 2-Gbps FC
Number of Emulated Tape Libraries	Up to 56
Number of Virtual Tape Cartridges	Up to 8192
Power (448 TB)	4.9 KW

http://www.copansys.com

*: COPAN's MAID platform provides multiple personalities: Virtual Tape Library, File and Disk Block

COPAN SYSTEMS

Reducing Storage Energy

Growing Energy Concerns

- **Data center energy consumption is growing out of control**
  - ". . . in 5 years, energy costs will consume up to 40% of IT budgets" Gartner, Nov 2006
  - "33% of data centers expect to be out of power and cooling capacity by end of '07; 96% would be out of capacity by '11"
- **Data storage energy demands growing higher than that of servers**
  - Today costs are 20%-30% vs. over 30% for servers
  - Expect storage energy costs to be at parity with server energy costs in few years
  - Disk storage consumption growing exponentially
  - Power consumption growing linearly (areal density increase)
- **Challenges in storage energy**
  - Data storage growth esp. for persistent data is much higher than server
  - More data is persistent and is retained longer
    - 80% or more data retained for long-term: existing disk still consumes power
    - *More data needs to accessed online: regulatory, compliance, corporate governance*

Heat Density Trends

"*15% annual increase for servers,* DASD (disk systems) and workstations." (through 2010)

"7% annual increase for tape storage (through 2005)

The Uptime Institute Inc.: 2005-2010 Heat Density Trends in Data Processing, Computer Systems and Telecommunications Equipment (2006)

Power Density and Storage Density Comparison

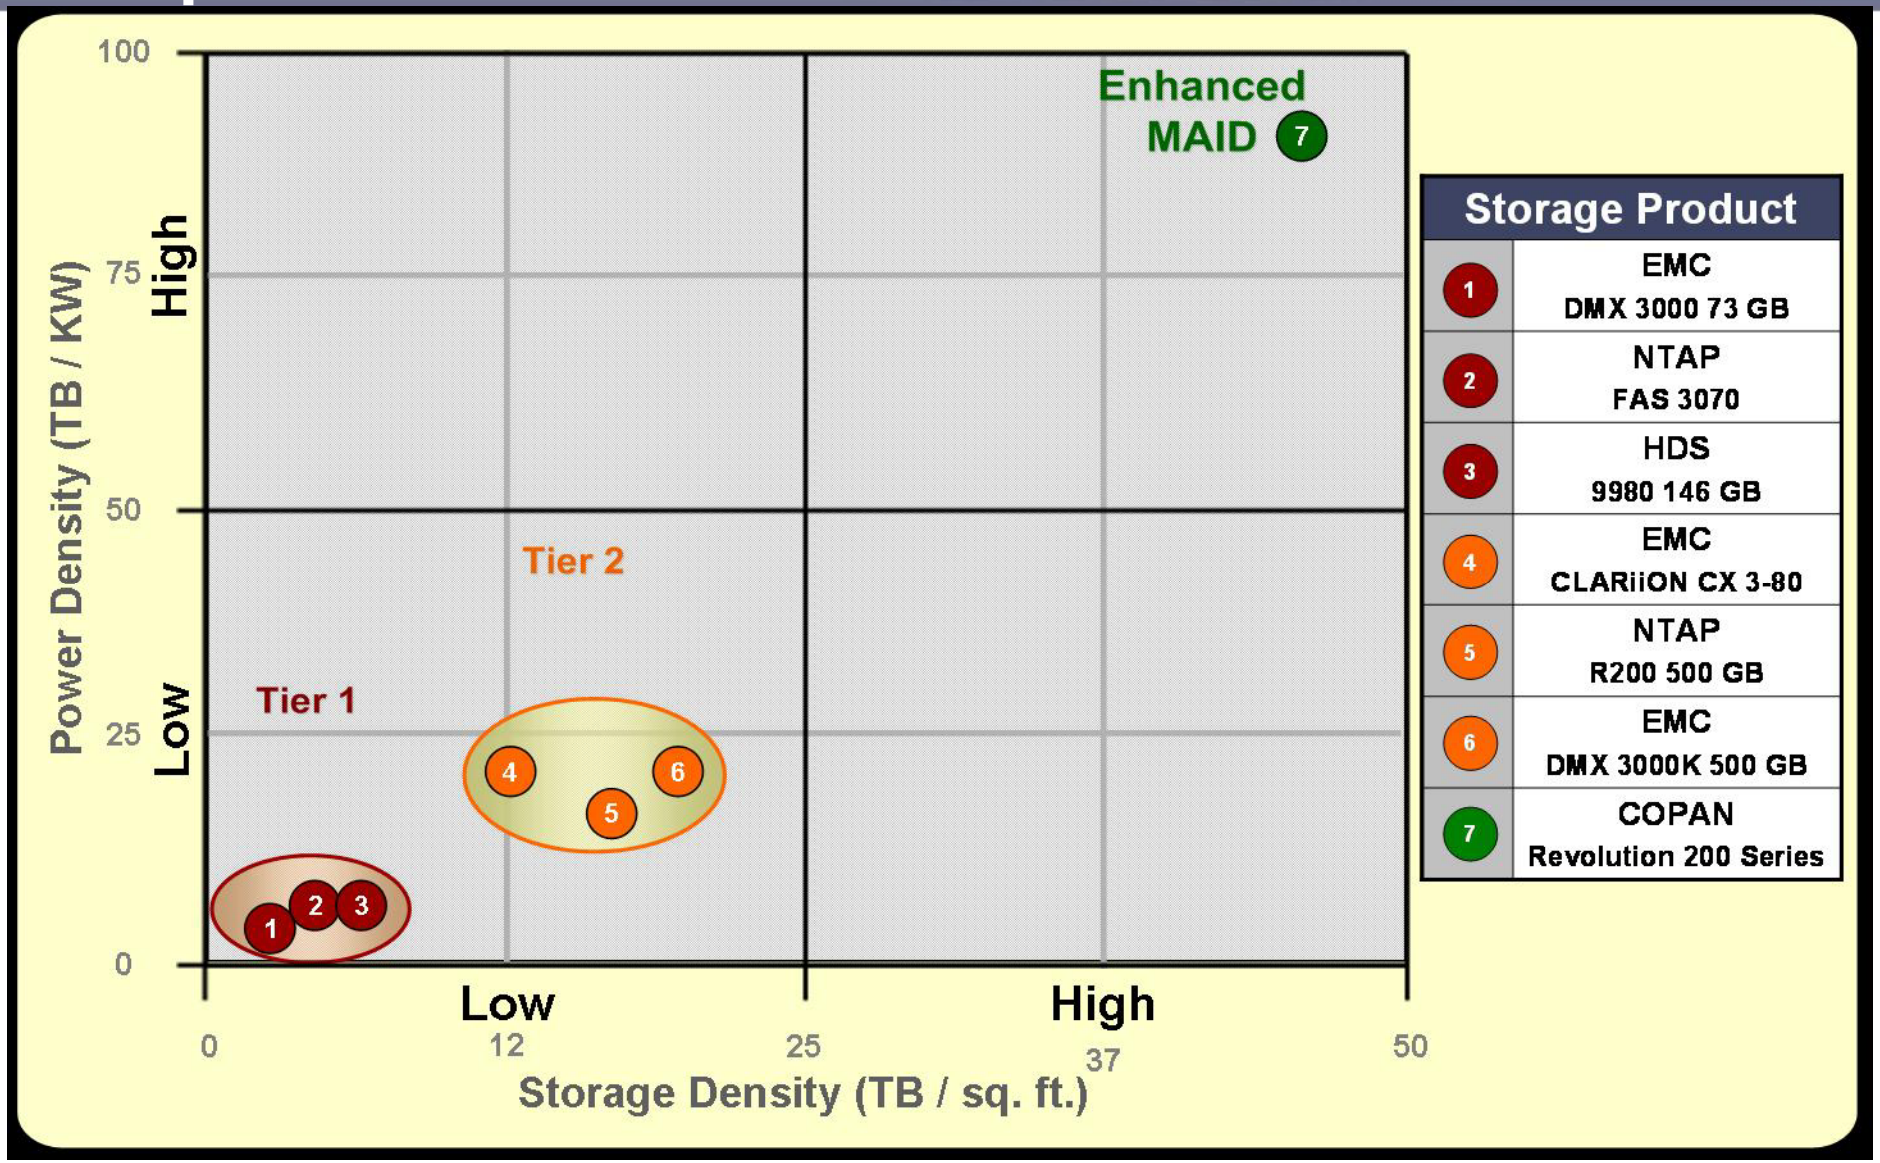

Based on vendors' currently published specifications

Average Fortune 1000 Data Center

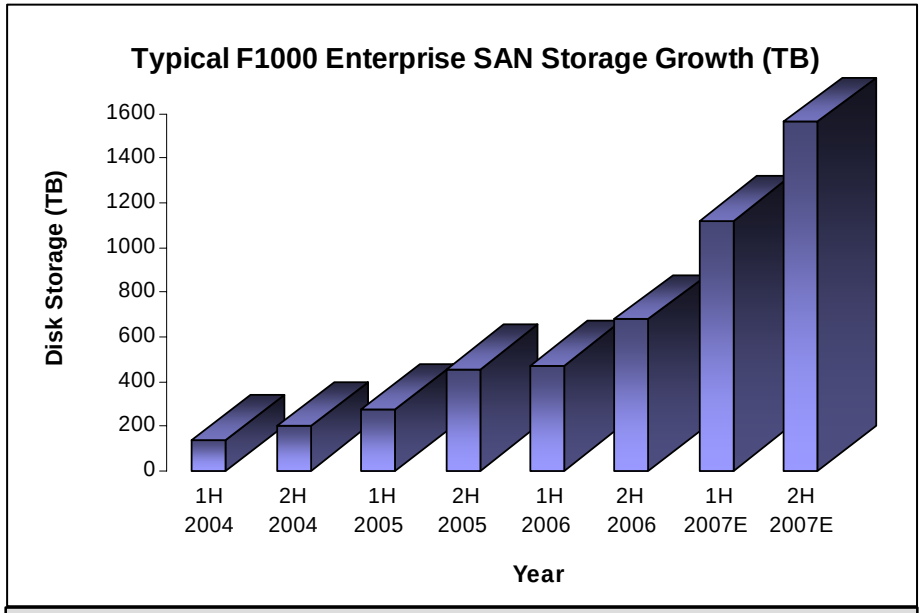

Average F1000 data stored on disk
(Source: *The InfoPro, Storage Survey Wave 8, 10/06*)

- Ave. F1000 data center has **680 TB** disk
  - **75% of this data is persistent**
- Typical mix of F1000 disk storage
  - **60% Fibre Channel**
  - **40% SATA**
- Ave. disk capacity growth rate: **230%**
  - Stored data is doubling every 10 months since '05
  - At this rate, storage for average F1000 company will grow from 680TB in '06 to over **8PB** by '09
- 1 year of power + cooling costs for **8 PB = $10,000,000**
  - Cost for traditional disk systems
  - 60%/40% FC to SATA mix

Enhanced MAID vs. Traditional FC and SATA Storage Products

Tier 1 storage (typ. Fibre Channel) and Tier 2 storage (typ. SATA) have low storage density and high power consumption compared to COPAN's enhanced MAID

	COPAN	Tier 1	Tier 2	Percentage Improvement
Power Density (TB / KW)	91.40	4.08	17.44	2140% / 424%
Storage Density (TB / sq. ft.)	44.80	3.62	13.68	1138% / 228%

COPAN enhanced MAID platform:

- **4x to 30x better in power density than traditional Tier 1 / 2 storage**
- **4x to 20x better in storage density than traditional Tier 1 / 2 storage**

- The power consumption advantage of enhanced MAID systems is in its power density and can be measured in terms of capacity per unit of power, or TB/KW.
- Enhanced MAID products maintain extremely low power requirements combined with massive density, which brings its power density to over **90 TB/KW**.

Ref.: A. Guha, Solving the Energy Crisis in the Data Center Using COPAN Systems' Enhanced MAID Storage, COPAN Systems White Paper, November 2006.

Migrate Your Persistent Data to COPAN's Enhanced MAID

Scenario 1: 73 GB Fibre Channel drives and 500 GB SATA			
Storage Mix	**Before Migrating Data**	**50% Migration to MAID**	**75% Migration to MAID**
Fibre Channel	60%	25%	12.5%
SATA	40%	25%	12.5%
Enhanced MAID	0%	50%	75%

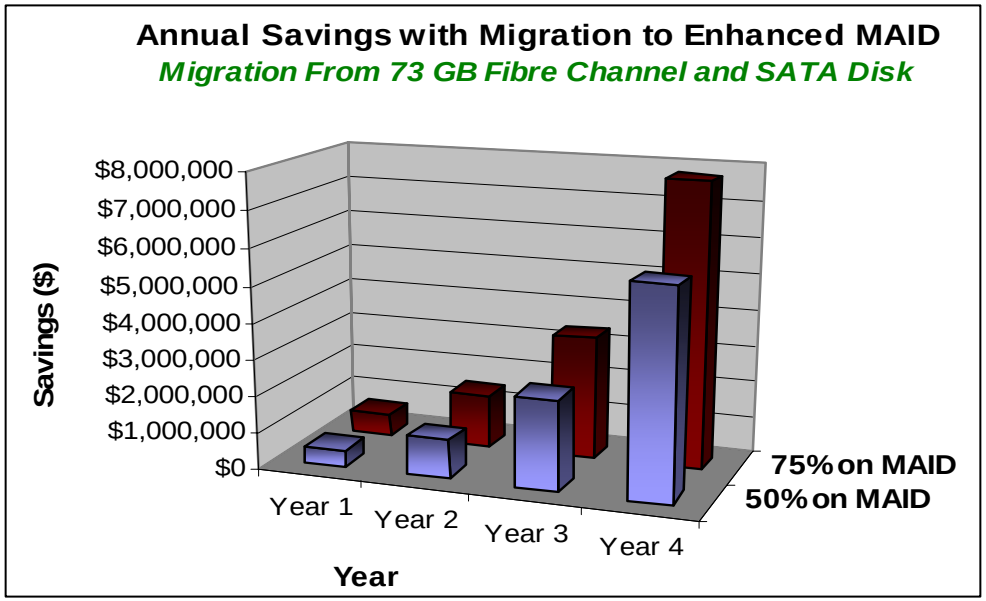

Potential savings of typical F1000 from migrating persistent data to COPAN's enhanced MAID

Assuming that 50% to 75% of all data currently residing on Fibre Channel and SATA disk arrays is persistent and can be migrated to enhanced MAID*:

If 50% of the data were migrated to enhanced MAID

- **First year annual savings = $470,000**
- **Fourth year annual savings = *$5,700,000***

If 75% of the data were migrated to enhanced MAID

- **First year annual savings = $640,000**
- **Fourth year annual savings = $7,800,000**

*Based on average F1000 SAN-attached disk and $0.11 per KW-hour

Migrate Your Persistent Data to COPAN's Enhanced MAID

Scenario 2: 146 GB Fibre Channel drives and 500 GB SATA			
Storage Mix	**Before Migrating Data**	**50% Migration to MAID**	**75% Migration to MAID**
Fibre Channel	60%	25%	12.5%
SATA	40%	25%	12.5%
Enhanced MAID	0%	50%	75%

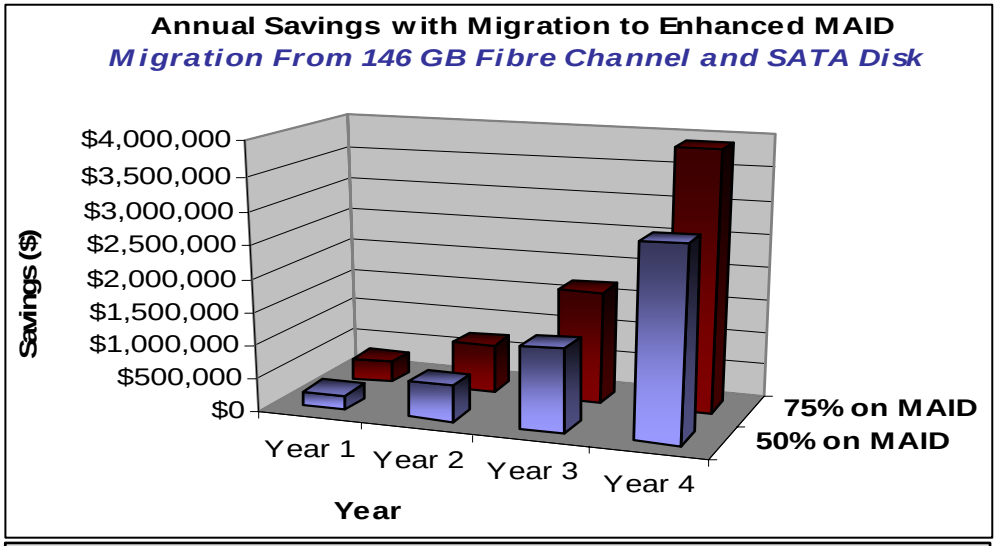

Potential savings of typical F1000 from migrating persistent data to COPAN's enhanced MAID

Assuming that 50% to 75% of all data currently residing on Fibre Channel and SATA disk arrays is persistent and can be migrated to enhanced MAID*:

If 50% of the data were migrated to enhanced MAID

- **First year annual savings = $235,000**
- **Fourth year annual savings = *$2,800,000***

If 75% of the data were migrated to enhanced MAID

- **First year annual savings = $322,000**
- **Fourth year annual savings = $3,900,000**

*Based on average F1000 SAN-attached disk and $0.11 per KW-hour

COPAN SYSTEMS

Improving Long-Term Storage Reliability

Disk Failure Modes

- **Multiple factors affect disk failure rates**

Factor	Controlling Entity	Effect on Annual Failure Rate (AFR)
Disk Technology: HDD platform maturity, etc.	Drive Vendor: Recording Technology, etc.	Lower AFR expected with more mature quality HDD Platform
Number of Platters per Drive	Drive Vendor: Areal Density, HDD Platform	More platters/drive increases AFR at higher POH
Power Duty Cycle (POH)	Storage System Vendor: Target Applications	Increased POH and Power Duty Cycle increases AFR
Ambient Temperature	Storage System Vendor: Packaging, Cooling	Higher ambient and operating temperatures increase AFR
Vibration and Drive Handling	Storage System Vendor: Drive Packaging,	Increased Vibration and Shocks to Drive increases AFR
I/O Workload	Application/User	Increased random I/O Access increases AFR

Ref.: A. Guha and S. Ouderkirk, "Disk Failure Rates and Implications of Enhanced MAID Storage Systems, COPAN Systems White Paper, April 2006.

Enhanced MAID Reliability

POWER-MANAGED RAID® Software

- Maximum of 25% of all drives are powered ON
- All data is RAID protected independent of their power state
- Contributes to increasing the service life of the drives

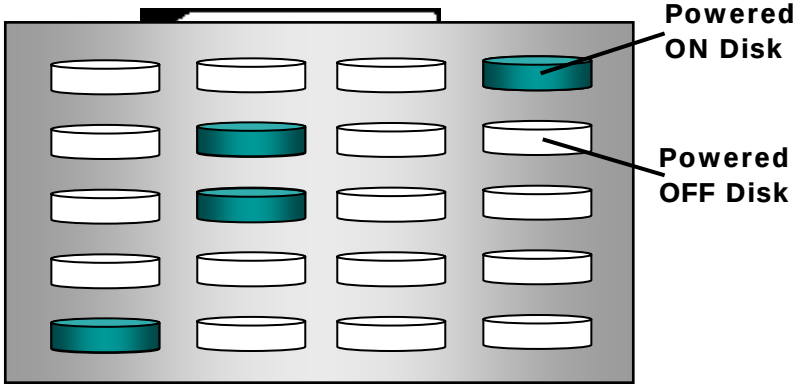

Benefits
▪ **Increased drive service life**
▪ **Reduced disk failures/time**
▪ **Power and cooling per TB**
▪ **The less the data is accessed, higher the service life and lower the power needs**

Enhanced MAID Benefits: Longevity

DISK AEROBICS® Software

Provides continuous data verification by exercising and testing idle drives at least once every 30 days to ensure data integrity and drive health

Benefits

- **Assures drive health and data integrity**
- **Predictive drive maintenance**
- **Avoids lengthy RAID rebuilds**
- **COPAN Enhanced MAID platform designed for 10+ years: reduces painful data migration and downtime**

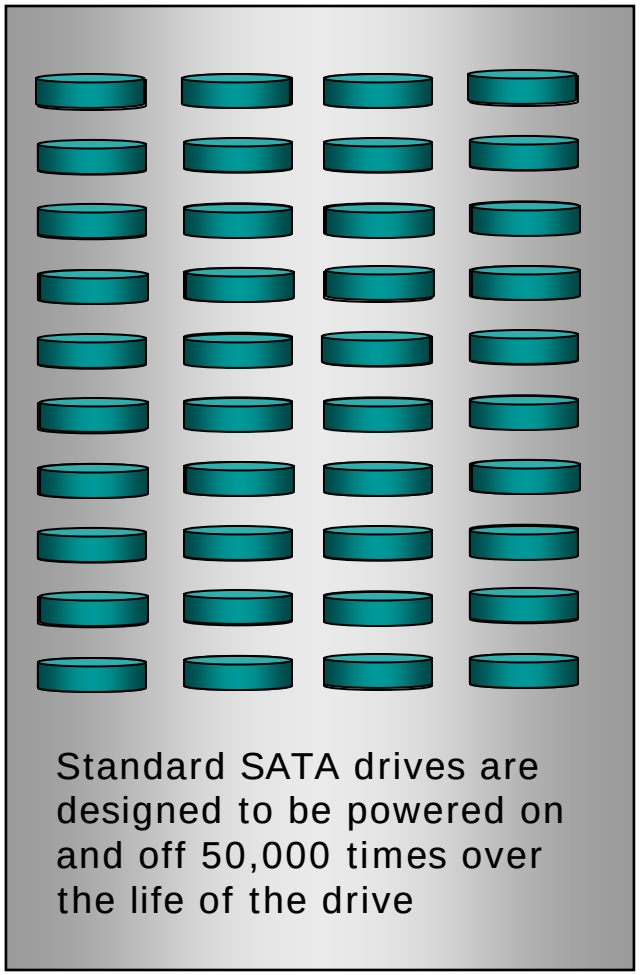

Standard SATA drives are designed to be powered on and off 50,000 times over the life of the drive

www.seagate.com/support/service/ - Reliability section - start-stop cycles specs at 50,000

DISK AEROBICS Software is a registered COPAN Systems trademark

DISK AEROBICS Software Features

- **Designed to disk failures and improve data integrity**

DISK AEROBICS™ Feature	Benefit	Mechanism
Proactive monitoring and management of drives	Predictive Drive Maintenance	Monitors SMART parameters and environmental data
Periodic Exercise of Idle Drives	Assures Drive Health	Every idle drive powered-on and tested at least once every 30 days
Power Management	Extends Drive Life	Drives spin only when necessary to meet application requirements
Disk Scrubbing	Assures Data Integrity	Background task that identifies any bad sectors on disk and copies data to new sector on drive
Data Migration	Avoids lengthy RAID rebuilds and data loss for long-term data.	Proactive failing of suspect drives – copies all data to spare drive and "fails-out" suspect drive. Inserts new drive into RAID set.

Drive Reliability Field Data

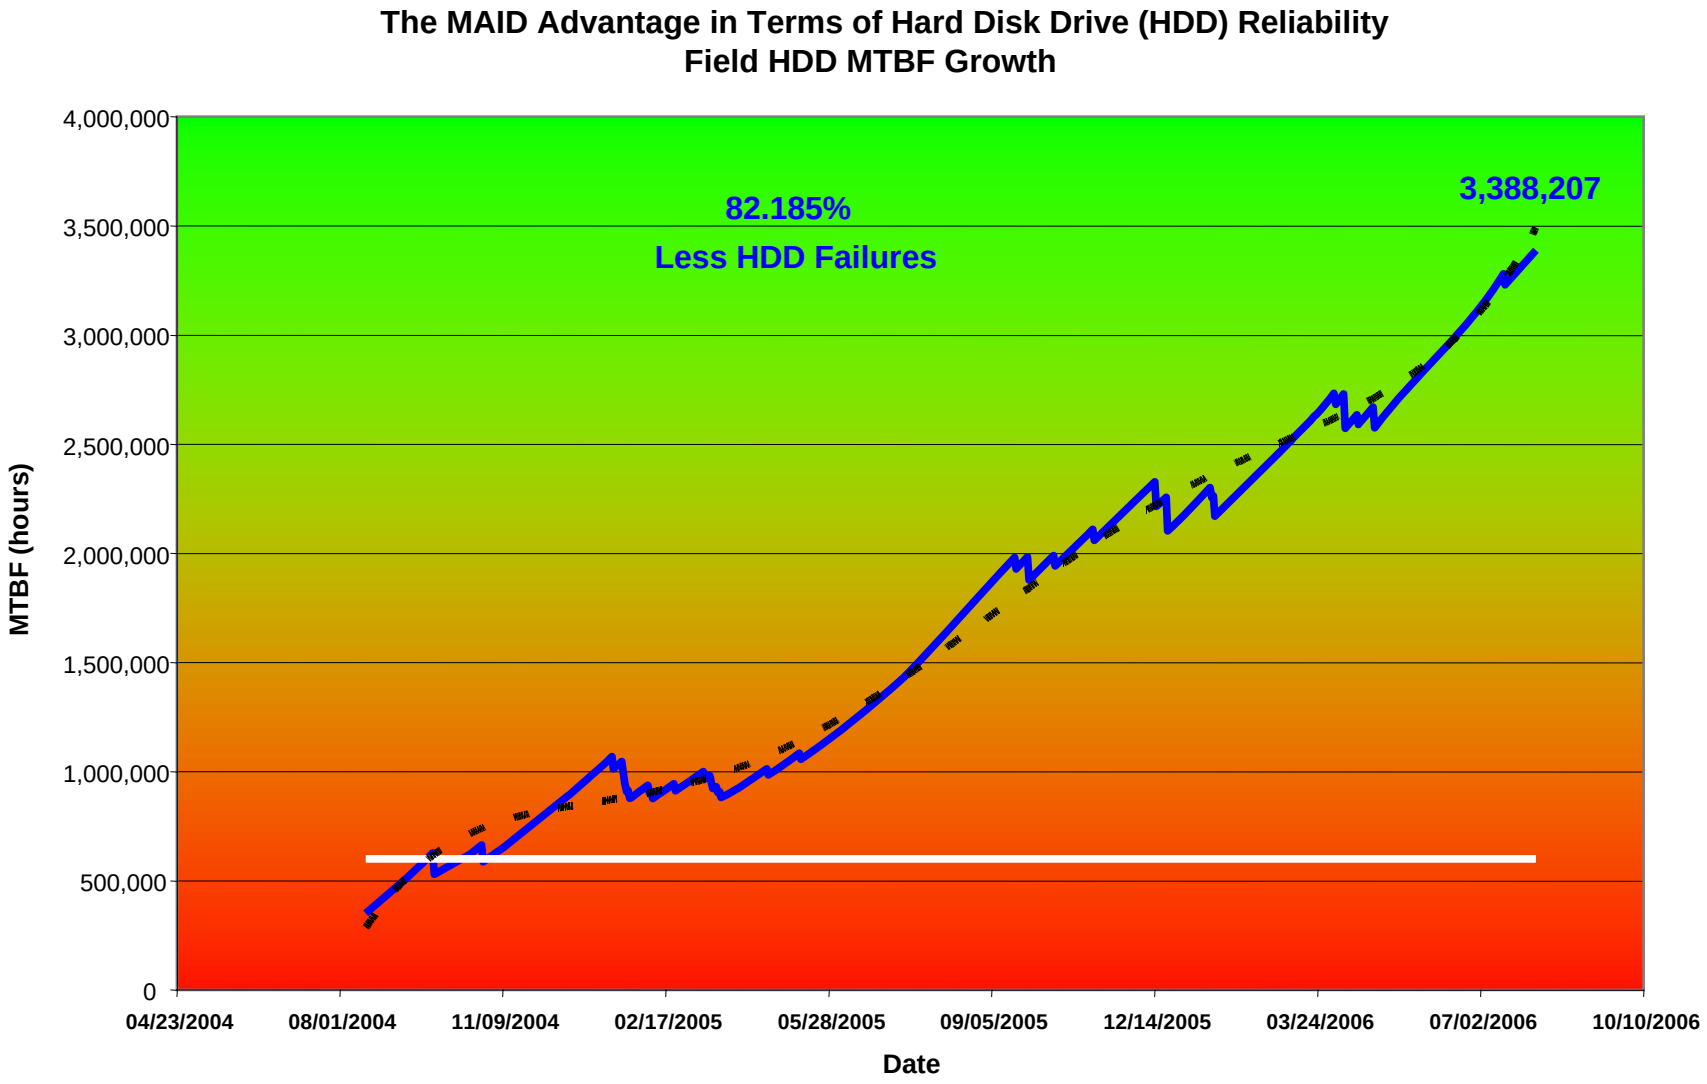

Drive Life and Reliability

- **POWER MANAGED RAID® software and DISK AEROBICS® software increase drive and data reliability**
  - Compared to standard SATA disk, enhanced MAID has < 1/5th failure rate
  - Field MTBF: more than 5X SATA disks, almost 3X FC disks
  - Service Life: expect almost 5X
- **Disk Reliability and TCO benefits**
  - Per 1000 drives, expect only 3 drives to fail/yr vs. 15 drives with std disk systems
    - Enhanced MAID: 0.26% failures/yr
    - SATA: 1.45% failures/yr
  - Competitors have
    - ~5X drive replacements
    - 17 touches versus 1 touch for COPAN
- **Significantly higher data reliability**

MTBF (hrs)	AFR (%)	Disk Specification
8,000,000	0.11%	
5,000,000	0.18%	
3,388,207	0.26%	COPAN - Aug 2006
3,000,000	0.29%	
2,400,000	0.36%	
2,000,000	0.44%	
1,200,000	0.73%	Fibre Channel
1,000,000	0.87%	Fibre Channel
800,000	1.09%	
600,000	1.45%	SATA
400,000	2.17%	SATA
200,000	4.29%	
100,000	8.39%	

600K hrs = 68 yrs
3.39M hrs = 387 yrs

Enhanced MAID Reliability: Summary

- **Failure Rate: Enhanced MAID has less than 1/5th failure rate of traditional storage systems and improves and drive service life by more than 5X**
- **Data Reliability: Enhanced MAID platform provides significantly greater data reliability – 32X standard SATA and 8X over Fibre Channel**
- **Lower TCO: traditional disk storage will have 17 system touches vs 1 scheduled touch for drive replacement**
- **Longevity: derive longer data life as required for persistent data**

MAID for Persistent Data

- Enhanced MAID: ideal Persistent Data Storage Platform
- Architecture and MAID OS tuned for Persistent Data
- Integrated storage services for Persistent Data application needs

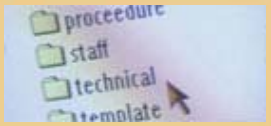

- Document Management
- E-mail

- Database
  - ERP
  - CRM
  - MRP

- PACS
- CAD/CAM

Data Management

Transactional

Primary
- Tier 1
- Tier 2

- Performance > 20%
- Reliability > 5X
- Footprint Svgs > 72%
- Power Svgs >75%
- TCO Svgs > 50%

Vaulted

Data Management

Enhanced MAID Impact

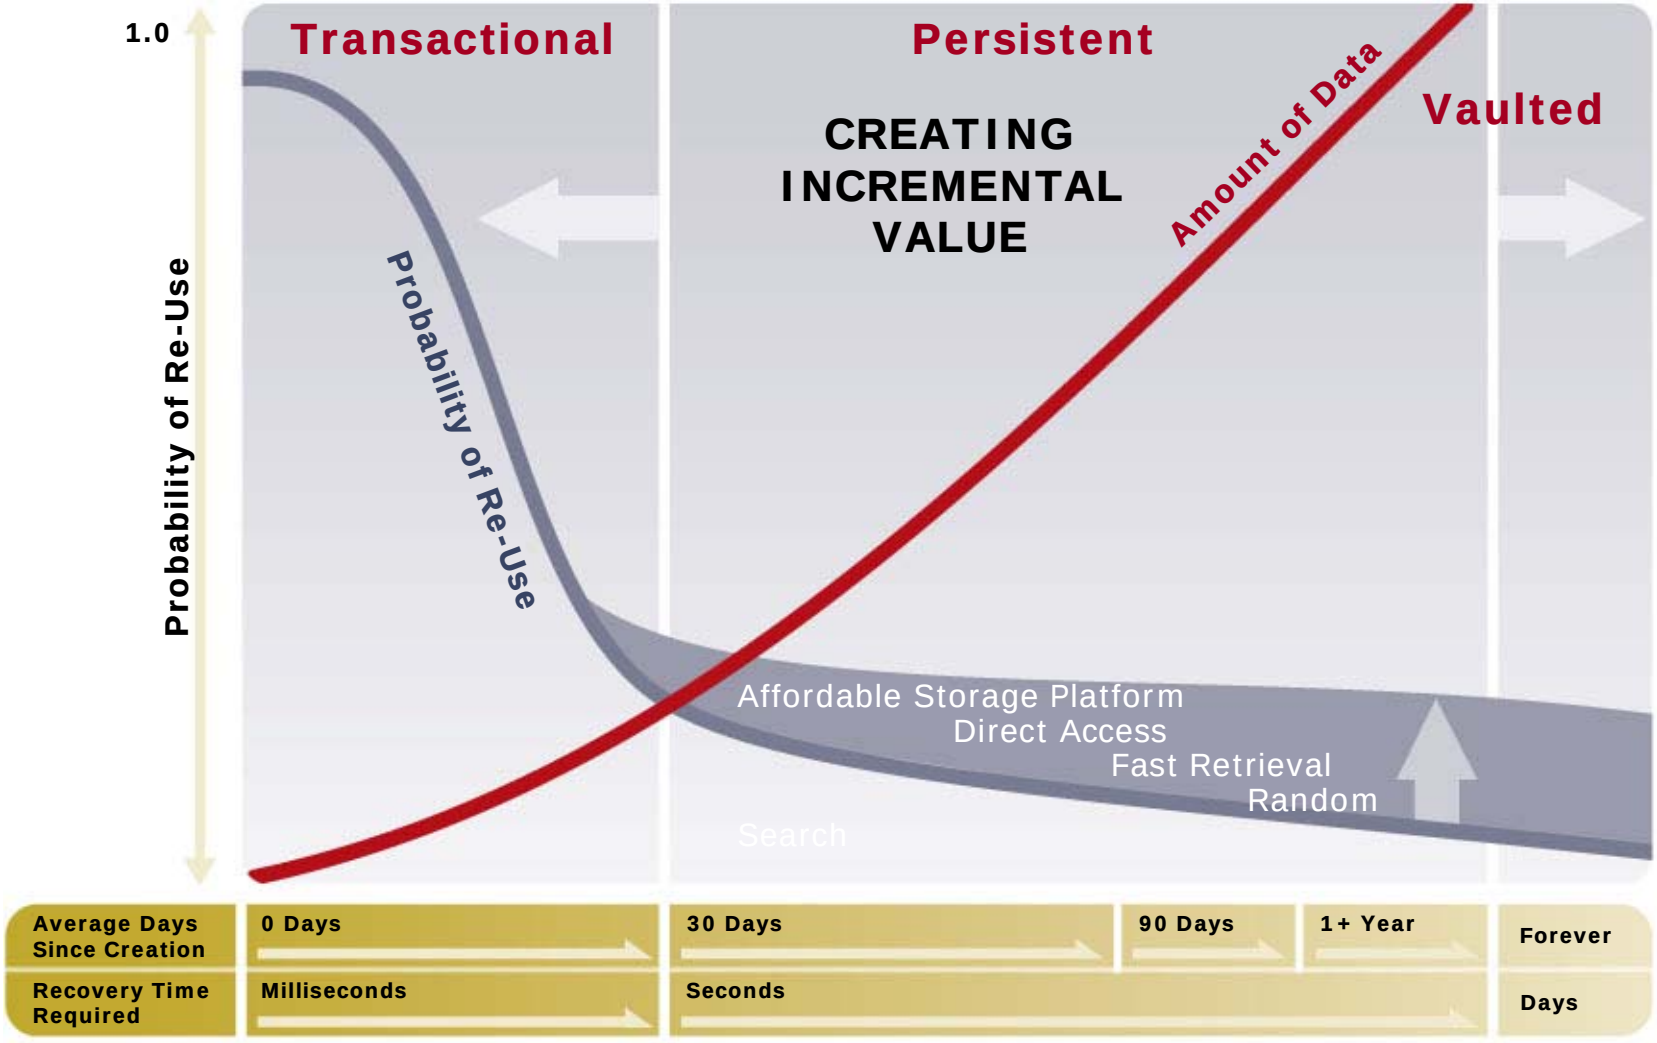

Conclusions

- **Growth of retained data creating many challenges: energy, reliability, footprint . . .**
- **Persistent Data has different access needs not met by traditional disk systems**
- **Potential benefits of Enhanced MAID in F1000 data center**
  - Energy Cost savings: > $500K/1st yr to $Ms later years
  - Reliability Improvement: Disks by 5X, Data by 32X
  - Footprint Savings: more than 80%
  - Other benefits: Performance, Longevity, Cost . . .
- **What can you do?**
  - Be proactive: identify your persistent data and move it to MAID

COPAN SYSTEMS

COPAN Systems provides intelligent storage solutions that unlock the value of your long-term persistent data.

Thank You

aloke.guha@copansys.com