Extending the Benefits of HDD

Breaking Down Walls All Storage Vendors Face

Keith Hageman

X-IO Technologies

The components of a computer

- Customers use a computer to solve a data problem…
- Storage is the foundation of the configuration

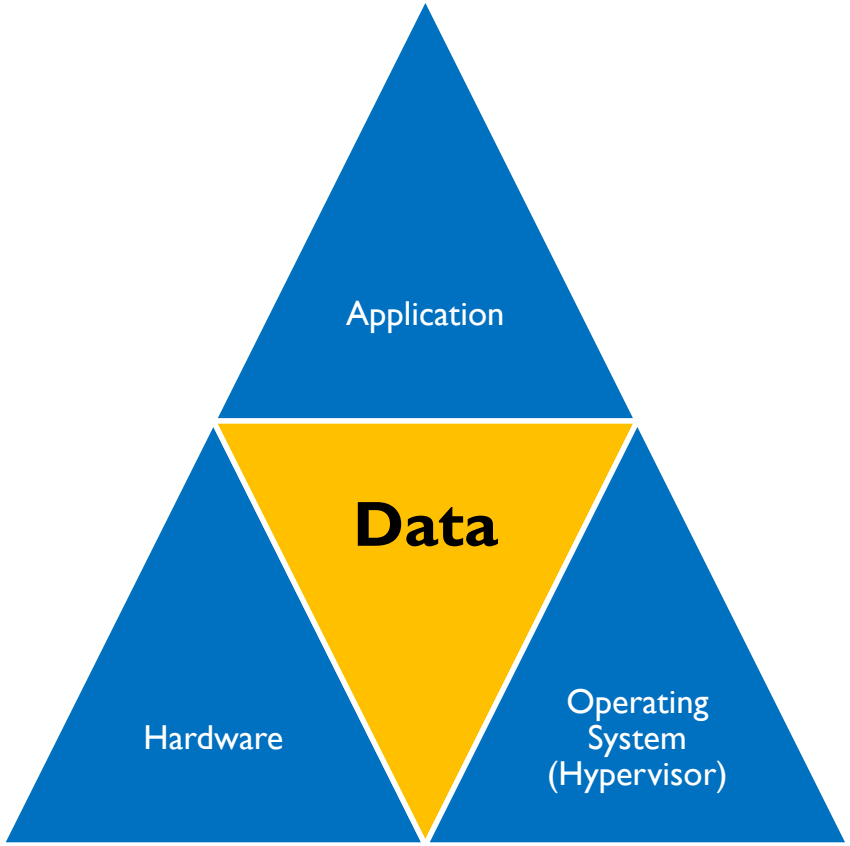

The Current Sad State of Affairs

- For over 20 years, customers have been conditioned by storage vendors to:
  - Buy on $/GB because all storage is same…
  - Capacity should not be fully used…
  - Arrays must be replaced every 3 years…
  - Assume drives just read/write and fail…
  - Availability is expensive…
  - Buy excess capacity for enough performance…

Tenants of hard drives

- Hard drives deliver:
  - Capacity
  - Reliability
  - Availability
  - Performance

- Capacity has increased over the last decade
- Reliability, Availability and Performance have *not* improved much over that time

Building a better foundation

- *The Problem* – Legacy storage has not kept pace with increasing performance and\or availability requirements that today's applications require
- Think - 24x7x365 x 5 or 6 years

Foundation problems

- Foundational issues are not being addressed by legacy SAN vendors
- AFA and Hybrid products seem to be repeating the same limitations

Capacity

The Problem:	The Impact:
Large datasets require 100's or 1,000's of drives	Too many points of failure increase likelihood of outage
Capacity over-provisioning required by legacy vendor "best practices"	More drives, more controllers use more power and cooling
At full capacity	Reduced performance

Reliability

The Problem:	The Impact:
Vibration & Cooling	Reduced performance, premature drive failure
Too few drives	Not reliable enough
Too many drives (100's to 1,000's)	Too many points of failure increase the likelihood of outages
Mixing different drive manufacturers and models	Controllers have to settle on least common denominator
50-85% of all drives replaced are NTF when they get back to the factory	Significant application disruption, cost and risk caused by drive replacements

Availability

The Problem:	The Impact:
100's or 1,000's HDDs	5% AFR is 4 drive failures per month… (1 rebuild per week)
At slightest hint of trouble, drives are failed. "When in doubt, throw it out…"	Premature drive failure, application outage, human service event
Drives hiccup	Data outage, data loss or reduced performance

Performance

The Problem:	The Impact:
Controllers inefficiencies & legacy code-base	Reduced performance
Drive IOPS less than rated	Legacy drive shelves reduce the IOPS a drive can deliver for the application
Drives hiccup	RAID rebuilds impact application performance until finished

Innovation

- This is a story of Innovation – *with great results…*
- Wikipedia roughly defines:
  - "Invention" is the first creation of an idea…
  - "Improvement" is doing the same thing better
  - "Innovation" is the application of a better solution that meets new requirements
    - Innovation differs from Invention in that Innovation refers to the use of a better, and as a result, novel idea
    - Innovation differs from Improvement in that Innovation refers to the notion of doing something different

Innovation - What's Possible?

- What if we no longer ascribe to the premise of "storage as the exception" to the "lights out" data center."
- What if we stop designing around the "removable drive"?
- What if we improve the capacity utilization, vibration, and thermals?
- What if we treat storage like servers and cluster storage elements for 'Pooled, Scalable' Data Centers

Innovation to challenge the status quo

- In 2002, Seagate created their Advanced Storage Architecture group
  - 4 years development and
  - ~100 patents
- In 2006, first customer units were shipped

A new generation “super drive”

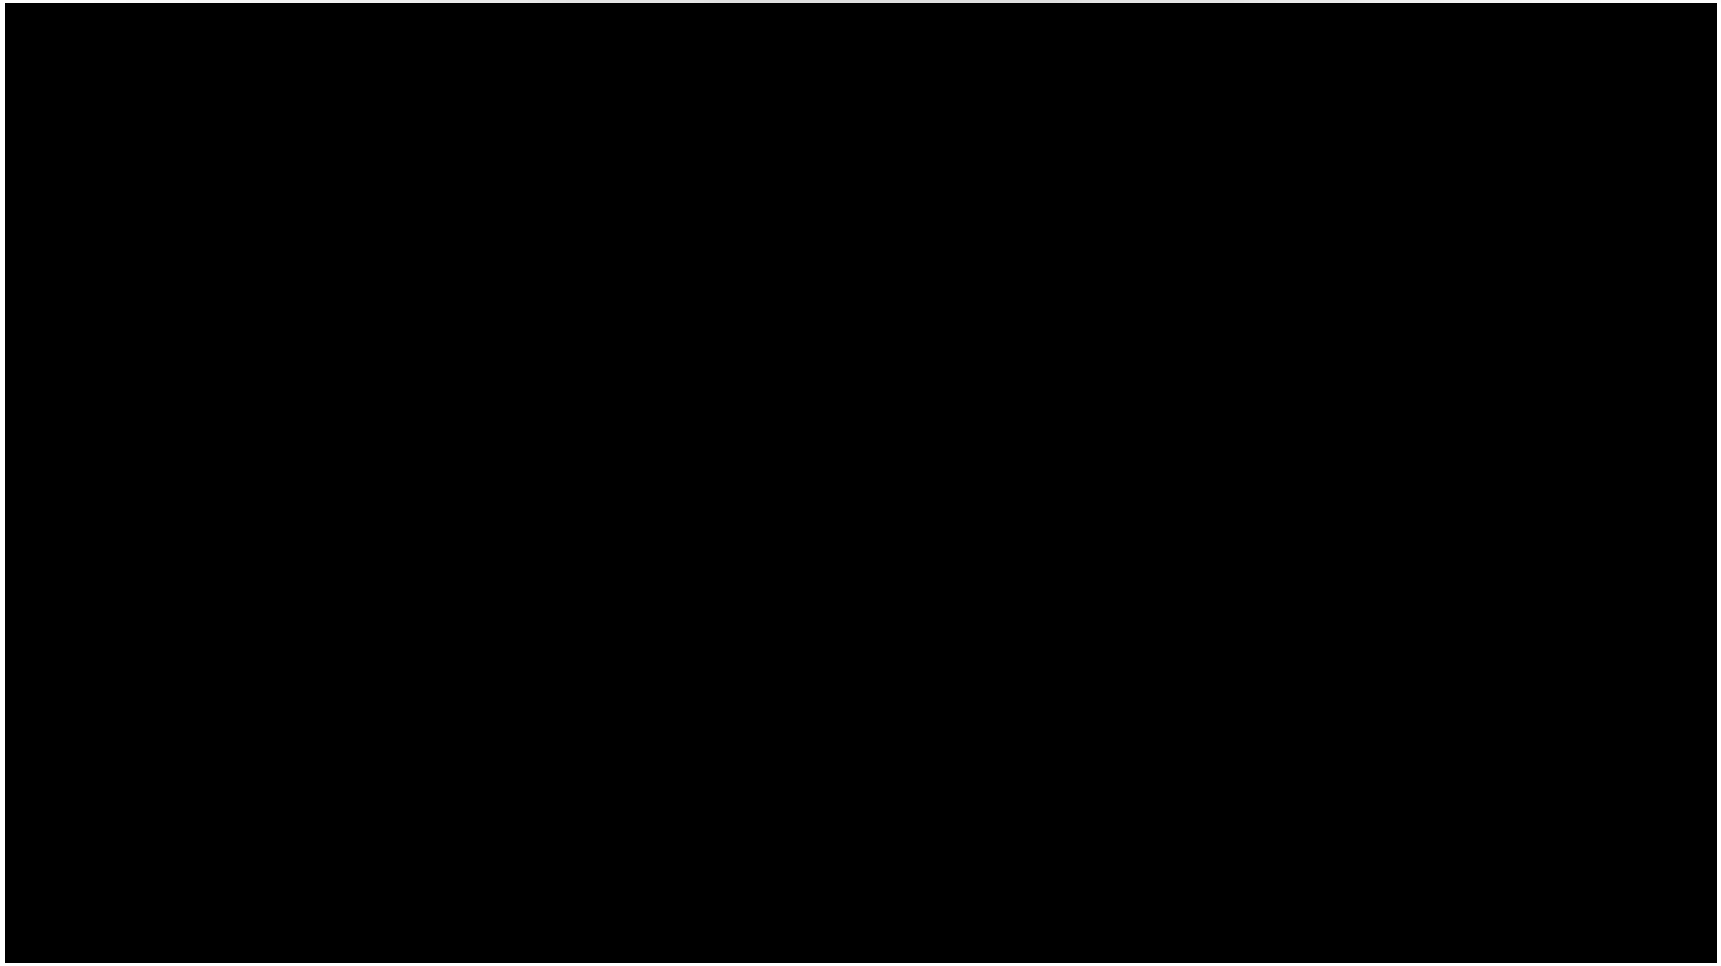

Capacity

The Solution:	The Advantage:
New design	- Enable full-capacity utilization for lower cost-per-GB - Capacity does not degrade over time - Performance does not degrade as capacity is consumed
Full capacity utilization	Fewer drives, fewer controllers requiring less power and cooling

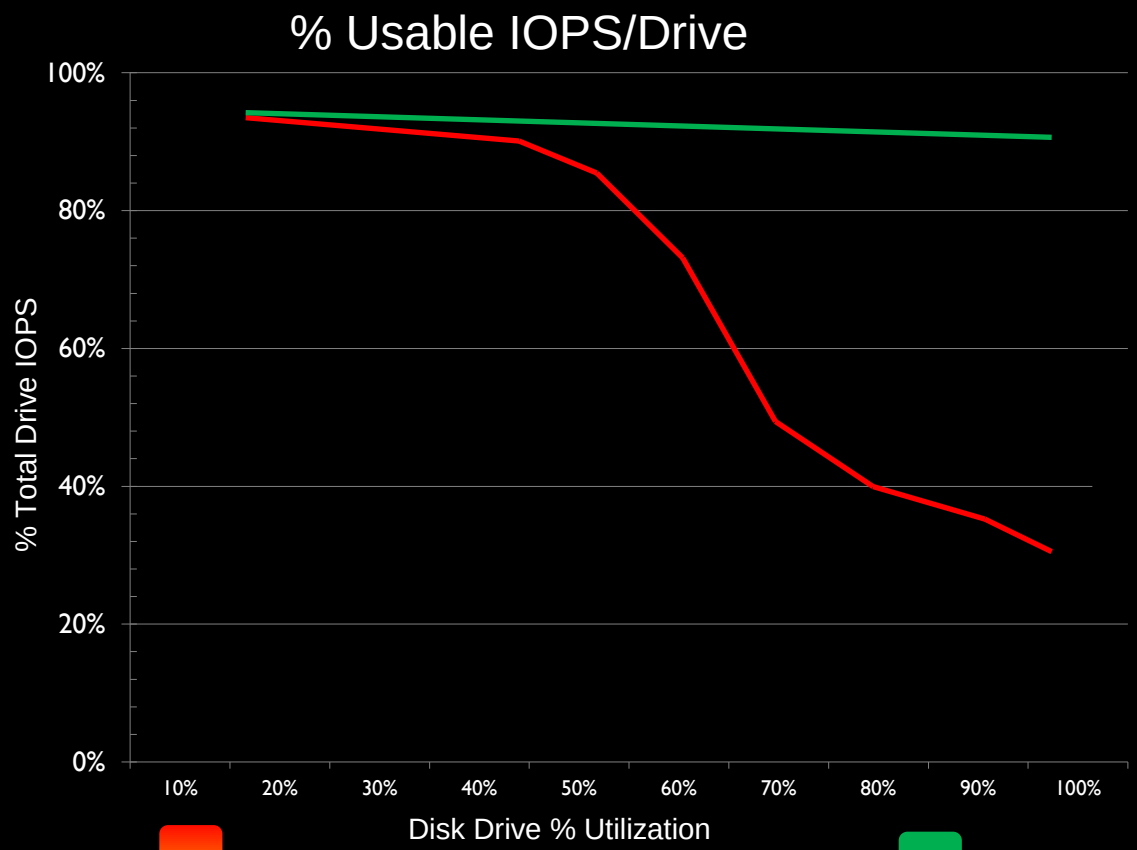

In Legacy Storage Arrays, drives typically slow down as they fill with data due to access latencies…

Unique Drive Firmware that changes the way the drives work!..

… grouping drives into a cohesive “Super Drive” with an efficient data layout pattern

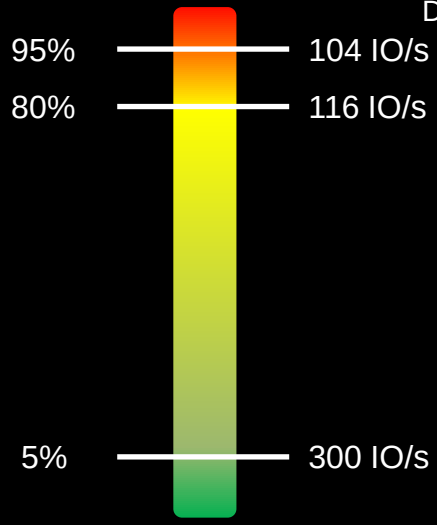

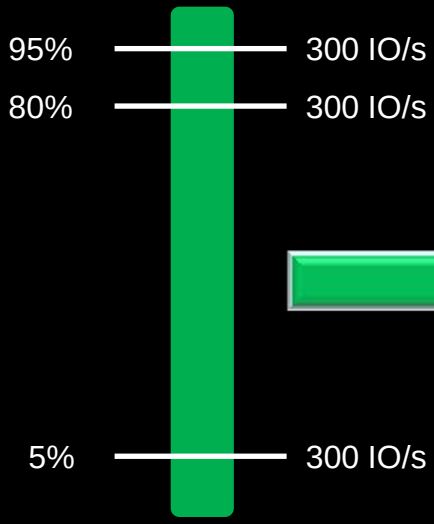

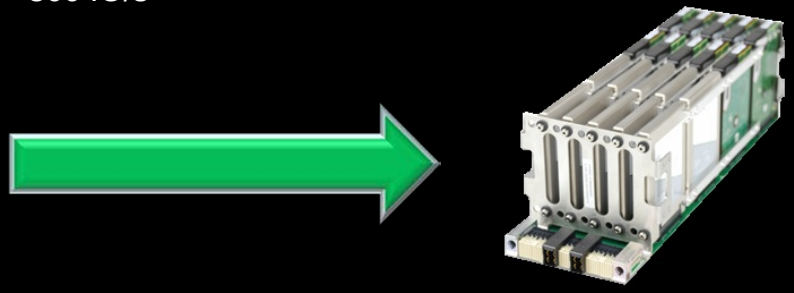

Super Drive

Reliability

The Solution:	The Advantage:
Reduced vibration & increased cooling	5 year reliability and significantly higher performance
20 or 40 drives combined together	5 years life cycles
Very large enterprise configurations	Predictable linear scalability to PBs with 5 year reliability
Single drive manufacturer and single drive models	Enables advanced analysis and repair before drive replacement required
Self-healing drives and controllers	No human service events for 5 years

Availability

The Solution:	The Advantage:
Fixed number of drives	Removes or reduces drive rebuilds, not a human service event
Drives telemetry monitored and acted upon	Self-healing drives. Drive failures are handled and repaired.
Managed Reliability. "When in doubt, check it out…"	Software determines if media issues need some automatic servicing (not a human event or a down-time event)

Performance

The Solution:	The Advantage:
Reduced vibration & increased cooling Better RAID implementation	Drives deliver 2x – 4x more IOPS
Drives issues handled with no service event	RAID rebuilds seldom impact application performance

Failed disk: SPC-1 (H450/R5)

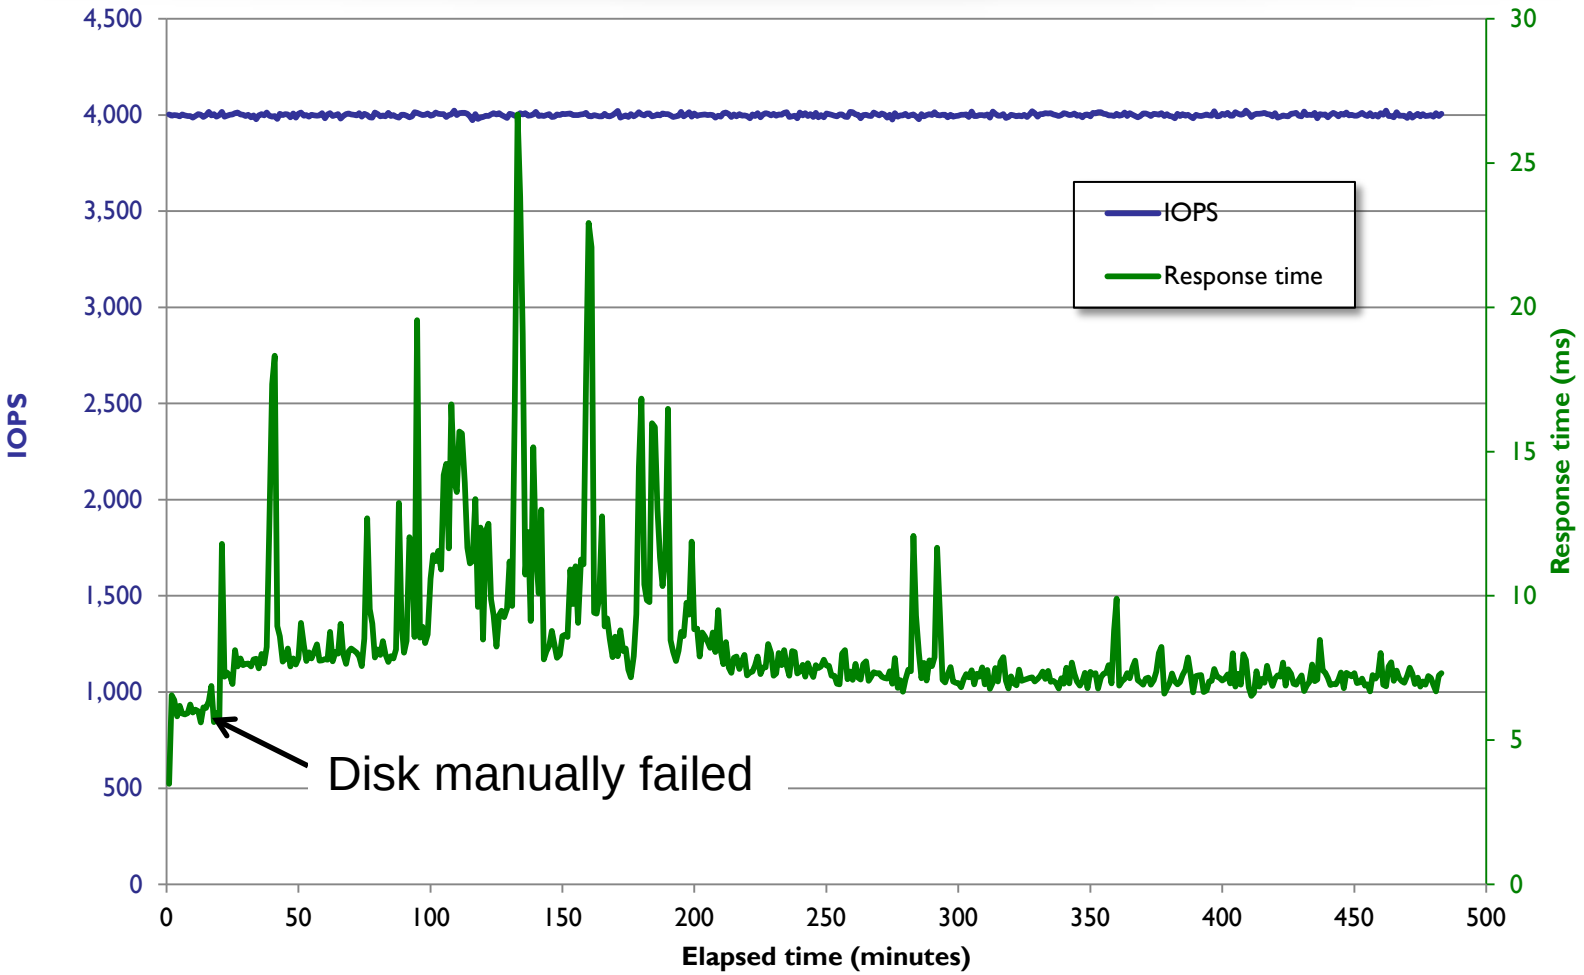

Next challenges

- Operating systems & hypervisors are implementing data management features
  - Enabling more, higher application aware capabilities
  - Cannot address low level media capabilities to maximize device utilization
- Very cool things are now possible together