

IT Infrastructure Matters

# Technical Computing Data Management

Matt Drahzal

Technical Computing Strategy



# Managing Change at the Heart of the Problem

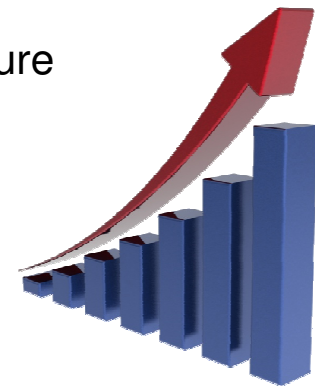
- **Managing Latency**

- As we have more and more data – information is too far from analysis
  - Keep Data as Close to Processing as Possible?
  - Keep Processing as Close to Data as Possible?
  - Both!



- **Managing Storage Evolution**

- Storage Devices are evolving quickly as new Tiers are added
- Customers asking to
  - Integrate with, extend, and enhance our current storage infrastructure
  - Evolve our architecture as technology evolves
  - Include new storage technologies as our needs change



- **Managing Collaboration and Globalization**

- How can we share data globally?
- How can we use collaboration to increase efficiency?
- How are we approaching the growing Technical Computing user base?

- **Reunification:** Breaking Down the Walls Between Compute and Persistent Storage
  - Distance to Data *counts*!
- **End of the Line for coming for traditional RAID 6**
  - Drive Size and Drive Count are now just too high
- **Abstraction:** Levels the Playing Field for all types of Storage
  - Nobody needs to ever know the TAR command line again
  - Storage Devices become invisible
- **Technology:**
  - Servers Getting More Powerful, and Less Expensive
  - Solid State is *finally* here

# Where is the Bottleneck?

In the past 10 years:

CPU speed performance increased ~8-10x

DRAM speed performance increased ~7-9x

Network speed performance increased ~100x

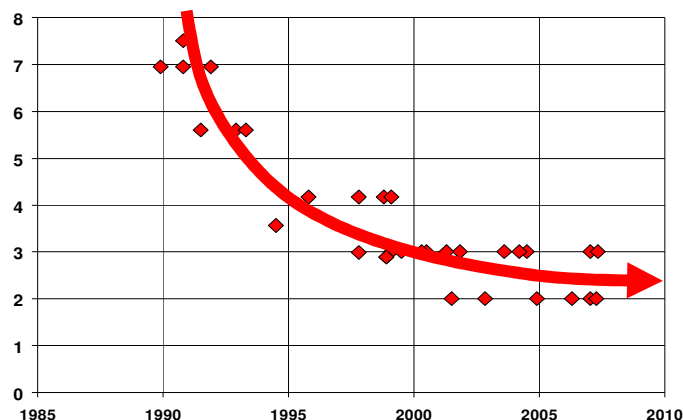
Bus speed performance increased ~20x

Disk speed performance increased **ONLY 1.2x**

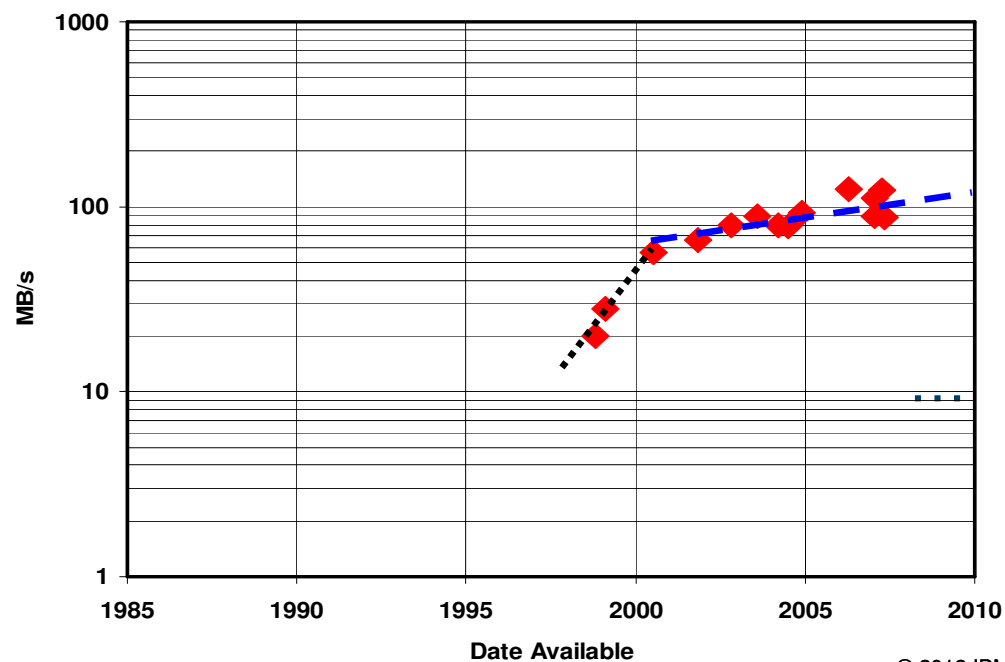
*Storage disk speed is the bottleneck  
that's slowing everything else in the IT stack*

# HDD Latency and Disk Transfer Speed – Still Little progress...

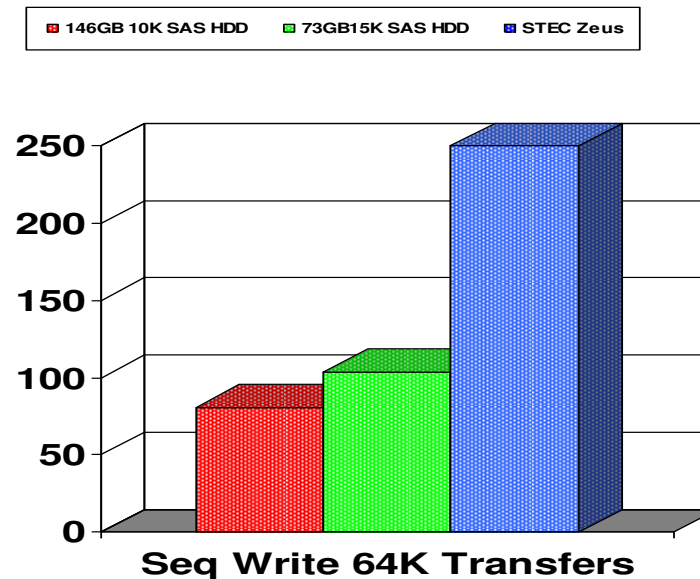
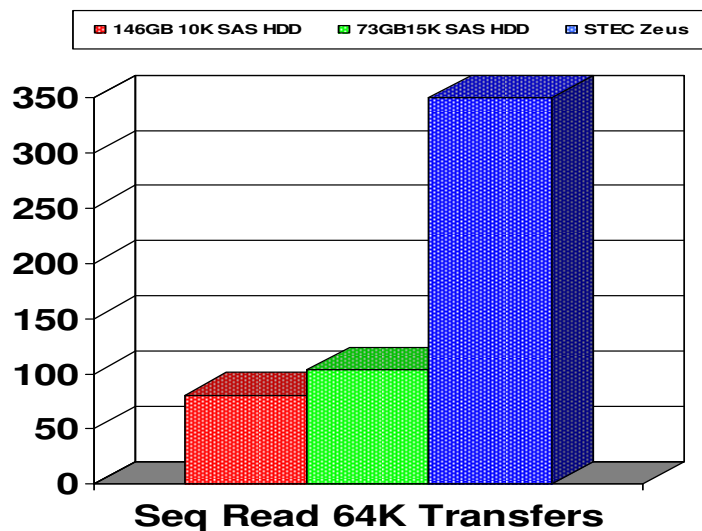
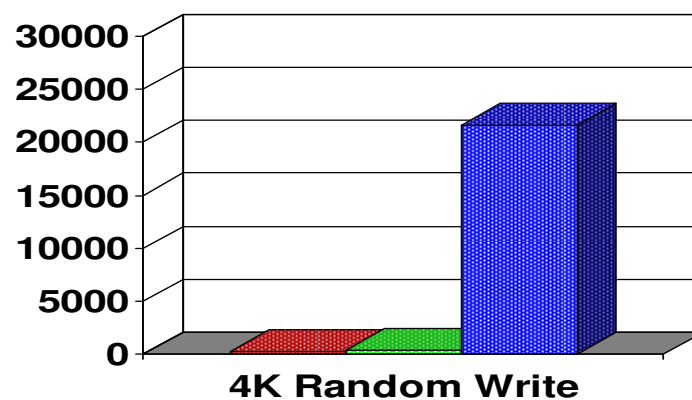
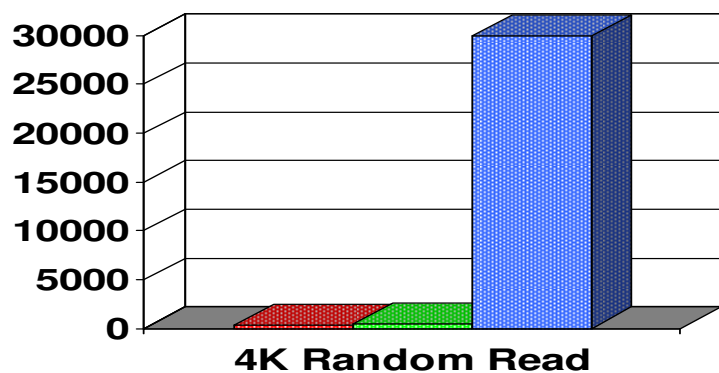
Latency  
[ms]



Individual Disk  
Transfer Speed  
[MB/s]



# SSDs Performance vs HDDs – Still True!



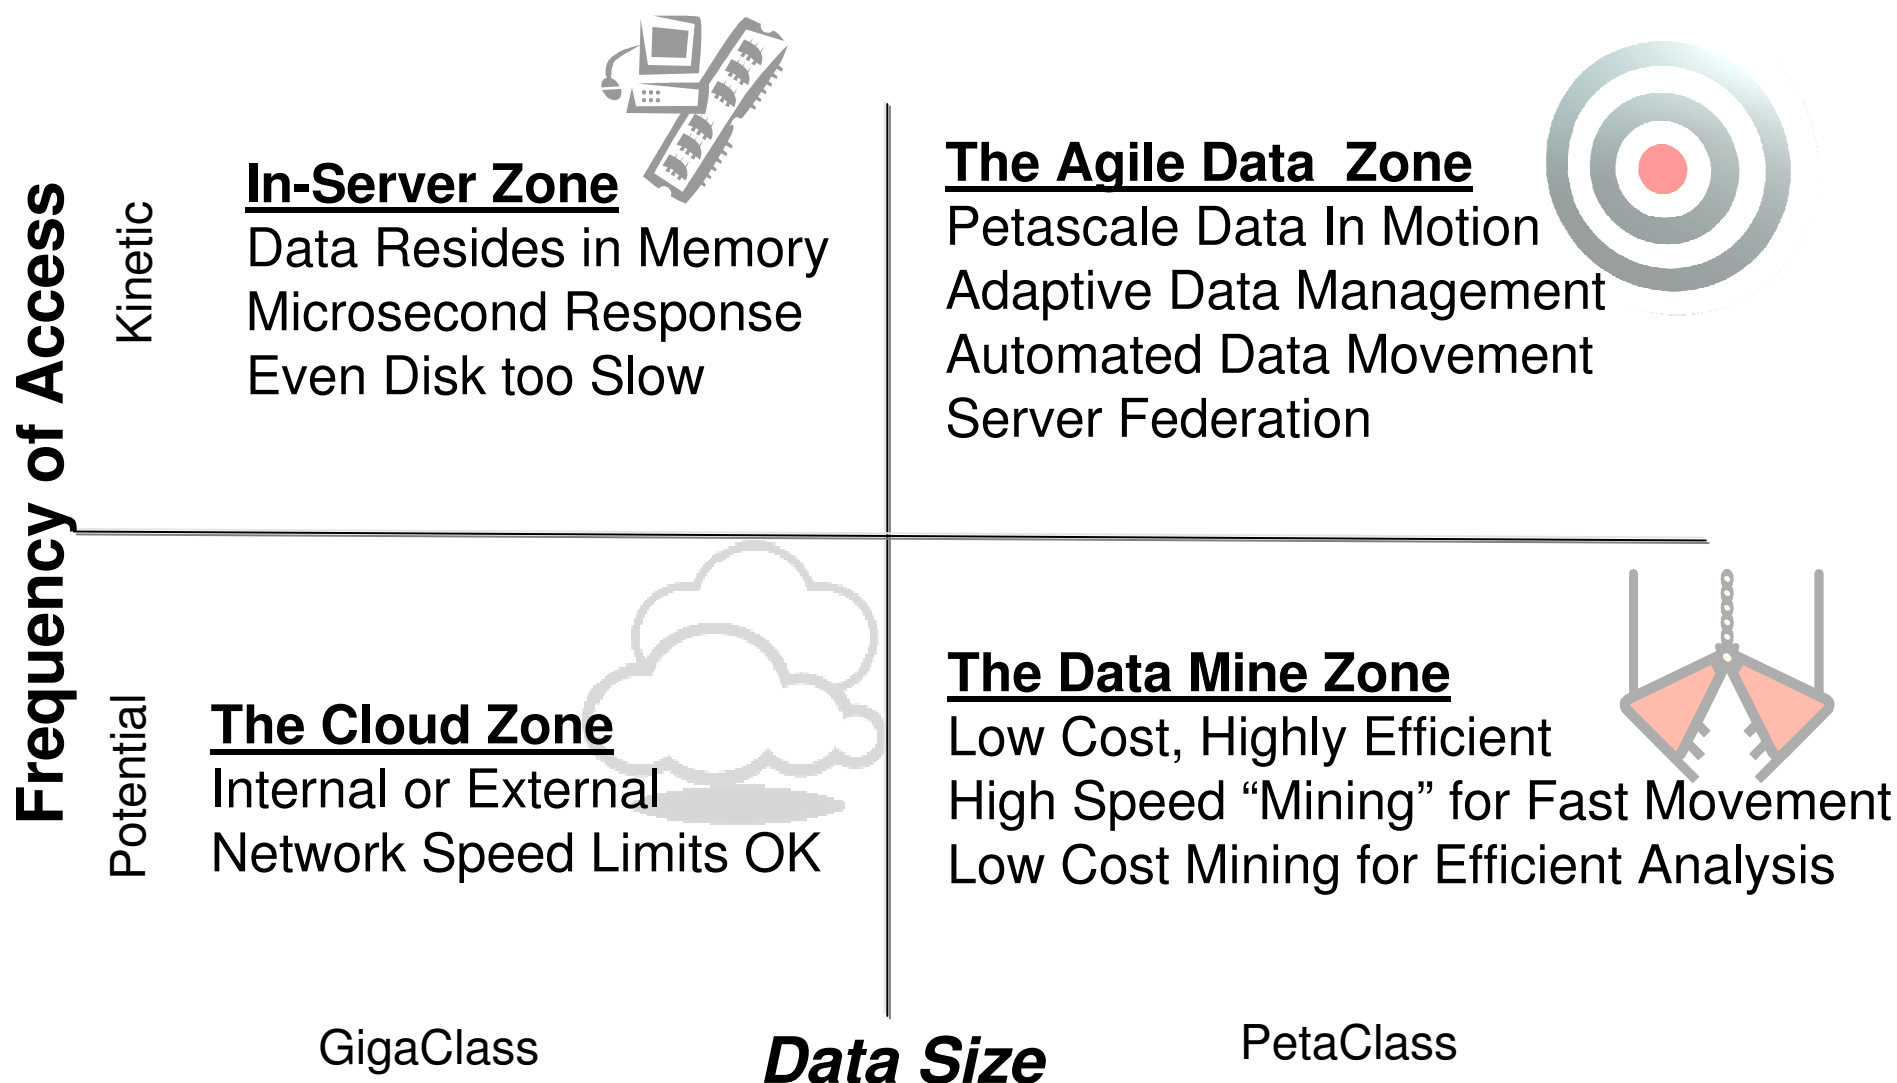
# Comparison of Technologies

<b>Technology</b>	<b>Latency (<math>\mu</math>S)</b>	<b>IOPs</b>	<b>Cost /IOPs(\$)</b>	<b>Cost / GB (\$)</b>
<b><i>Capacity HDDs</i></b>	12,000	600	13.3	3
<b><i>Performance HDDs</i></b>	7000	1,200	16.6	28
<b><i>Flash SSDs</i></b>	200	500	140	100
<b><i>Flash SSDs (read only)</i></b>	45	50,000	1.4	100
<b><i>DRAM SSDs</i></b>	3	200,000	0.5	400

## You're So Far Away From Me...

<b>Data Size</b>	<b>1 Mb/Sec</b>	<b>10 Mb/Sec</b>
<b>1 GB</b>	2.8 Hours	0.3 Hours
<b>10 GB</b>	1.2 Days	2.8 Hours
<b>50 GB</b>	5.8 Days	13.9 Hours
<b>1 TB</b>	16.5 Weeks	1.7 Weeks
<b>3 TB</b>	49.6 Weeks	5 Weeks
<b>100 TB</b>	31.7 Years	5.2 Years

# Data Management Quadrants



## RAID Controller Evolution

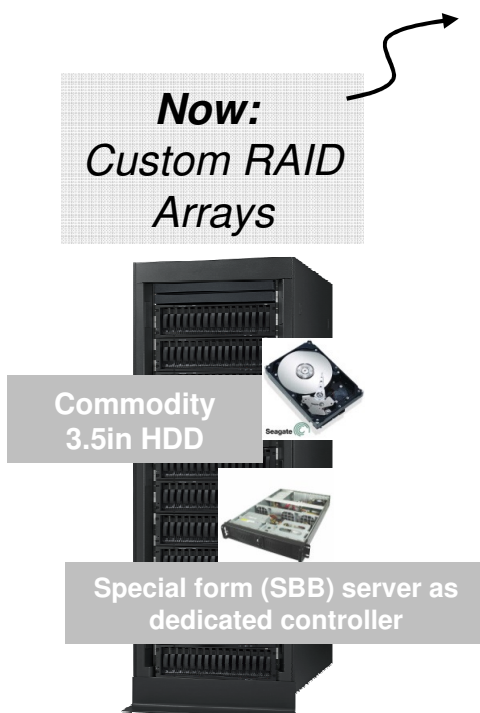
- **Traditional RAID has Evolved**
- **At one point RAID 5 was “Good Enough”**
  - NOW Mean Time to Data Loss is WAY TOO LOW
- **Now, we Deploy Classical RAID 6 everywhere**
  - Is it good enough?
- **Yet, Traditional External RAID controllers remain**
  - Costly
  - Slow to Evolve
  - Far, Far away from Processors



***Where Do We Go Next?***

# High-End “Big Data” Disk Storage Design: Evolution to Simplicity

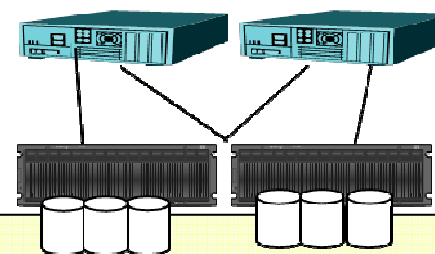
**Yesterday:**  
*Custom hardware*



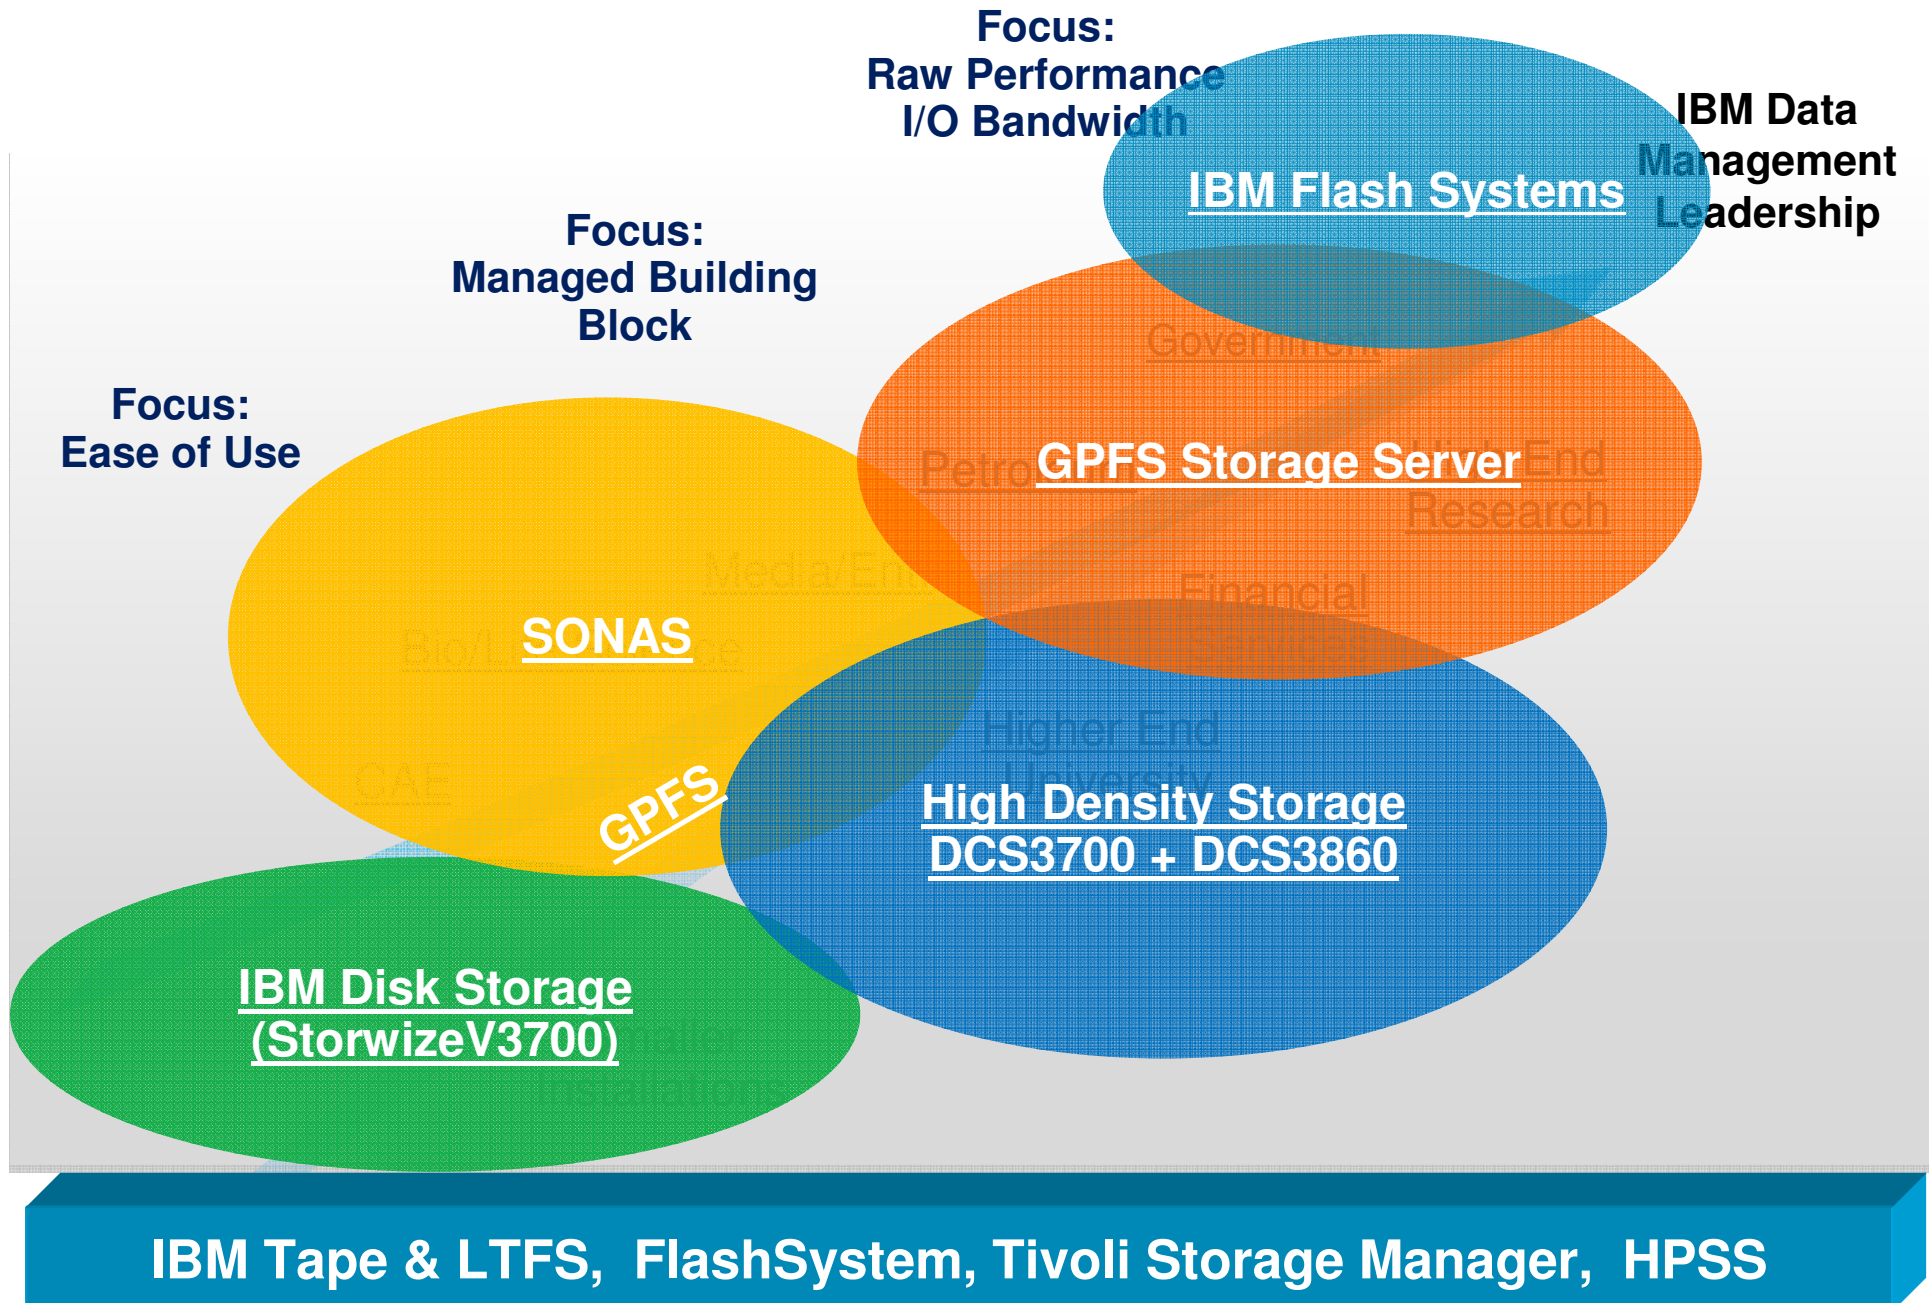
**Tomorrow:**  
*Simplified Infrastructure, Focus on the Data!*

*Controller Optional:*  
**Software RAID in Host nodes spare cores**

**Solid State Devices Change Paradigms**

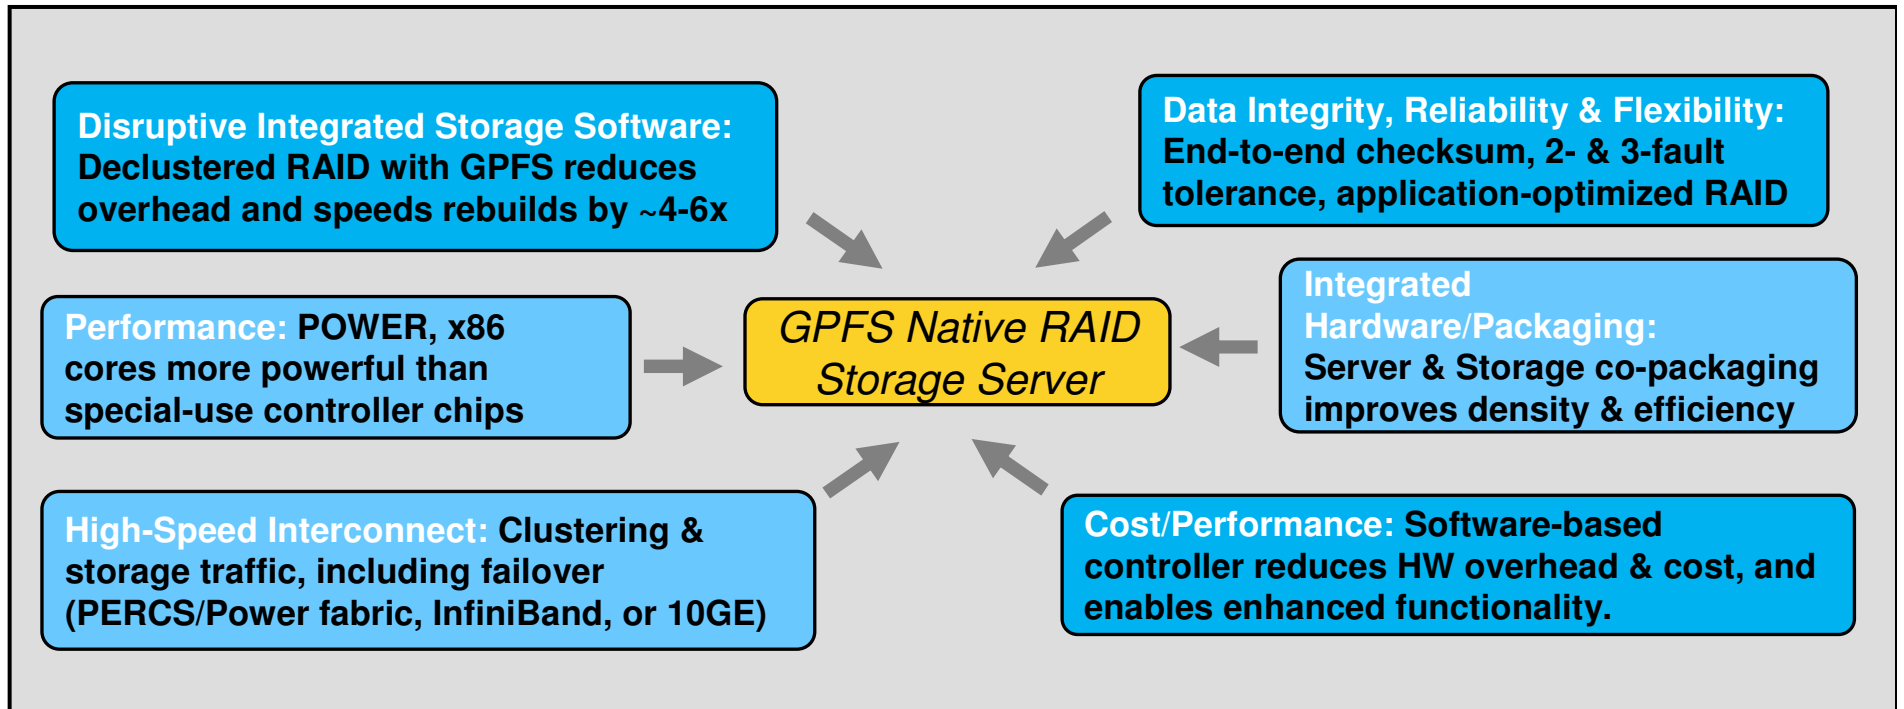


# IBM offers a wide range of storage and data management



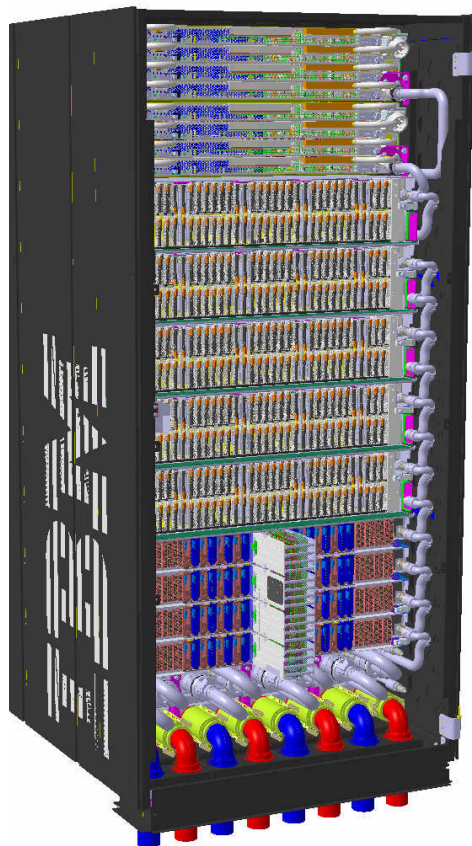
# GPFS Storage Server

“Perfect Storm” of Synergetic Innovations



Big Data Converging with HPC Technology  
Server and Storage Convergence

Shipping NOW  
from POWER



**1 Rack performs a 1TB Hadoop  
TeraSort in less than 3 minutes!**

## ***IBM GPFS Native RAID p775: High-End Storage + Compute Server***

- **Based on Power 775 / PERCS Solution**
- **Basic Configuration:**
  - 32 Power7 32-core high bandwidth servers
  - Configurable as GPFS Native RAID storage controllers, compute servers, I/O servers or spares
  - Up to 5 Disk Enclosures per rack
    - 384 Drives and 64 quad-lane SAS ports each
- **Capacity:** 1.1 PB/rack (900 GB SAS HDDs)
- **Bandwidth:** >150 GB/s per rack Read BW
- **Compute Power:** 18 TF + node sparing
- **Interconnect:** IBM high-BW optical PERCS
- **Multi-rack scalable, fully water-cooled**

## GPFS Storage Server Goals

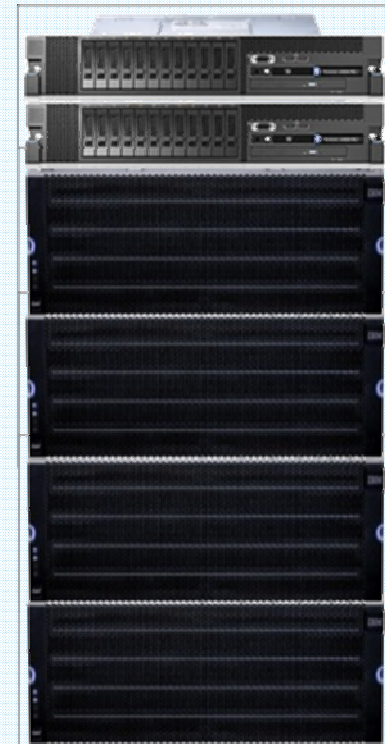
- Better, Sustained Performance
  - The GPS Storage Server provides industry-leading throughput using efficient de-clustered RAID Techniques
- Better Value
  - GPFS Storage Server leverages System x servers and commercial JBODS
- Better Data Security
  - New Data Protection insures data is written, read, and delivered correctly and precisely, from the disk platter to the client.

- Affordably Scalable
  - Start Small and affordable, scale via incremental additions, adding capacity and bandwidth with each change.
- Data Management
  - All of this with the enhanced commercial-class data and lifecycle management capabilities which are part of GPFS Today!
- IT Facility Friendly
  - GPFS Storage Server fits in industry-standard 42u 19 inch rack mounts – no special height requirements
  - 3 Year Warranty

***All with the enhanced commercial-class data/lifecycle management capabilities which are part of GPFS today!***

## Introducing IBM System x GPFS Storage Server: Bringing HPC Technology to the Mainstream


- **Better, Sustained Performance**
  - Industry-leading throughput using efficient De-Clustered RAID Techniques
- **Better Value**
  - Leverages System x servers and Commercial JBODS
- **Better Data Security**
  - From the disk platter to the client.
  - Enhanced RAID Protection Technology
- **Affordably Scalable**
  - Start Small and Affordably
  - Scale via incremental additions
  - Add capacity AND bandwidth
- **3 Year Warranty**
  - Manage and budget costs
- **IT-Facility Friendly**
  - Industry-standard 42u 19 inch rack mounts
  - No special height requirements
  - Client Racks are OK!



16 • **And all the Data Management/Life Cycle Capabilities of GPFS – Built in!**


# A Scalable Building Block Approach to Storage

*Complete Storage Solution  
Data Servers, Disk (NL-SAS and SSD), Software, InfiniBand and Ethernet*




x3650 M4

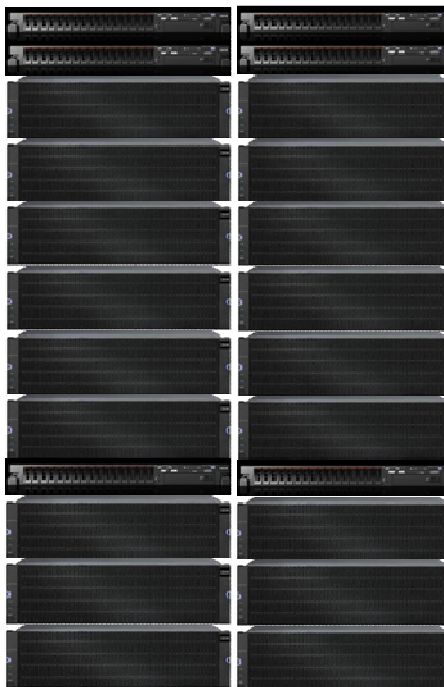
"Twin Tailed"  
JBOD  
Disk Enclosure



**Model 24:**  
**Light and Fast**  
4 Enclosures, 20U  
232 NL-SAS, 6 SSD  
**10 GB/Sec**

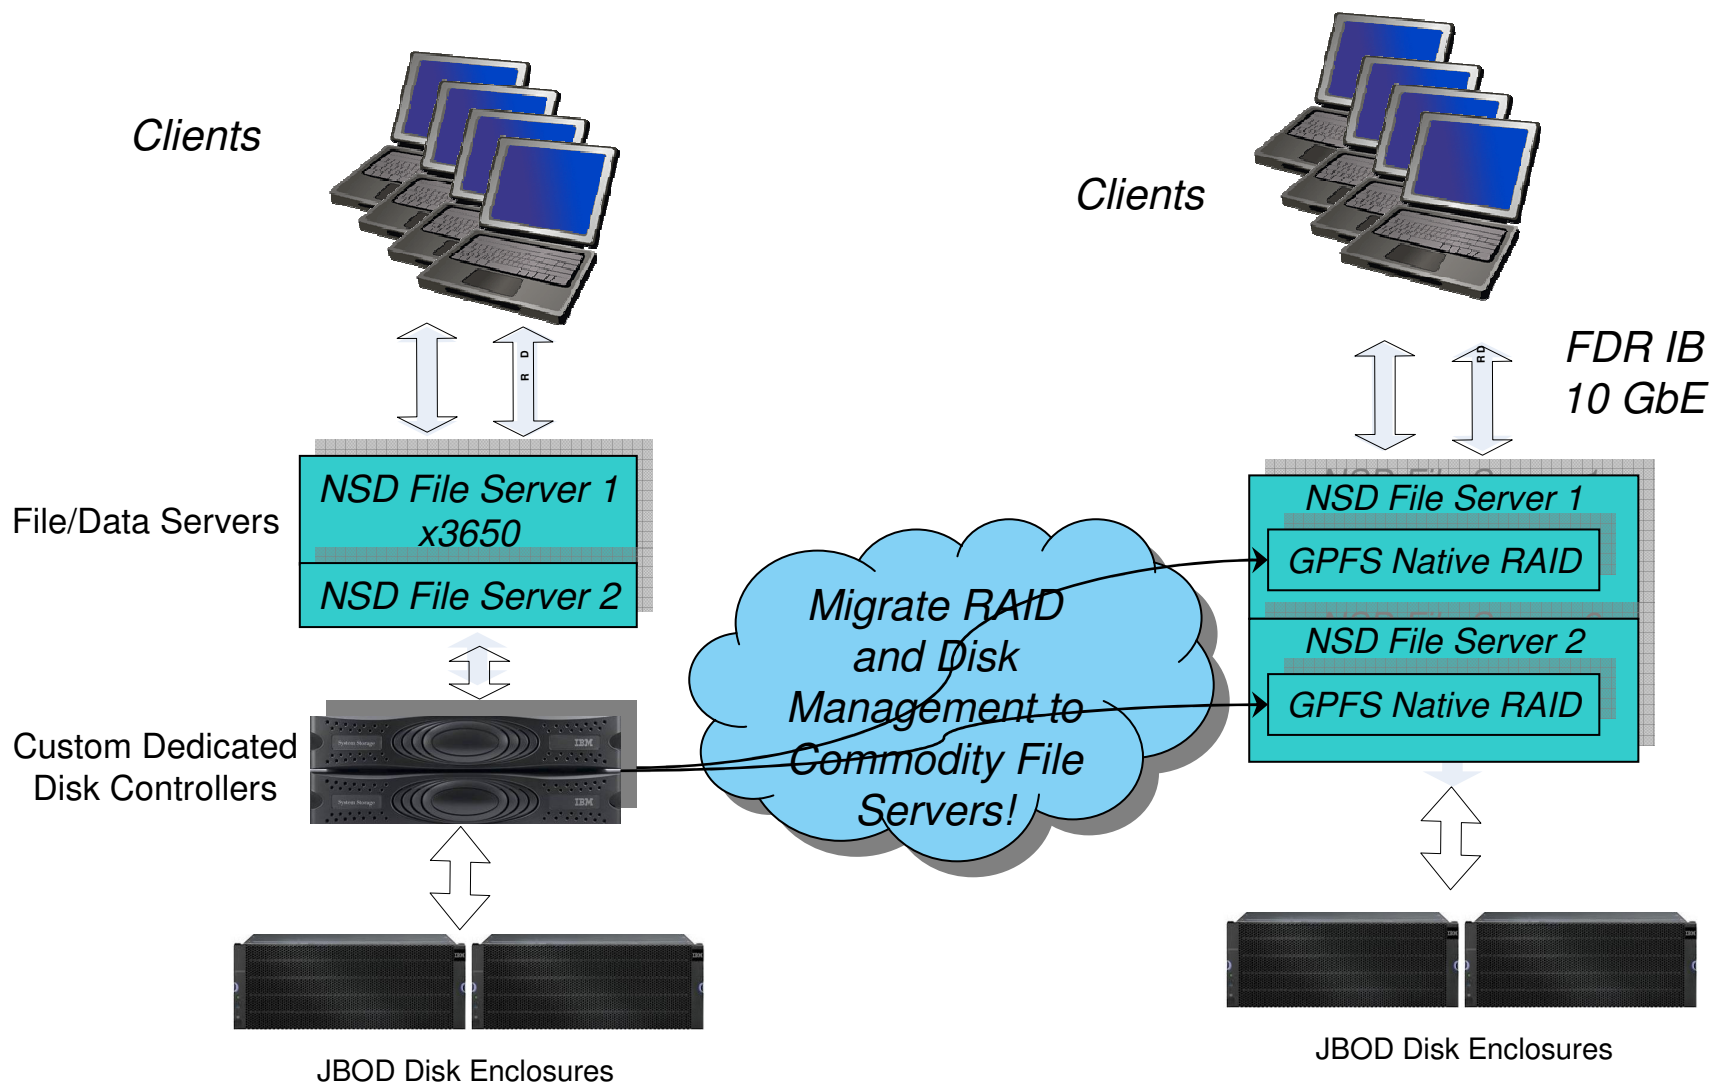


**Model 26:**  
**HPC Workhorse!**  
6 Enclosures, 28U  
348 NL-SAS, 6 SSD  
**12 GB/sec**



**High-Density HPC Option**  
18 Enclosures  
2 - 42U Standard Racks  
1044 NL-SAS 18 SSD  
**36 GB/sec**

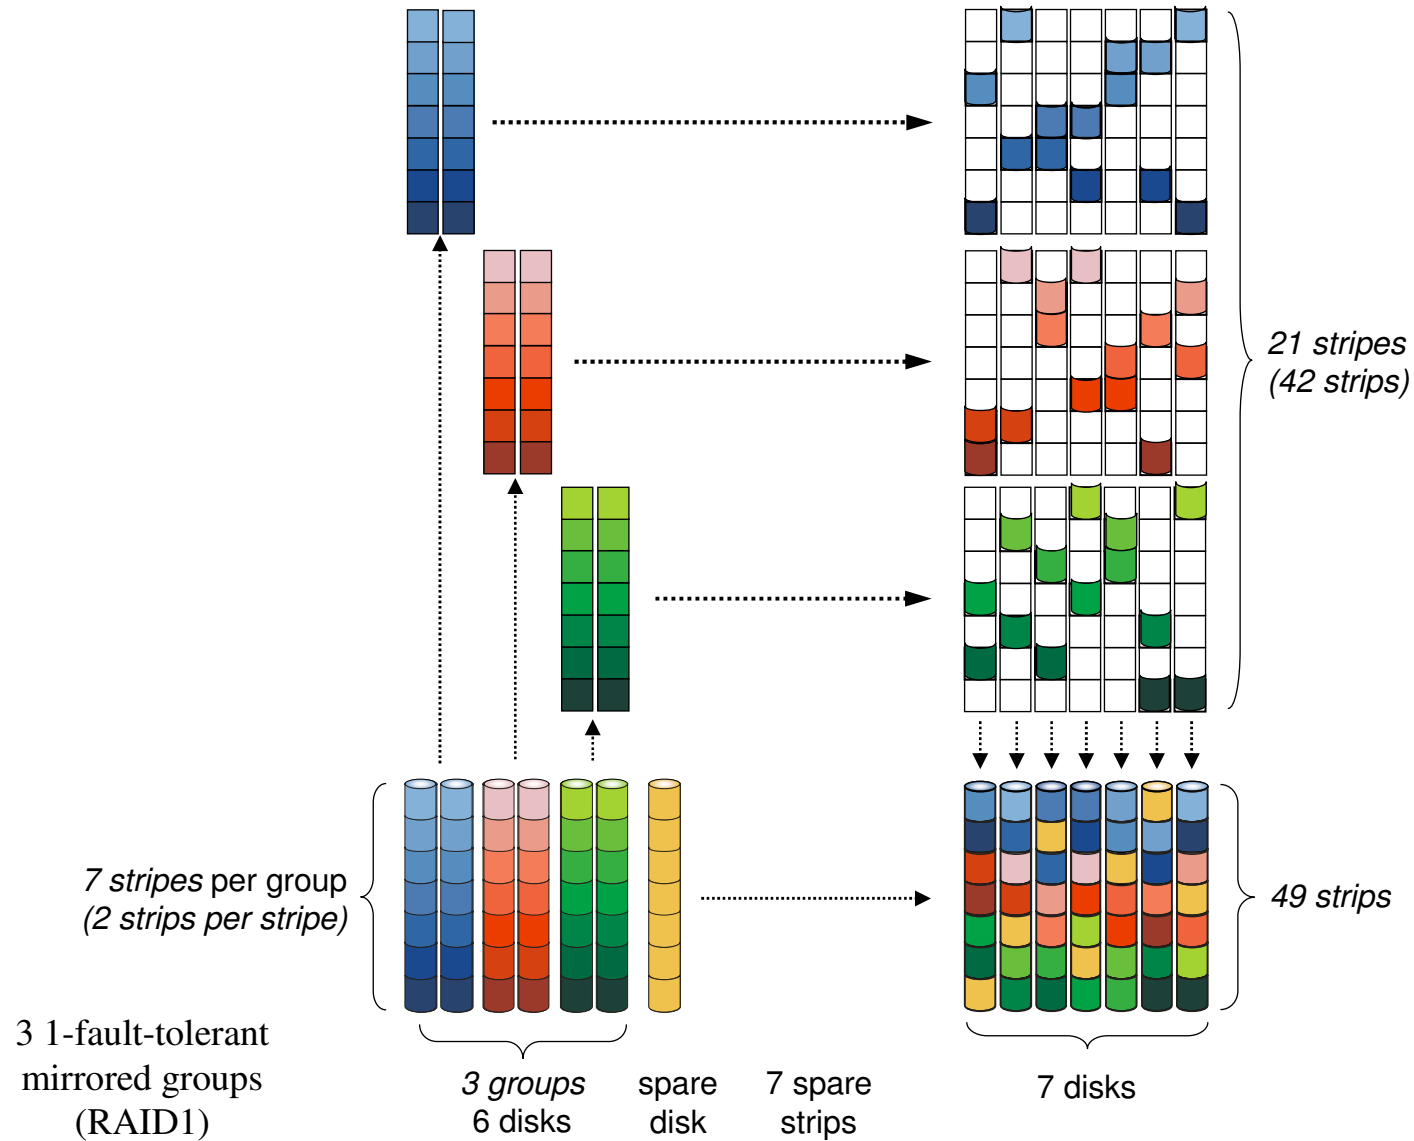
## How We Did It!



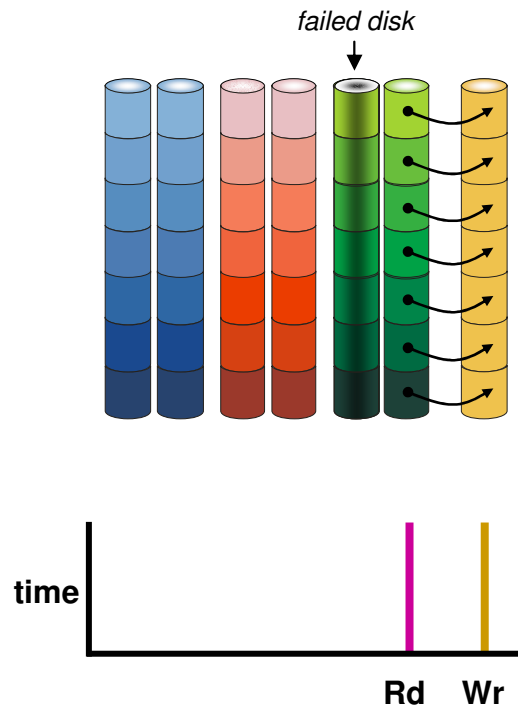
## GPFS Native RAID Feature Detail

- **Declustered RAID**
  - Data and parity stripes are uniformly partitioned and distributed across a disk array.
  - Arbitrary number of disks per array (unconstrained to an integral number of RAID stripe widths)
- **2-fault and 3-fault tolerance**
  - Reed-Solomon parity encoding
  - 2 or 3-fault-tolerant: stripes = 8 data strips + 2 or 3 parity strips
  - 3 or 4-way mirroring
- **End-to-end checksum & dropped write detection**
  - Disk surface to GPFS user/client
  - Detects and corrects off-track and lost/dropped disk writes
- **Asynchronous error diagnosis while affected IOs continue**
  - If media error: verify and restore if possible
  - If path problem: attempt alternate paths
- **Supports live replacement of disks**
  - IO ops continue on for tracks whose disks have been removed during carrier service

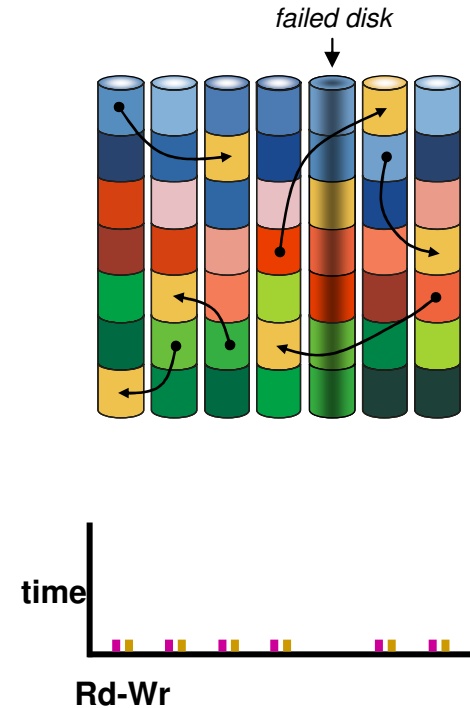
# Declustered RAID Example



# Rebuild Overhead Reduction Example



Rebuild activity confined to just a few disks – slow rebuild, disrupts user programs

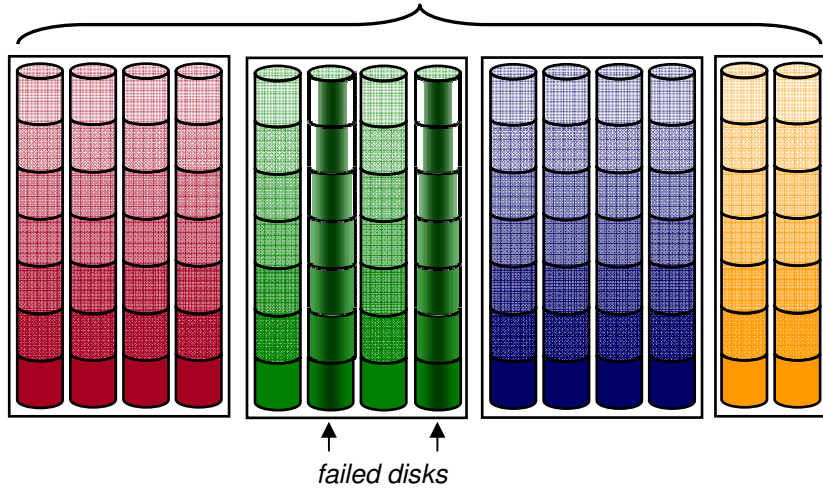


Rebuild activity spread across many disks, less disruption to user programs

Rebuild overhead reduced by 3.5x

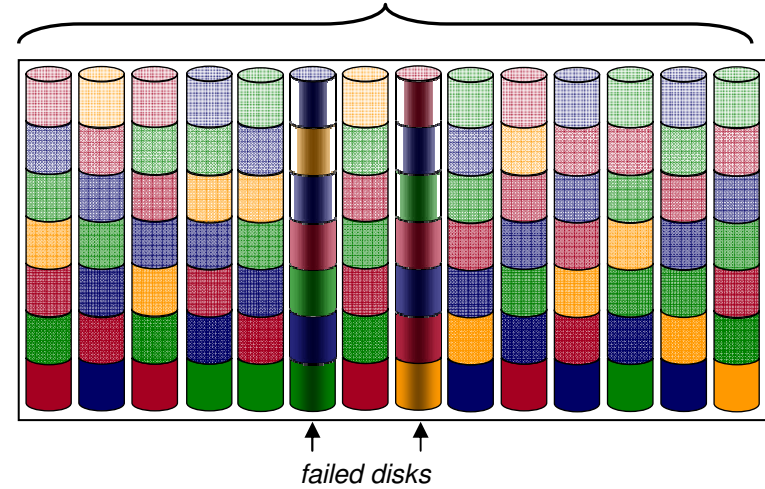
# Declustered RAID6 Example

14 physical disks / 3 traditional RAID6 arrays / 2 spares

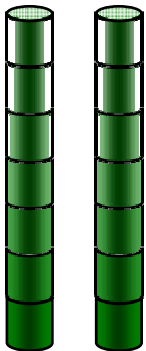


Decluster data, parity and spare  
→

14 physical disks / 1 declustered RAID6 array / 2 spares



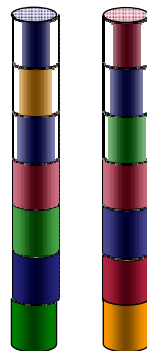
failed disks



Number of faults per stripe		
Red	Green	Blue
0	2	0
0	2	0
0	2	0
0	2	0
0	2	0
0	2	0
0	2	0

Number of stripes with 2 faults = 7

failed disks



Number of faults per stripe		
Red	Green	Blue
1	0	1
0	0	1
0	1	1
2	0	0
0	1	1
1	0	1
0	1	0

Number of stripes with 2 faults = 1

- Platter-to-Client Protection
  - Multi-level data protection to detect and prevent bad writes and on-disk data loss
  - Data Checksum carried and sent from platter to client server
- Integrity Management
  - **Rebuild**
    - Selectively rebuild portions of a disk
    - Restore full redundancy, in priority order, after disk failures
  - **Rebalance**
    - When a failed disk is replaced with a spare disk, redistribute the free space
  - **Scrub**
    - Verify checksum of data and parity/mirror
    - Verify consistency of data and parity/mirror
    - Fix problems found on disk
  - **Opportunistic Scheduling**
    - At full disk speed when no user activity
    - At configurable rate when the system is busy

## Non-Intrusive Disk Diagnostics

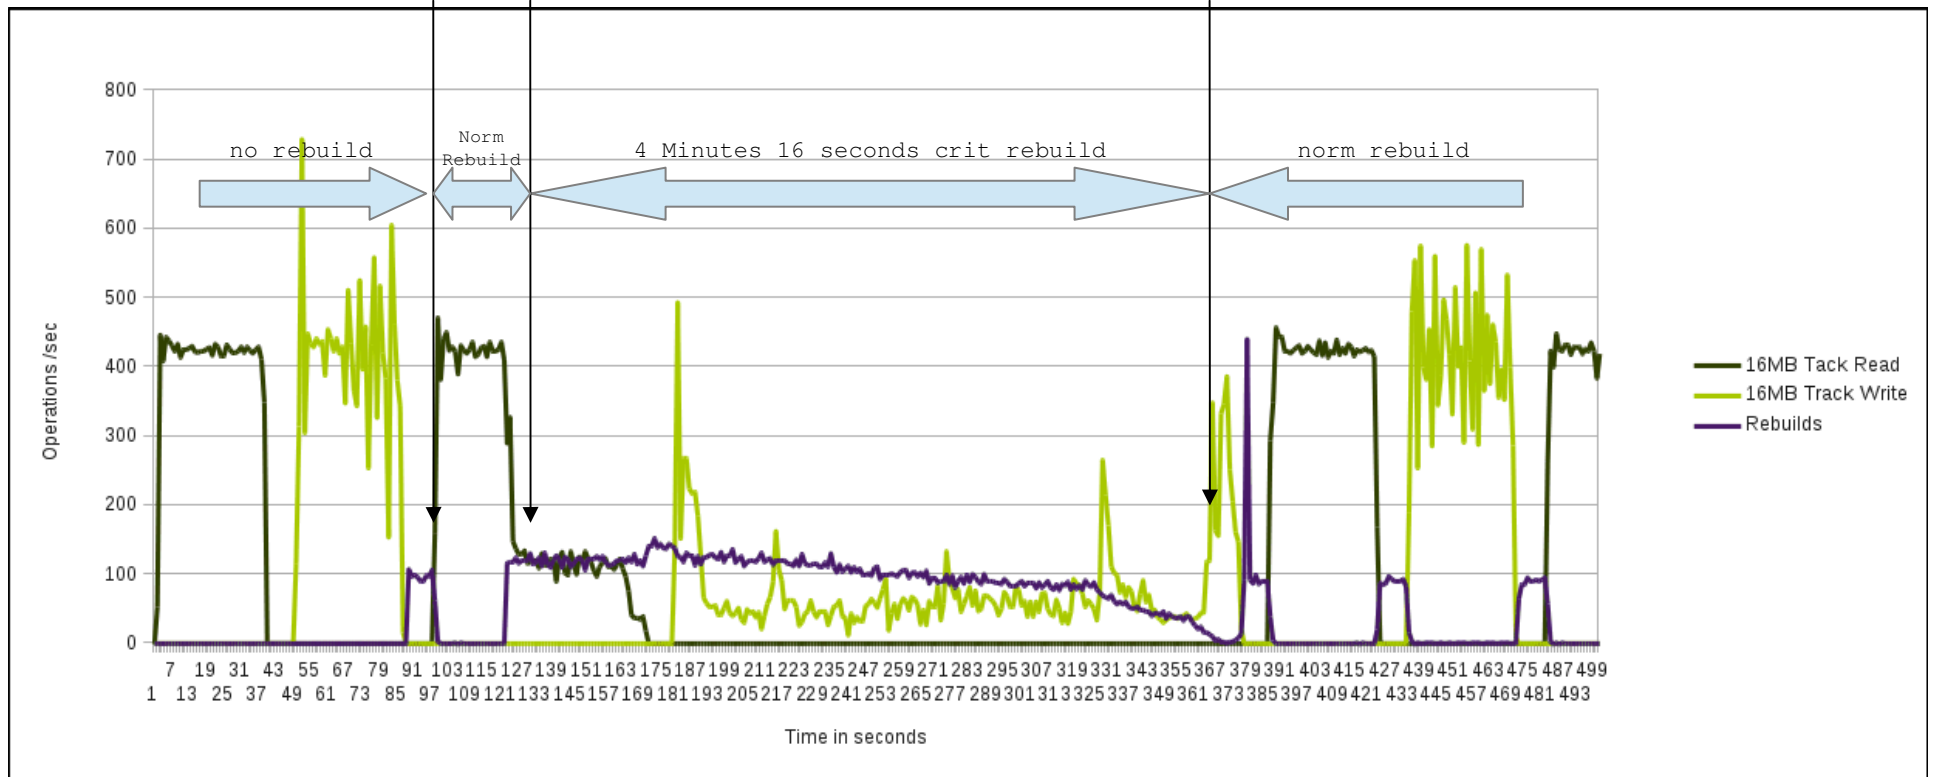
- **Disk Hospital: Background determination of problems**
  - While a disk is in hospital, GNR non-intrusively and *immediately* returns data to the client utilizing the error correction code.
  - For writes, GNR non-intrusively marks write data and reconstructs it later in the background after problem determination is complete.
- **Advanced fault determination**
  - Statistical reliability and SMART monitoring
  - Neighbor check
  - Media error detection and correction

# Rebuild Test

1<sup>st</sup> disk failure

3<sup>rd</sup> disk failure -  
start of critical  
rebuild

critical rebuild  
finished, continue  
normal rebuild



As one can see during the critical rebuild impact on workload was high, but as soon as we were back to parity protection (no critical data) the impact to the customers workload was less than 5%

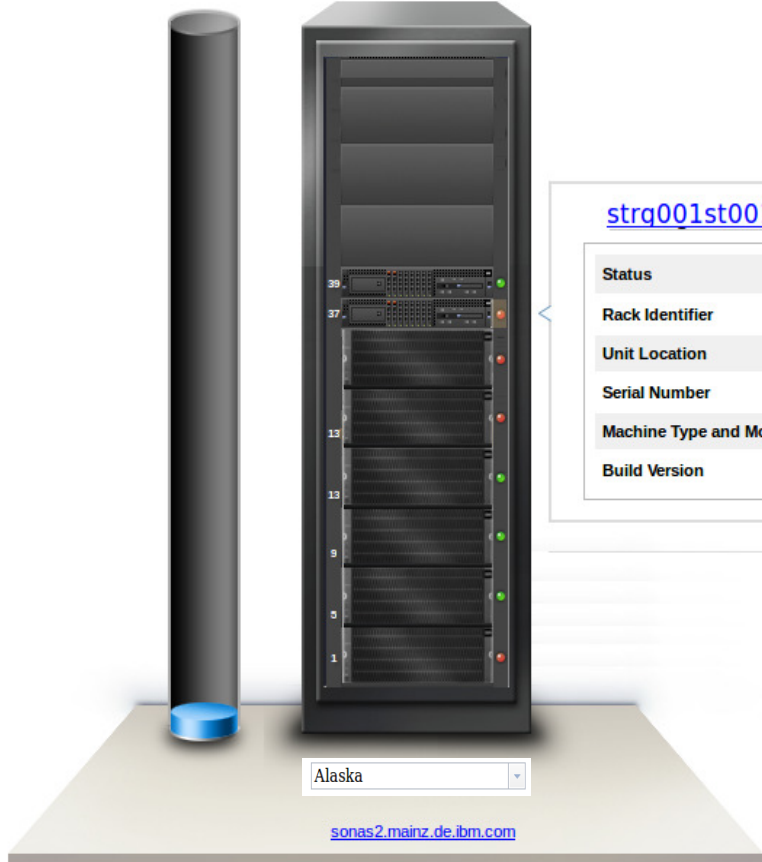
# Monitoring – System: Information for nodes



IBM GPFS Storage Server

Welcome, admin   About   Logout   Help   IBM

gss01-gui > Monitoring > System ▾



[strq001st001](#) Actions ▾

Status	✖ Critical Error
Rack Identifier	1
Unit Location	
Serial Number	99A1399
Machine Type and Model	x3650m2
Build Version	1.4.1.0-12

See Demo in IBM Booth!

Allocated Capacity: 5.4 GB / 42.3 TB (0%)   Running Tasks (0)   ✖ Health Status

# Error on Drawer – Health Flyover



sonas2.mainz.de.ibm.com - Monitoring - IBM Scale Out Network Attached Storage - Mozilla Firefox

File Edit View History Bookmarks Tools Help

STGC STO... sonas2.... GSS GUI - P... SourceForg... IBM Redbo... IBM Storwi... Google-Erg... Objekt nich... Main Page ...

https://sonas2mgmt.mainz.de.ibm.com:1081/gui#monitor-system

IBM GPFS Storage Server Welcome, admin About Logout Help IBM.

gss01-gui > Monitoring > System

[strq001st001](#)

Actions

Alerts 1

- 6m The enclosure has gone offline.
- 6m Fibre cable on port 4 is faulty.
- 6m Fan #3 has stopped working.

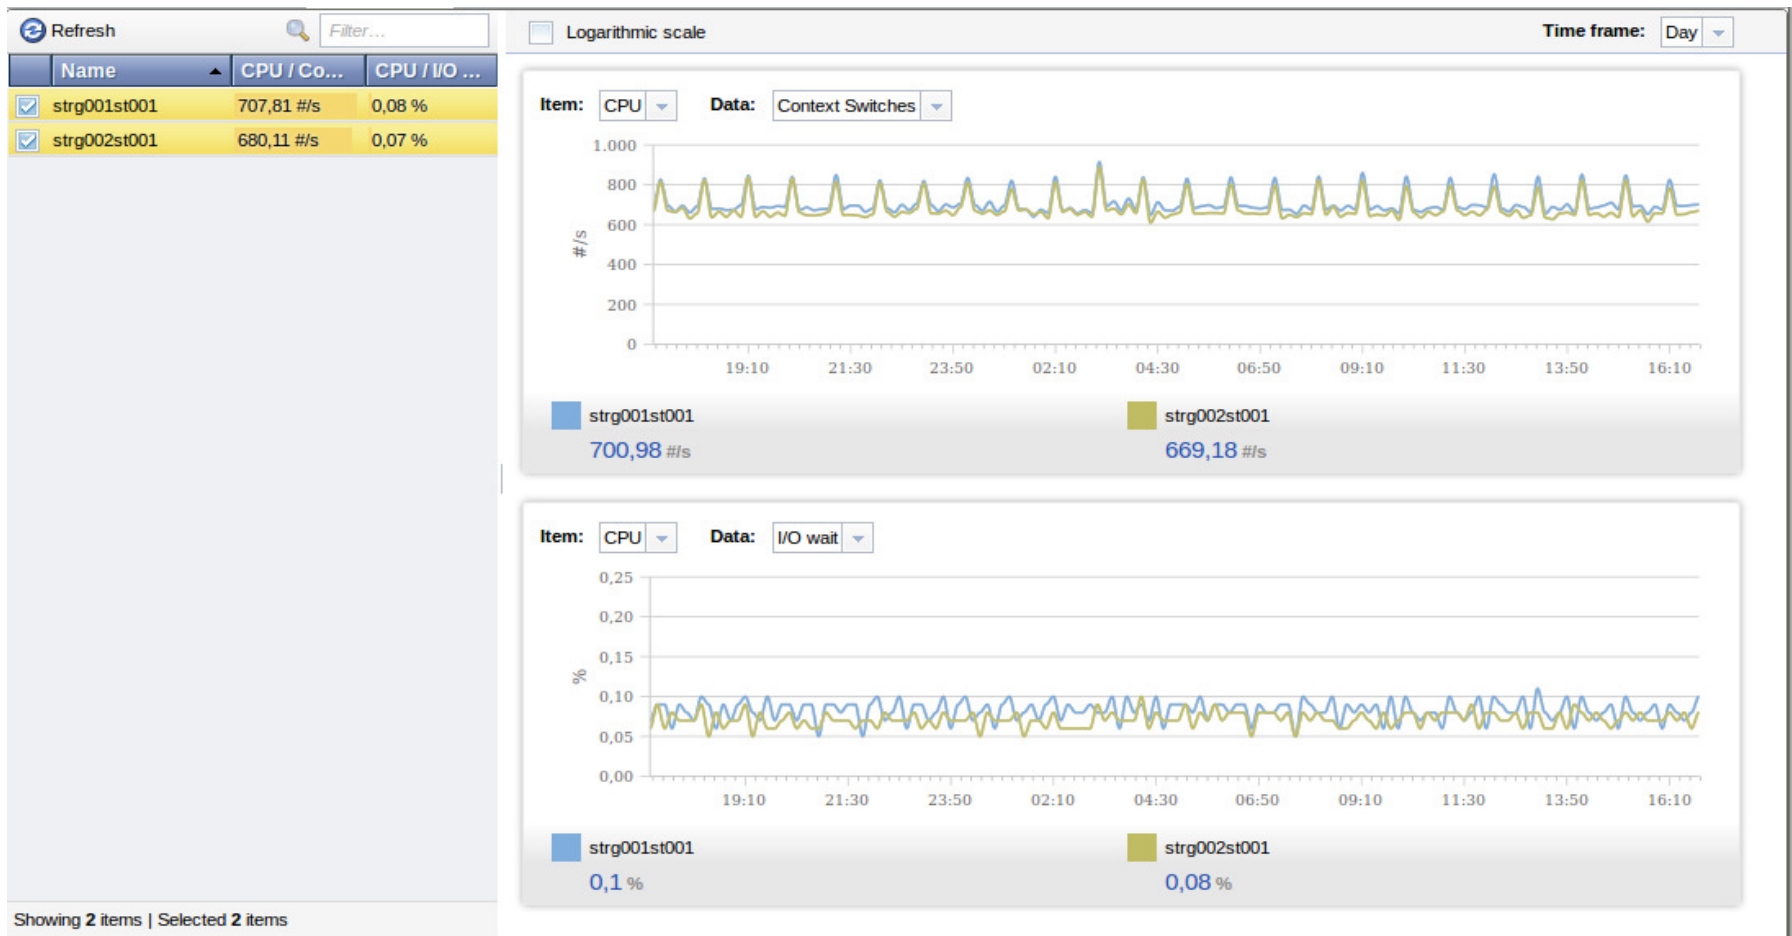
sonas2.mainz.de.ibm.com

Allocated Capacity: 5.4 GB / 42.3 TB (0%) Running Tasks (0) Health Status

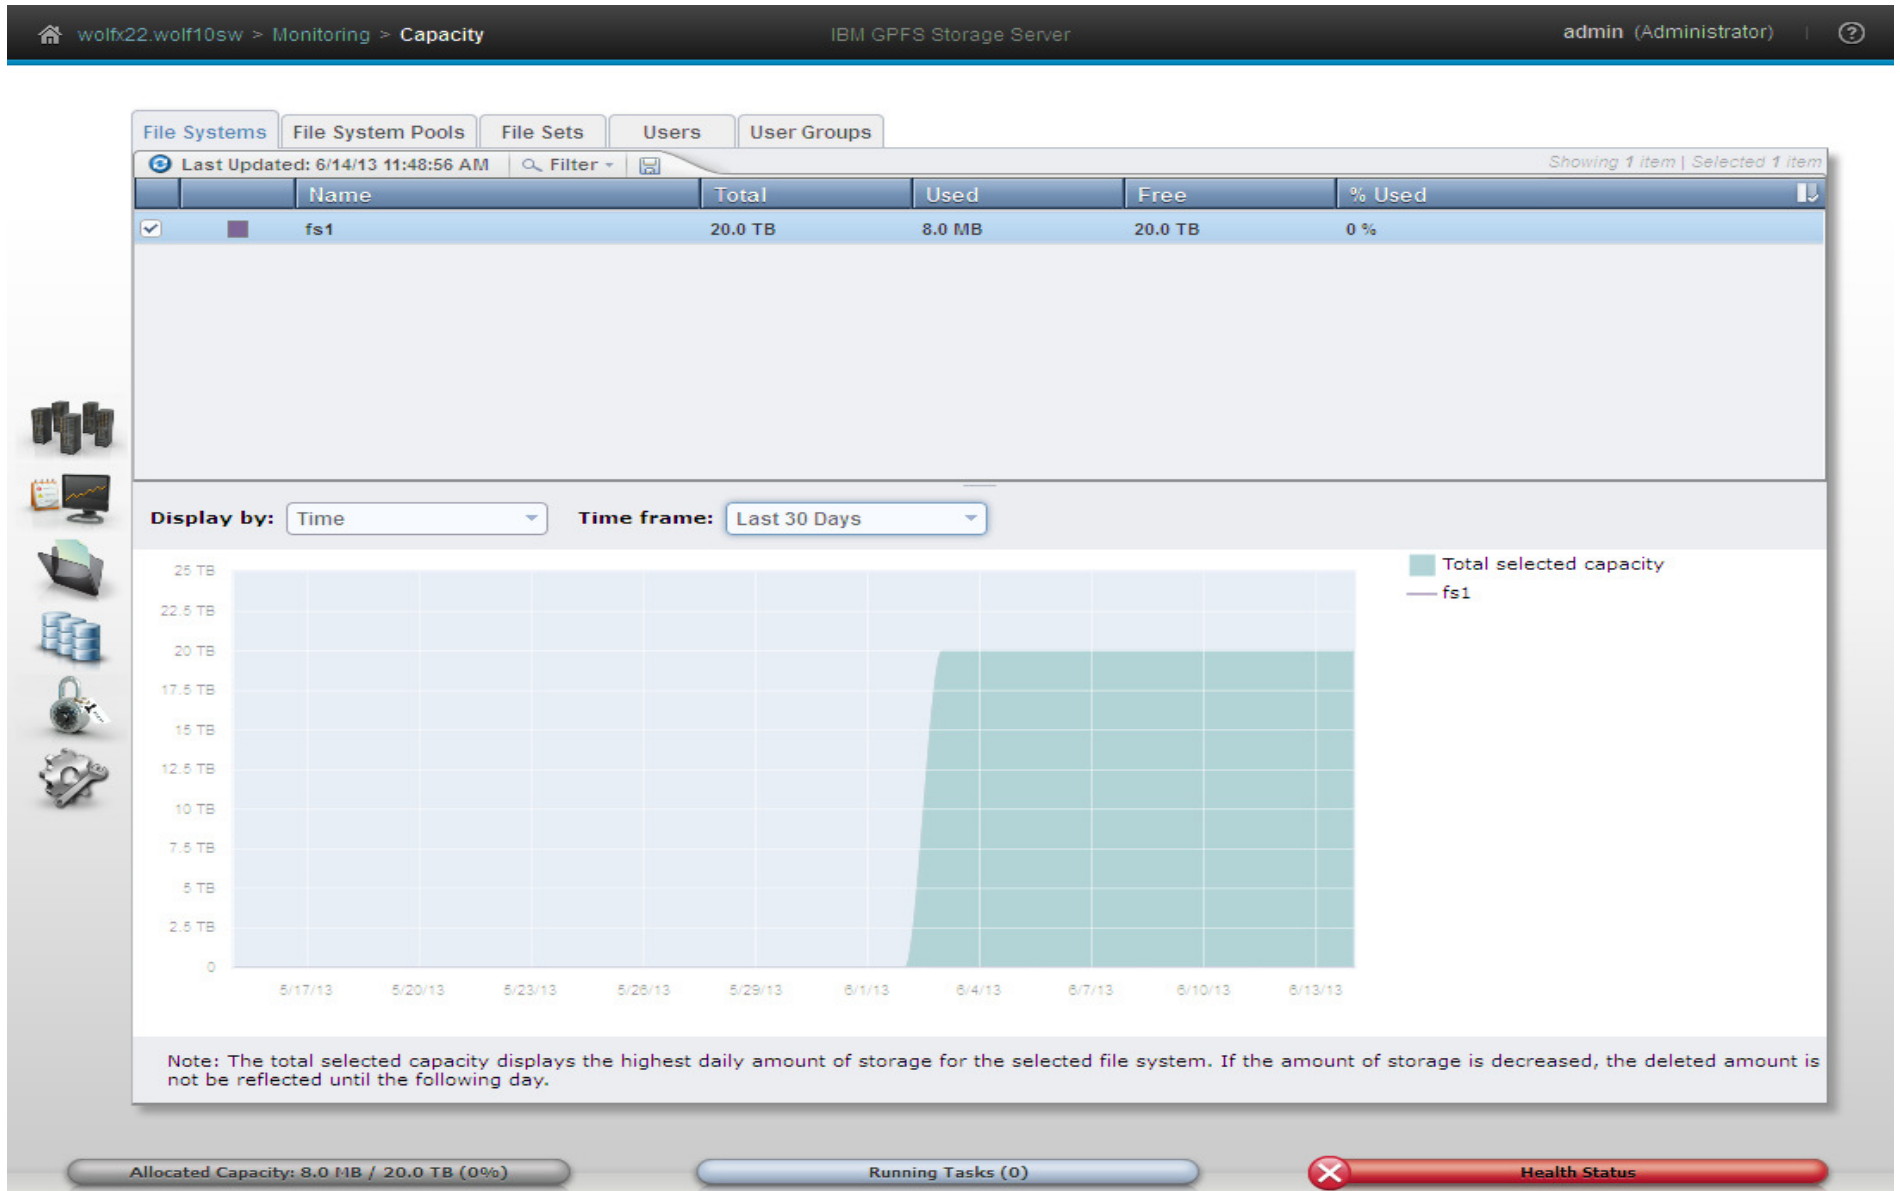
# Monitor>System – Select Drawer



# Monitor - Performance



# Monitor - Capacity



# Files – File Systems



wolf22.wolf10sw > Files > File Systems IBM GPFS Storage Server admin (Administrator) ?

Selected 1 file system: mounted on 2 of 2 nodes

Name	Capacity	Status
fs1	20.2 TB	OK
system (metadata disks)	200.7 GB	OK
G1Metadata	100.3 GB	OK
G2Metadata	100.3 GB	OK
normal	20.0 TB	OK
G1Data	10.0 TB	OK
G2Data	10.0 TB	OK

New File System

Single Pool Migration-ILM Custom

Name and Pool

\* File system name: Block size: 256 KB System pool disk usage: ☐ Use separate disks for data and metadata

\* Owner: \* Group: Owner or Group is required Owner or Group is required

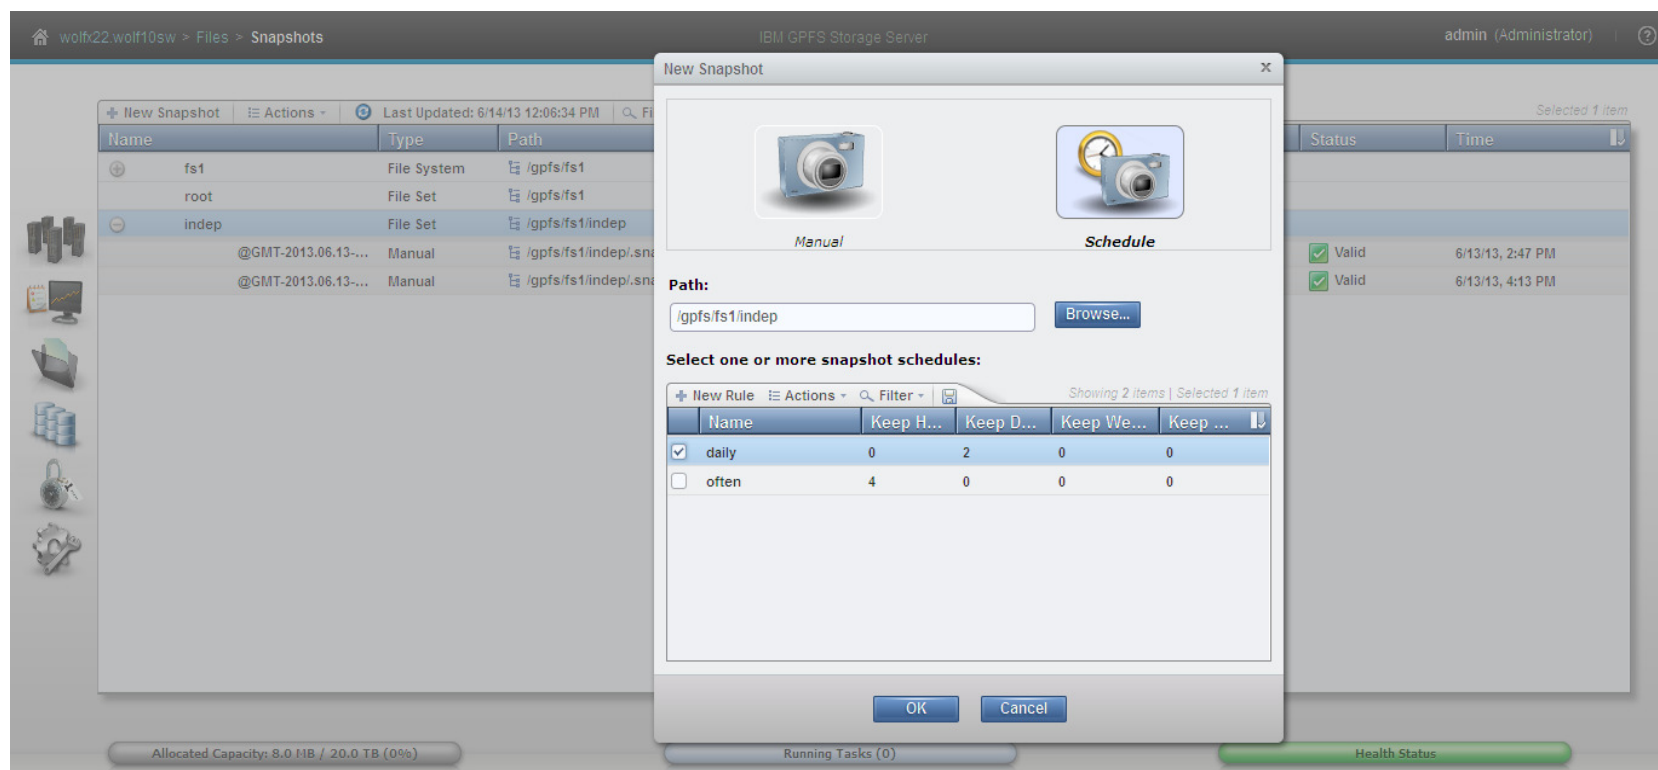
File system pool settings

\* Name: system \* Disk type: Size: There are no available disk types.

OK Cancel

Health Status

# Files -Snapshots



GPFS snapshot capability is exposed and are enriched by snapshot schedule capabilities which are derived from SONAS/V7000 Unified

## Files - Quotas



The "New Quota" dialog box is shown. It has three tabs: "File Set", "Group ID", and "User ID". The "User ID" tab is selected. Below the tabs, the "Quota scope:" is set to "/gpfs/fs1". Below that, the "Specify user IDs and quota sizes:" section contains a table with two rows. The first row shows User ID 143 with a Soft Limit of 100 GB and a Hard Limit of 120 GB, with a green plus icon. The second row shows User ID 234 with a Soft Limit of GB and a Hard Limit of GB, with a green plus and a red X icon. At the bottom are "OK" and "Cancel" buttons.

User ID	Soft Limit	Hard Limit	
143	100 GB	120 GB	+
234	GB	GB	+ X

GPFS Quota Management capabilities for file sets, users and groups are inherited from SONAS/V7000 Unified

# Access – Users and Audit Log

The screenshot displays the IBM GPFS Storage Server web interface. The top navigation bar shows the path 'wolf22.wolf10sw > Access > Users' and the user 'admin (Administrator)'. The main content area is divided into two panels. The left panel, titled 'User Groups', lists 'All Users', 'Administrator', and 'securityAdmin'. The right panel, titled 'All Users', shows a table of users with columns 'Name', 'User Groups', and 'Active'.

Name	User Groups	Active
admin	Administrator	Yes
superuser	securityAdmin	Yes

Below the user list, there are buttons for 'Running Tasks (0)' and 'Health Status'. The bottom panel, titled 'Audit Log', shows a table of system events with columns 'Date and Time', 'Originator', 'Command', and 'Result'. A modal dialog box is open over the audit log, displaying the command 'chfs fs1 --remove G2Data -c 1857390657570266891' and a 'Close' button.

Date and Time	Originator	Command	Result
6/14/13 11:12:58 AM	GUI	chdisk G1SystemData,G2SystemData --pool GUI_label_SystemData --usagetype dataOnly --cluster 1857390657570266891	SUCCESS
6/14/13 11:12:58 AM	GUI	chdisk G2SystemData --failuregroup 13 --usagetype dataOnly --cluster 1857390657570266891	SUCCESS
6/14/13 11:12:58 AM	GUI	chdisk G1SystemData --failuregroup 11 --usagetype dataOnly --cluster 1857390657570266891	SUCCESS
6/14/13 10:56:09 AM	CLI	chfs fs1 --remove G2Data -c 1857390657570266891	INPUT_ERROR
6/13/13 4:14:26 PM	GUI	mksnaprule daily --frequency daily --minutesAfterHour 15 --hour 13 --maxDays 2 --maxWeeks 0 --maxMonths 0 --cluster 185...	SUCCESS
6/13/13 4:13:45 PM	GUI	mksnapshot fs1 @GMT-2013.06.13-14.13.37 --fileset indep --cluster 1857390657570266891	SUCCESS
6/13/13 2:48:25 PM	GUI	mksnaprule often --frequency _none_ --minutesAfterHour 7,47,57 --hour all --maxHoursOrMinutes 4 --maxDays 0 --maxWeeks...	SUCCESS
6/13/13 2:47:31 PM	GUI	mksnapsh Command	SUCCESS
6/13/13 2:47:13 PM	GUI	mksnapsh	SUCCESS
6/13/13 1:04:20 PM	CLI	mkuser su	SUCCESS
6/13/13 1:03:47 PM	CLI	mkusergrp	SUCCESS
6/13/13 11:49:02 AM	GUI	mksnapsh	SUCCESS
6/13/13 11:48:20 AM	GUI	mksnapsh	SUCCESS
6/13/13 11:47:17 AM	GUI	chfs fs1 c	SUCCESS
6/13/13 2:15:40 PM	CLI	cfppercen	SUCCESS
6/12/13 2:14:42 PM	CLI	cfppercen --restart -c 1857390657570266891	SUCCESS
6/12/13 12:33:50 PM	CLI	cfppercen --start -c 1857390657570266891	SUCCESS
6/12/13 12:33:17 PM	CLI	cfppercen --stop -c 1857390657570266891	SUCCESS
6/12/13 11:38:33 AM	CLI	cfppercen --start -c 1857390657570266891	SUCCESS
6/12/13 11:33:19 AM	CLI	cfppercen --stop -c 1857390657570266891	SUCCESS
6/12/13 11:25:27 AM	CLI	cfppercen --start -c 1857390657570266891	SUCCESS
6/12/13 11:23:23 AM	CLI	cfppercen --stop -c 1857390657570266891	SUCCESS
6/12/13 11:18:03 AM	CLI	cfppercen --status -c 1857390657570266891	SUCCESS
6/12/13 11:17:44 AM	CLI	cfppercen --status -c 1857390657570266891	SUCCESS

Role based security for administrators and an Audit log of GUI activity

## Settings - Event Notifications

wolfx22.wolf10sw > Settings > Event Notifications IBM GPFS Storage Server admin (Administrator) ?

Event Notifications

Email Server

Email Recipients

SNMP Server

Email Server

Configure an email server that is used by your site. Ensure that the email server is valid.

☒ Enable email notifications

IP address:  
1.2.3.4

Sender's email address:  
gss01@techcomputing.org

Sender's name:  
Admin

Subject:  
...

Header:  
GSS Storage Server gssrmx007 Notification

Footer:

Test email address:

☐ Maximum emails sent per hour: 5

Allocated Capacity: 8.0 MB / 20.0 TB (0%)

Running Tasks (0)

Health Status

Configure and send events per Mail or SNMP

## Galileo Performance Explorer (GSS Partner Product)

- Monitors and stores Performance, Capacity, Configuration info for OS, Storage and Clusters
- Cloud based SaaS model
- 5 minute intervals, stored and viewable for 1+ years
- Expose true
- facts about performance and capacity in seconds

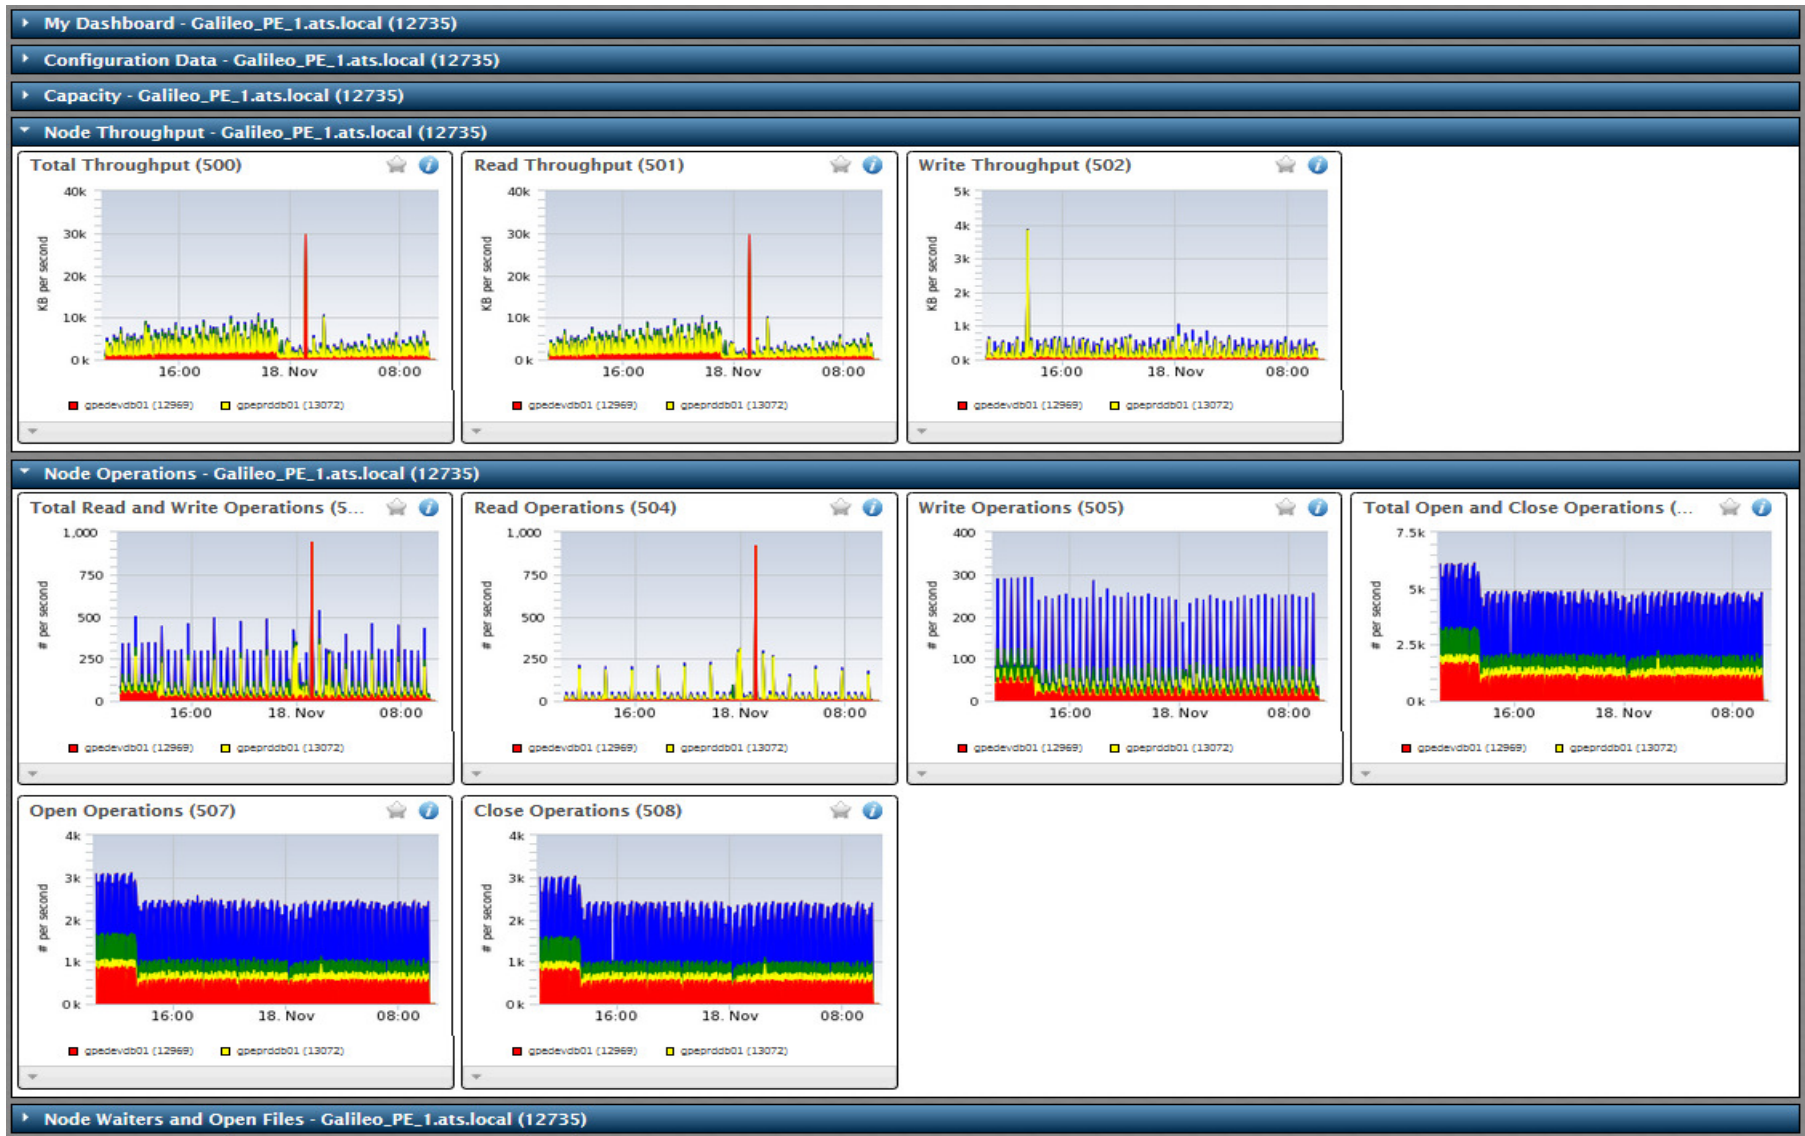


*"All truths are easy to understand once they are discovered; the point is to discover them."*

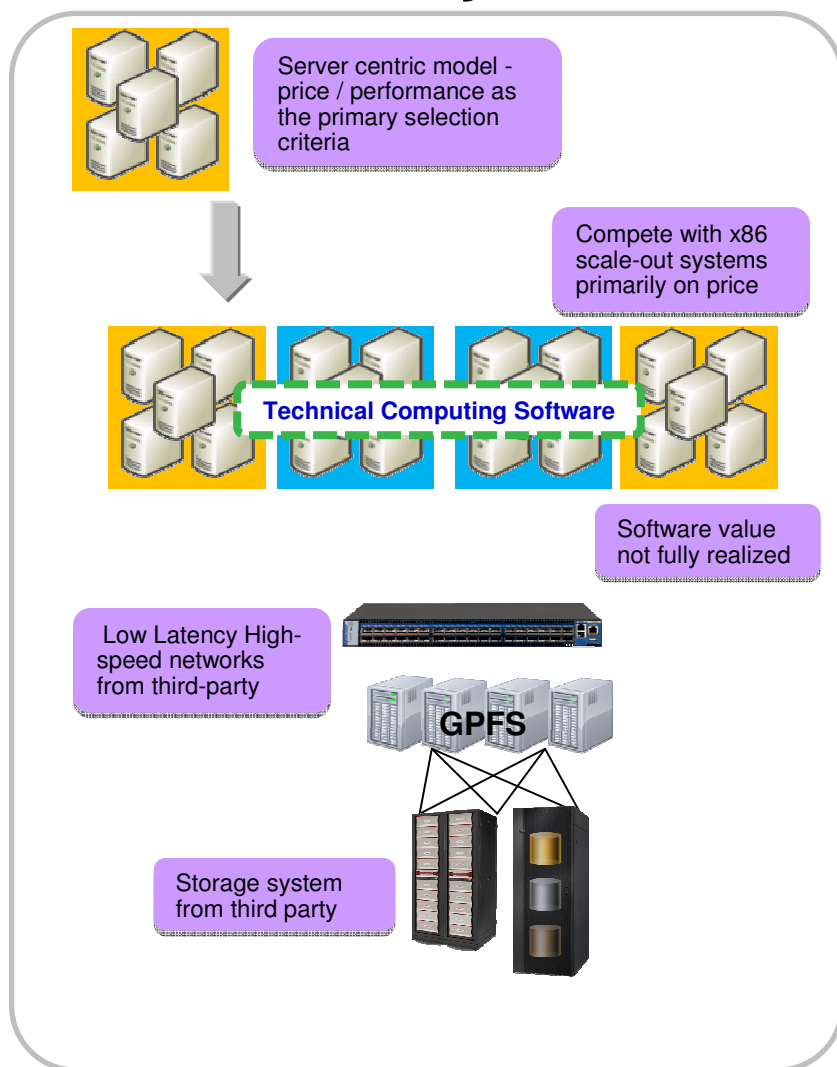
SaaS-built innovations by **ATS** to empower IT administrators to C-level management

© 2007-2013 ATS Group. All rights reserved. | [Privacy](#)

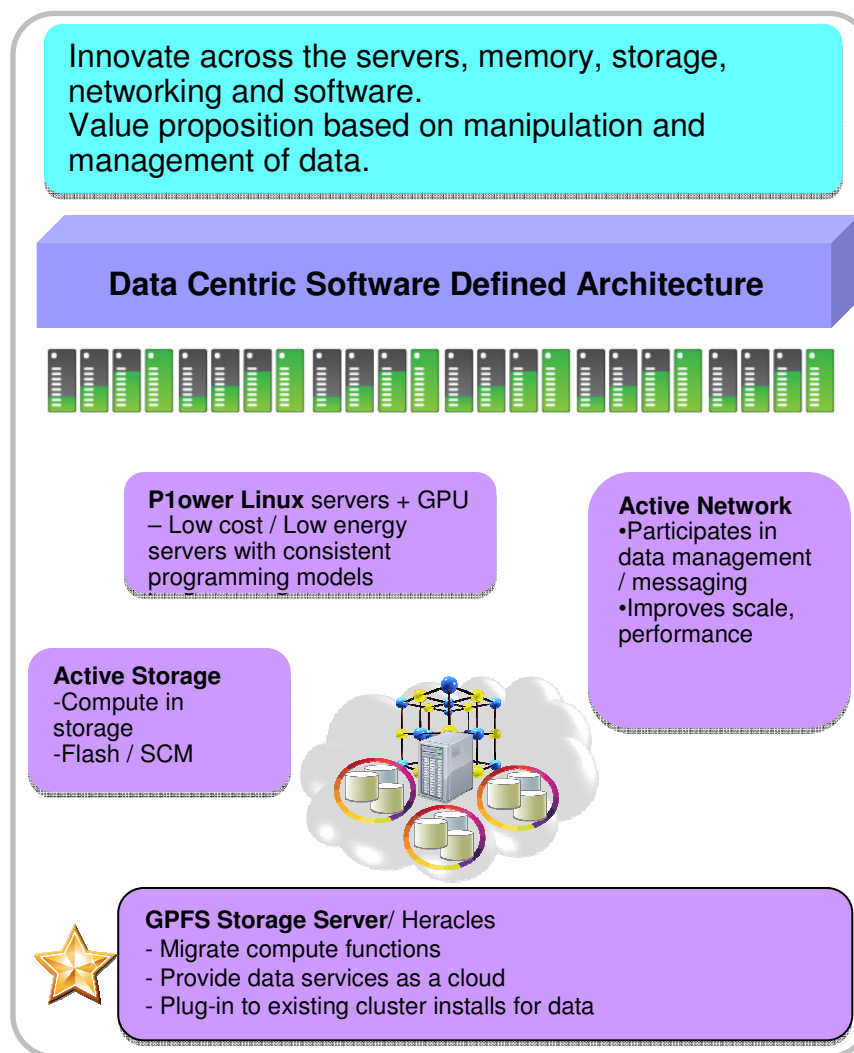
# Detailed Cluster/Node metrics for current and historical



## Today



## Future



# IBM Technical Computing Building Block w/ DCS3XXX Storage



GPFS  
File System  
Client Servers



High Speed  
Network

## GPFS NSD Server

- x3550, x3650,
- Power 7R2

## Client Server farm

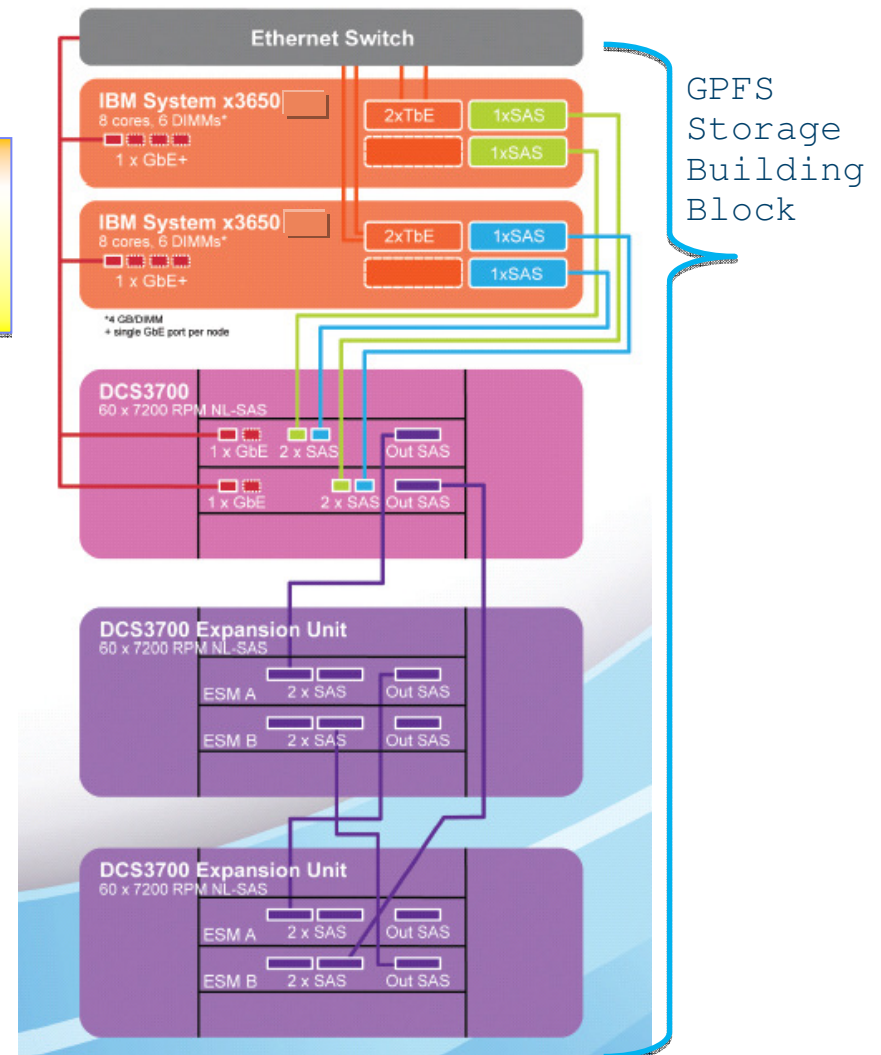
- Handful to thousands of servers
- Each server with 1-n cores
- CPUs & GPUs
- GBs of memory
- Shared Application

## DCS3XXX0 & Disk Drives

- SAS, NL-SAS, SSD
- 7.2K, 15K, SSD
- Capacities to 4TB each

## HighSpeed Network

- >= 10GB/sec
- InfiniBand or Ethernet



# New! DCS3860



*IBM System Storage DCS3860: next generation of high performance controller module to deliver enhanced performance, scalability and simplicity*

- **What's new:**

- New member of the IBM Technical Computing storage family: reliable and affordable high performance storage
- 6Gb SAS host connectivity and scalability up to 360 drives (1.4 PB)
- Flexible intermixing drive capability: SAS, NL-SAS, and SSD
- 2X Sequential Read Improvement
- 2X Sequential Write Improvement



IBM System Storage DCS3860

- **Client Value:**

- Improved streaming performance to satisfy HPC needs
- Performance Read Cache enables the utilization of Solid State Drives to significantly improve read performance
- Unprecedented data availability and dynamic recovery with Dynamic Disk Pools
- T10 PI standard to ensure data integrity
- Intuitive storage management that doesn't sacrifice control

# IBM Disk Storage for Technical Computing – introducing a new member of the DCS series family

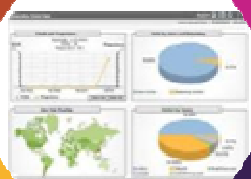
## IBM System Storage DCS3860

Up to Twice as Fast on Streaming Workloads<sup>1</sup>



**DCS3700  
Performance Model**  
Higher performance,  
density, and scalability

**DCS3700**  
High performance, high  
density, scalability



Technical  
Computing  
for Big Data



HPC Cloud

**Storwize V3700**  
Built-in efficiency enhancers

<sup>1</sup>Compared to DCS3700 with Performance Module

# IBM midrange high-performance storage

## **DCS3700**

- Host: FC, SAS, iSCSI
- 180 SAS, NL-SAS, SSD
- DCS3700 Expansion
- 4G Read, 2.1K Write



## **DCS3700 Performance**

- Host: FC, SAS, iSCSI
- 360 SAS, NL-SAS, SSD
- DCS3700 Expansion
- 6G Read, 4.1G Write



## **DCS3860**

**New**

- Host: SAS
- 360 SAS, NL-SAS, SSD
- DCS3860 Expansion
- 12G Read, 9G Write



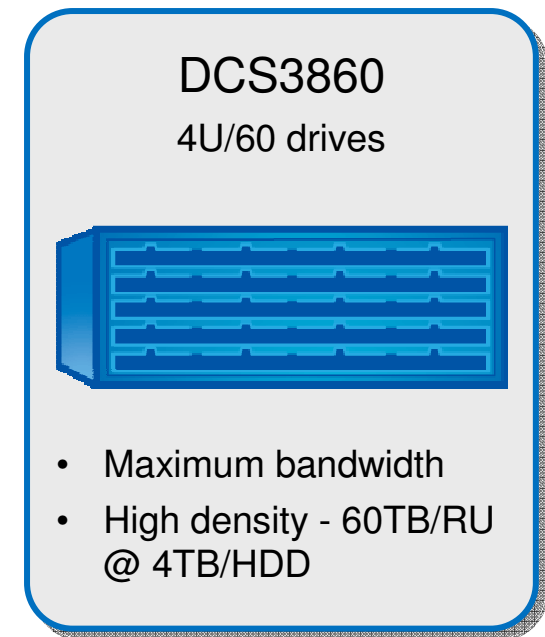
Performance metrics for sustained disk read, 512K MB/s,  
disk write, 512K MB/s- CMD

## DCS3860 Overview

*Next generation high performance storage subsystem*

### Key features

- Dual controller system design
  - 12 GB cache each (24 GB cache system total)
- 60 drive, 5 drawer enclosure (same profile as DCS3700)
  - Two 6 Gb SAS expansion connections per controller
  - Supports up to 5 expansion (EXP3800) enclosures, 360 drives
- 6Gb SAS drive ports support up to 360 SAS disk drives
  - SAS HDDs
  - Up to 24 SSD's
- Uses IBM DS Storage Manager 10.86 and firmware 7.86
  - Linux OS Support



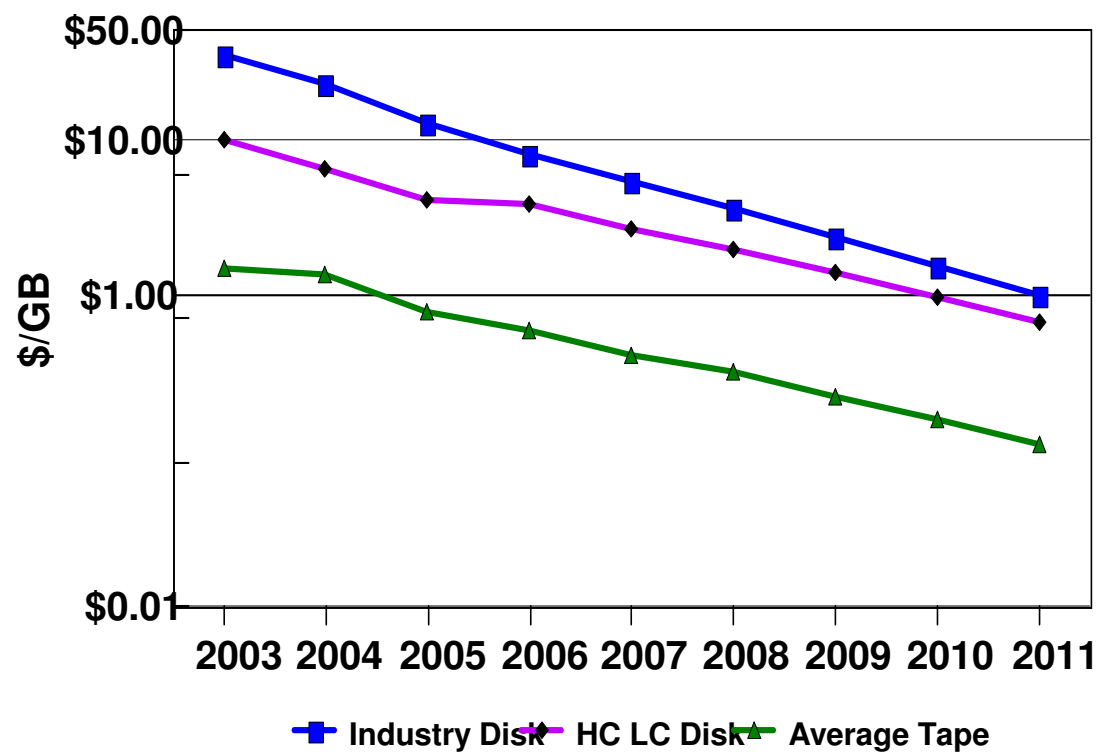
## DCS3860 Drive Capacities

Drive Type	Supported Drives
6 Gb SAS 2.5"	300 GB 15,000 rpm
6 Gb NL-SAS 3.5"	4 TB 7,200 rpm

*\* Max of 24 SSDs per system*

## Average Storage Cost Trends

### Projected Storage Prices

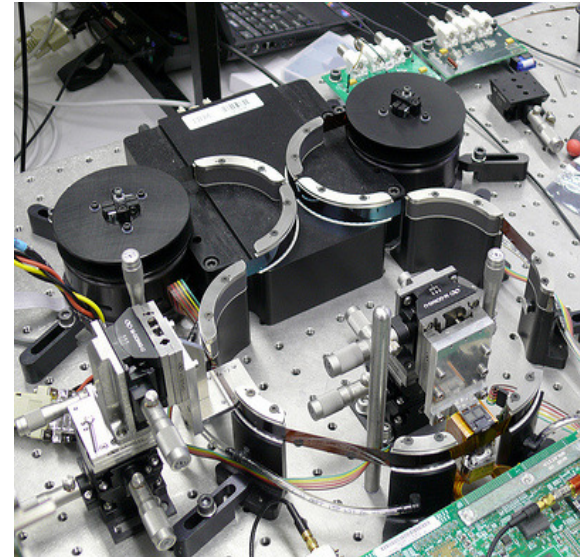


Source: Disk - Industry Analysts, Tape - IBM

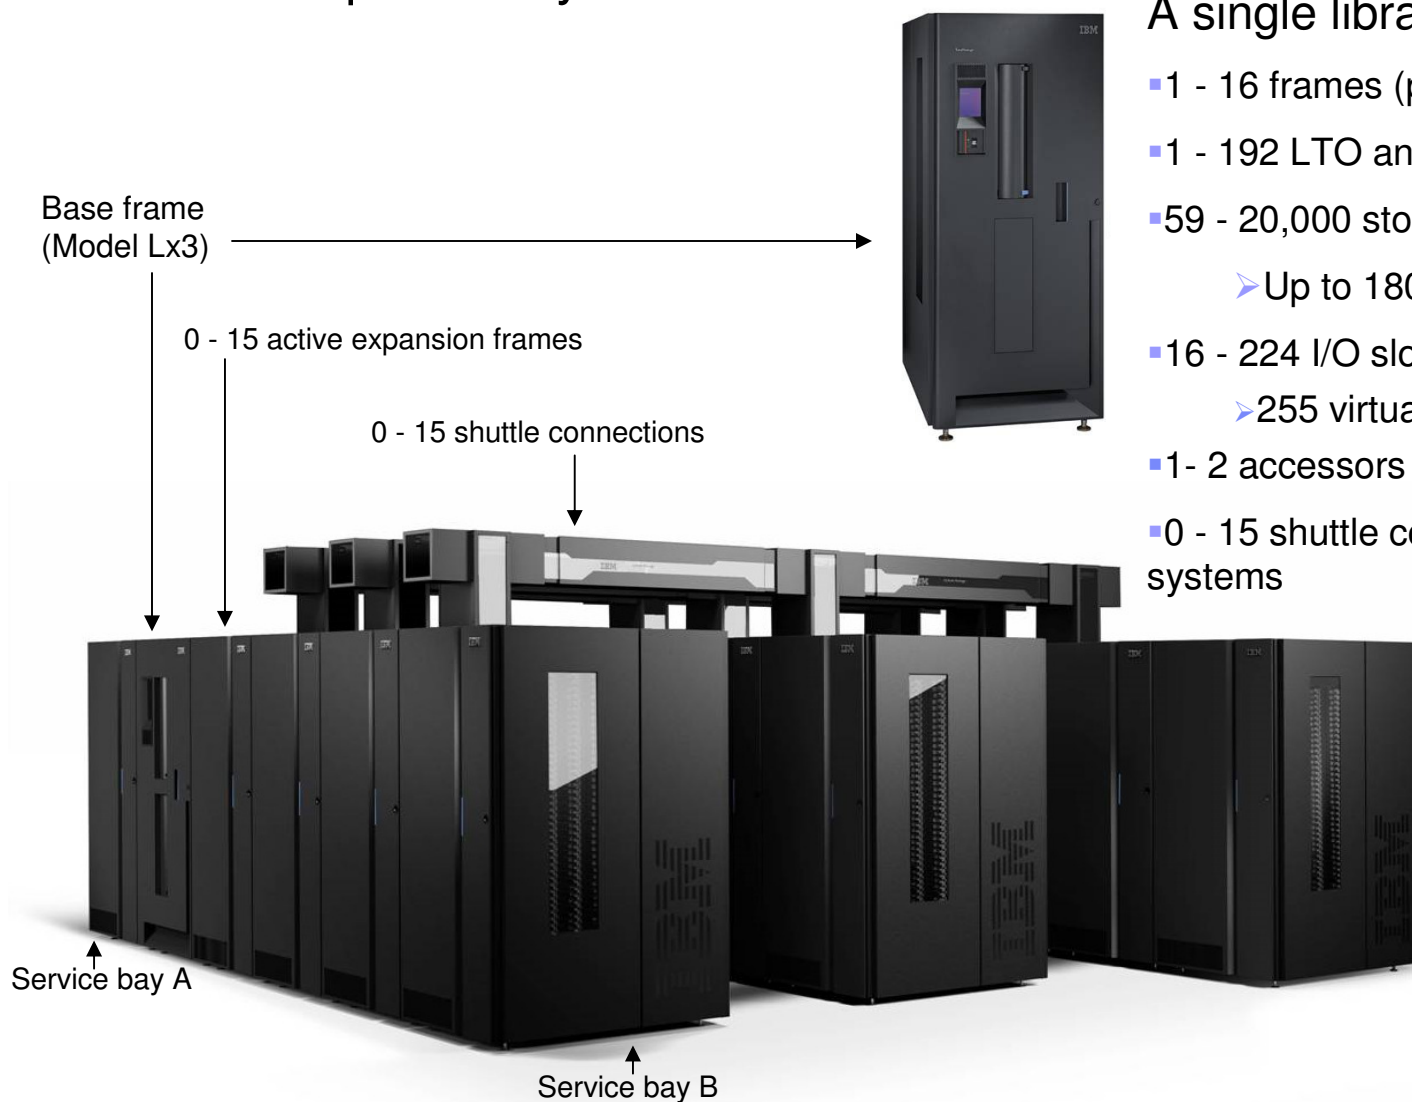


# IBM Tape Drive Strategy

- IBM has demonstrated a Tape Technology Pipeline
  - 1 TB in 2002, 8TB in 2006 , 35 TB in 2010
  - Demonstrates unconstrained capacity and performance growth path for tape technology
  - Two product lines based on this technology pipeline
- 3592 enterprise tape product line
  - Reliability, Performance and Function differentiation
    - Move to 32-channel technology and higher capacities
  - Enterprise media cartridge with reuse
  - Longer support cycles for media, format and hardware
    - Can still read the first Jaguar tape written in 2002
- LTO midrange product line
  - Streaming device model
    - Cost-centric model
  - New media each generation
  - TPC Consortium driven development/function
- Over 1.5M combined units shipped



## TS3500 Tape Library Overview

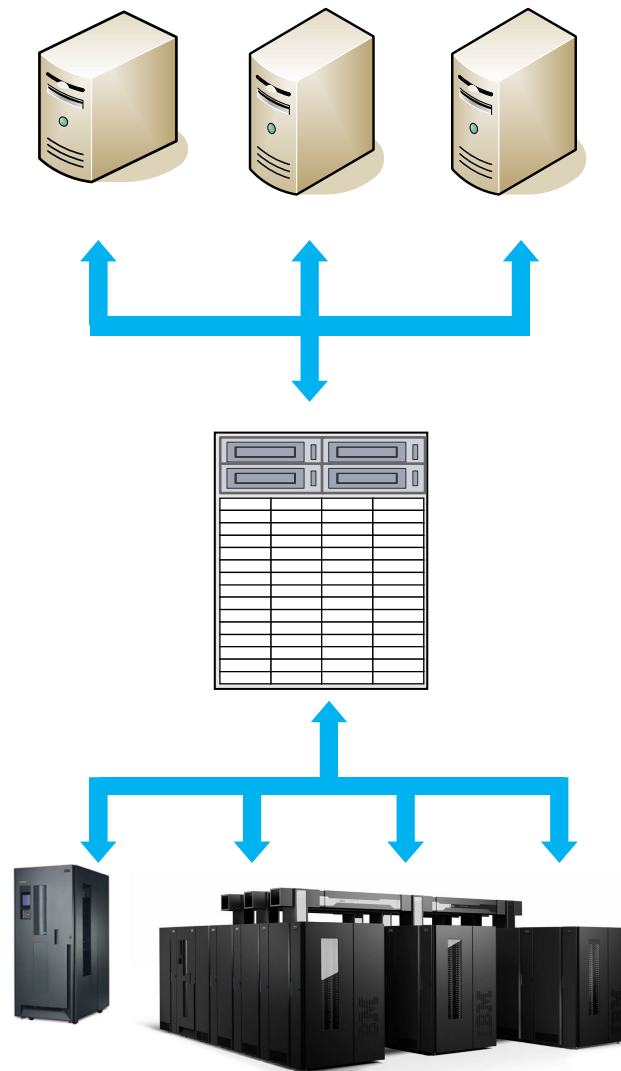


A single library system includes:

- 1 - 16 frames (plus two service bay frames)
- 1 - 192 LTO and/or 3592 tape drives
- 59 - 20,000 storage slots
  - Up to 180 PB\* Storage Capacity
- 16 - 224 I/O slots (1 - 14 I/O Stations)
  - 255 virtual I/O slots per logical library
- 1- 2 accessors (robotics)
- 0 - 15 shuttle connections to other library systems

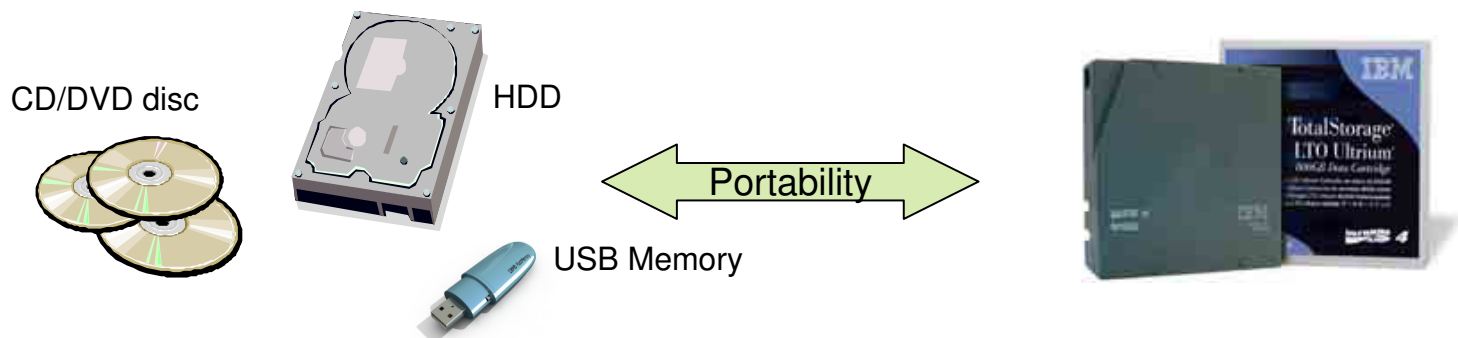
## TSLM – Tape System Library Manager

- Consolidate
- Simplify
- Exploit the benefits of the TS3500 Tape Library Shuttle complex by enabling IBM Tivoli Storage Manager and other ISVs



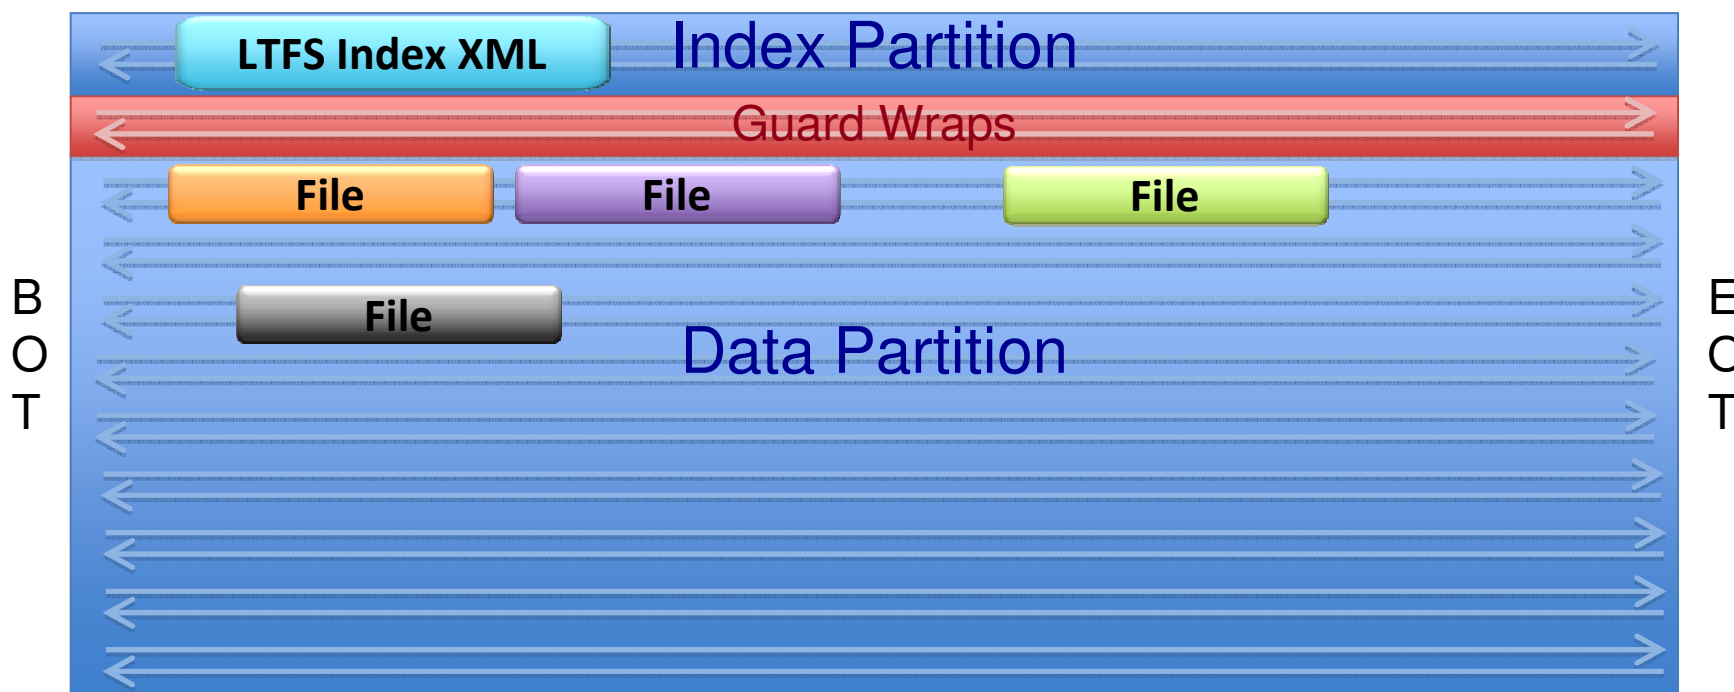
## LTFS: What is the Linear Tape File System?

- Self-describing tape format to address tape archive requirements
- Implemented on dual-partition linear tape (LTO-5)
- Makes tape look and work similar to other removable media
  - File and directories show up on desktop and directory listing
  - Share data across platforms
  - Drag&Drop files to/from tape
  - Self Describing Tape Format (SDTF) in XML-Architecture
  - Simple, one-time installation
- Developed by IBM



## LTFS on tape

- LTFS enables File System access against tape device
- LTFS utilizes media partitioning (new to LTO Gen 5 and Jag 4)
- The tape is logically divided “lengthwise”
  - (think C: & D: drives on single hard disk unit)
- LTFS places the index on one partition and data on the other

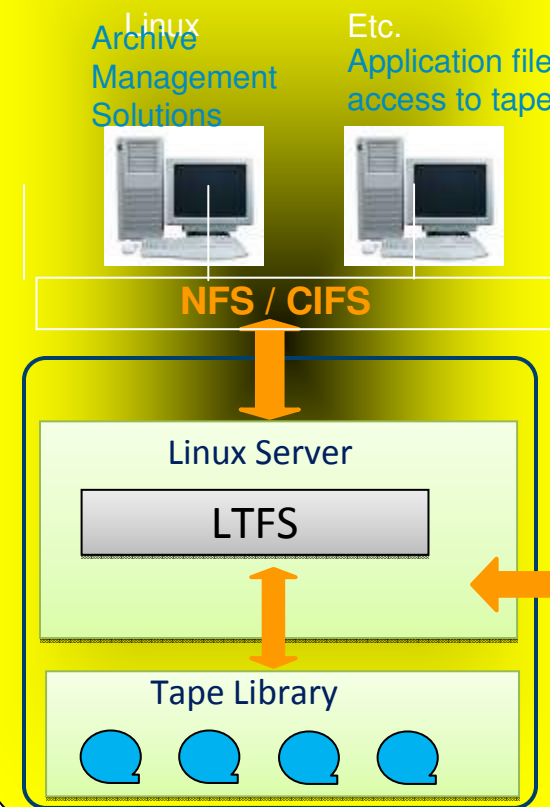


# LTFS – Product Roadmap

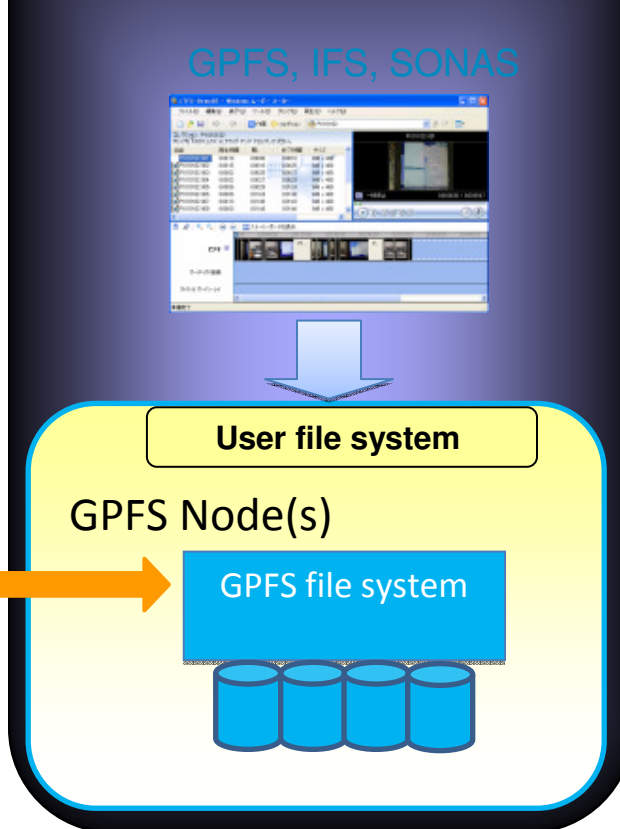
## LTFS Phase 1 LTFS Format Enablement Single Drive Support (2010)



## LTFS Phase 2 Digital Archive Enablement Tape Automation Support (2011)



## LTFS Phase 3 Integrated Solution Enablement GPFS / LTFS Support

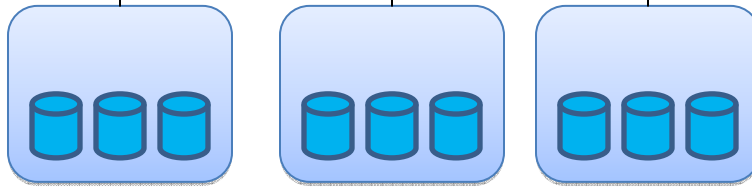


## The Problem – Network Disk Growth...

Operational

- Manageability
- Cost
- Data mix - Rich media & databases, etc
- Uses – active, time sensitive access & static, immutable data

C:/user defined namespace



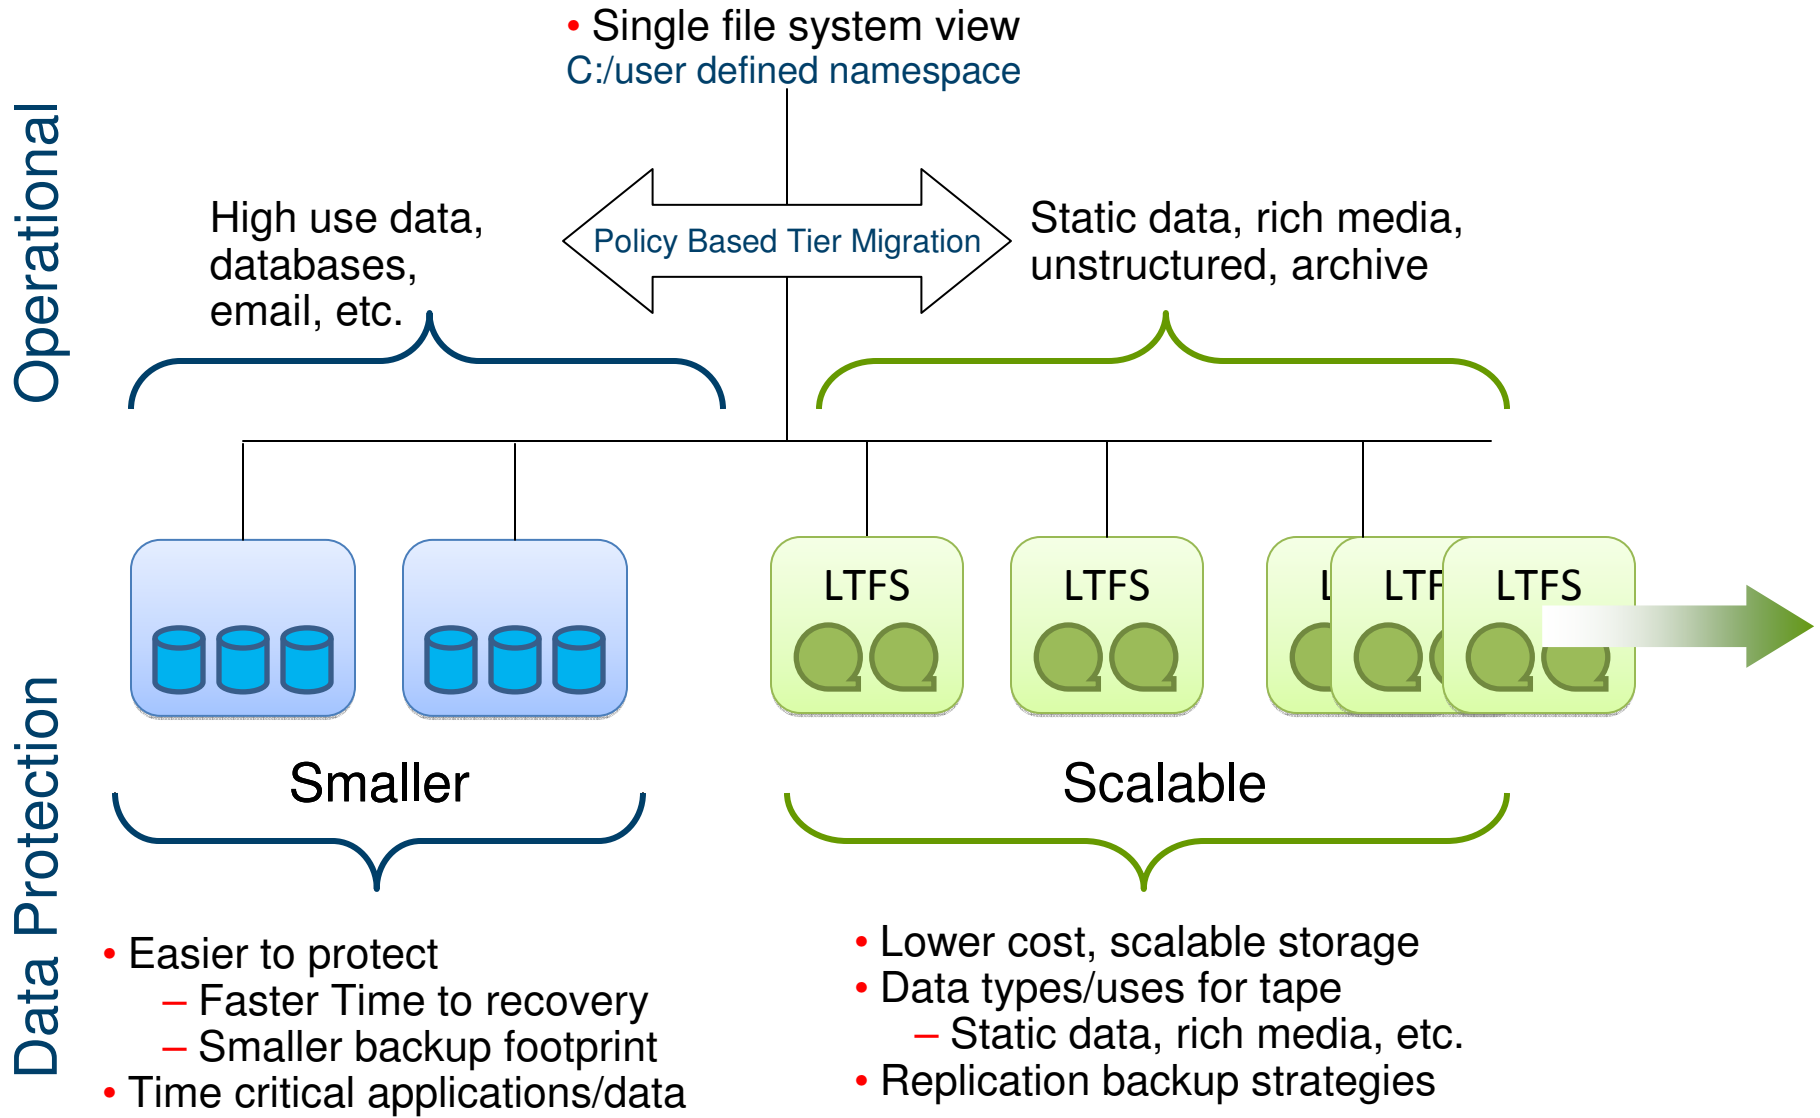
Large

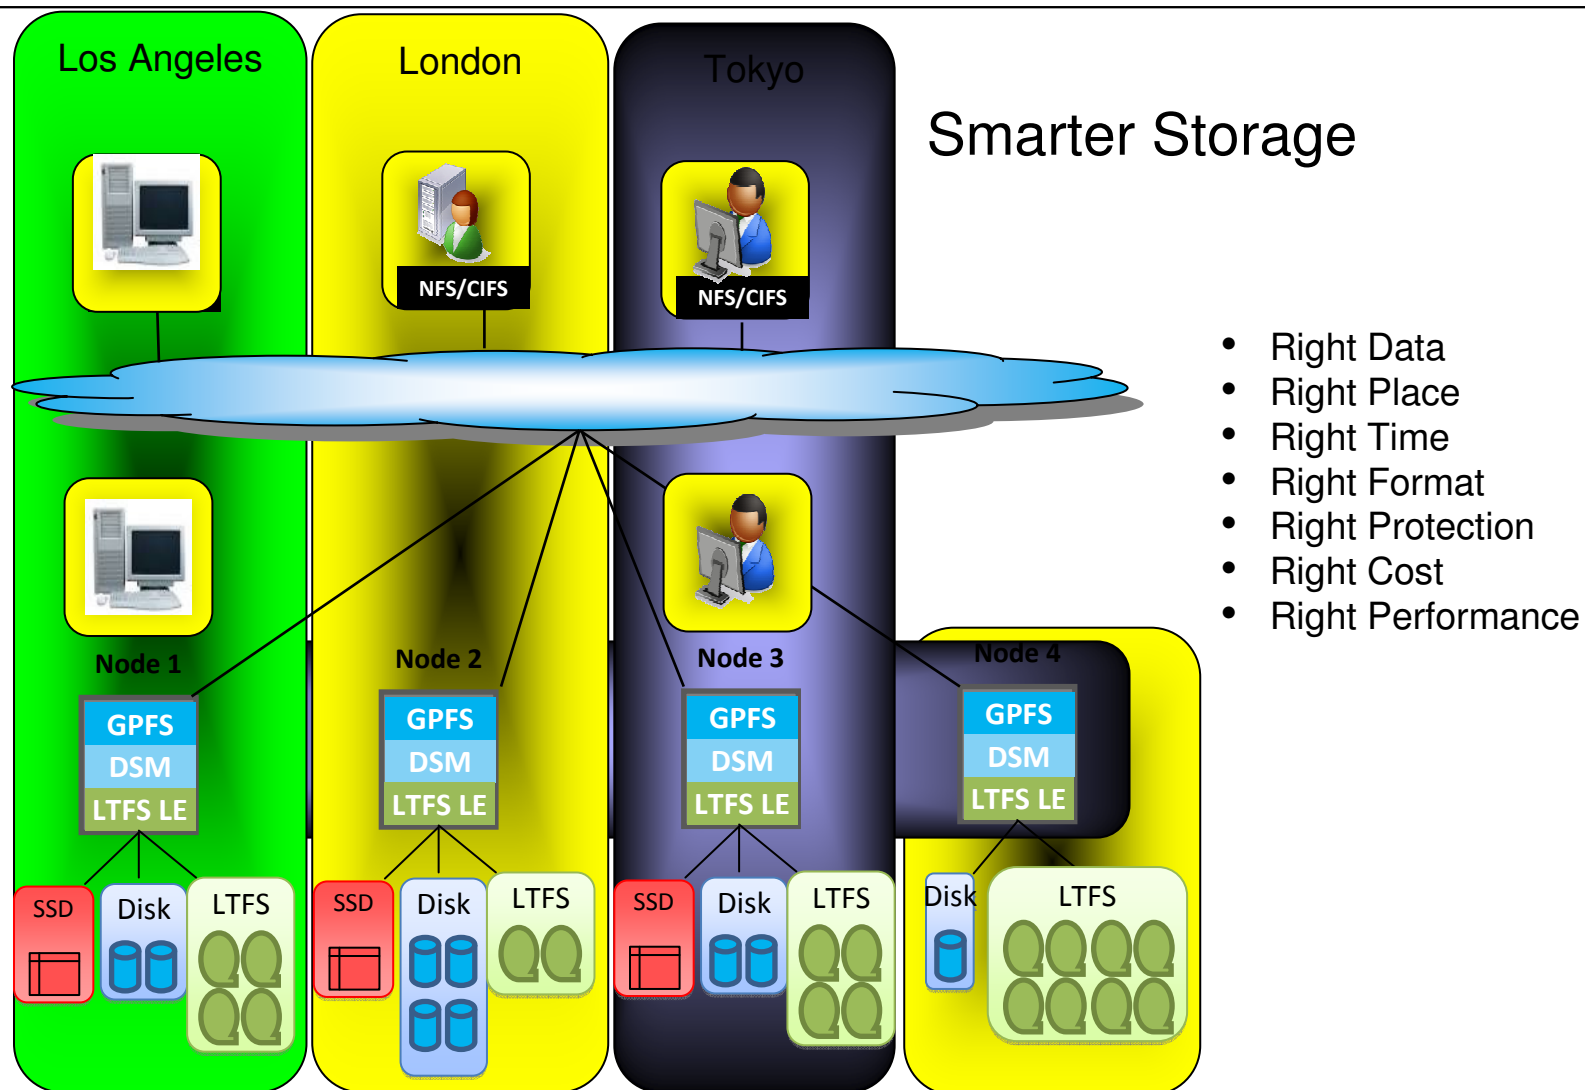
And Growing Bigger

Data Protection

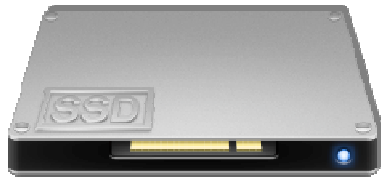
- Difficult to Protect / Backup
  - Cost
  - Backup windows
  - Time to recovery
- Data mix reduces effectiveness of compression/dedupe

# The Solution – Tiered Network Storage





## A Future with No Spinning Disk at All?

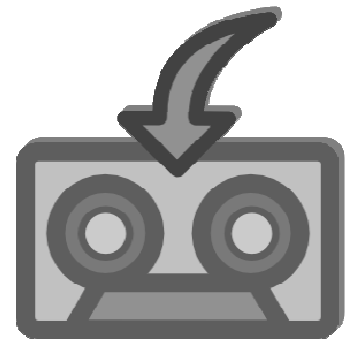


Flash

- Best overall \$/IOP for performance
  - Highest Performance
  - Immediate Access
  - Most Critical Data
  - Low Power Consumption



- Best overall \$/GB for long term data retention
  - Lowest cost
  - Easily accessible
  - Highly scalable
  - Low Power Consumption



Tape

