Guaranteeing enterprise-class availability in a flash environment

Phill Gilbert

HPE 3PAR WW Product Manager

What is enterprise-class?

Enterprise-class flash attributes

Optimization

Critical day-one architecture decisions that make the most of flash

Work with flash vendors to implement optimizations at the system level

Use these to increase the performance and endurance of flash media

Reliability

ARR of NAND flash is significantly lower compared to HDD

Rebuilds for SSD are faster so systems are degraded for less time; rebuilds place less pressure on SSDs

Failure modes for flash are different to HDDs

MTBF are greater for SSDs

Performance

Enterprises expect predictable performance

'Ring-fenced' performance for critical applications

Provide low-latency that flash makes possible

Increased performance (IOPS / bandwidth) to enable workload consolidations

Cost

Capacity efficiency technologies to reduce initial costs

Reduced system overheads reduces the cost of flash

Increased density reduces DC requirements

Lower power and cooling costs

Ease of use

Flash in an enterprise environment

User experience

Increased revenue

Consolidation

Predictability

Datacenter footprint

Power and cooling costs

Host licensing costs

Complexity

Reliability: the cost of downtime

99.999%
5 minutes 15 seconds downtime per year

99.9999%
32 seconds downtime per year

The annual difference between 5x9s and 6x9s is 283 seconds

44,746
Items sold

1,698,000
Tweets

11,320,000
Searches

Just 5 minutes of downtime from Google in August 2013 resulted in a 40% drop in local Internet traffic

In today's always-on environment, every second counts!

Sources: http://www.internetlivestats.com/google-search-statistics/, http://www.fastcompany.com/3040445/fast-feed/amazon-sold-5-billion-items-in-2014, http://www.internetlivestats.com/twitter-statistics

Latency: the flash differentiator

0.1ms
CPU process

5ms
CPU waiting for IO from spinning disk

0.1ms
CPU process

Total: 5.2ms

0.5ms
CPU waiting for IO from flash

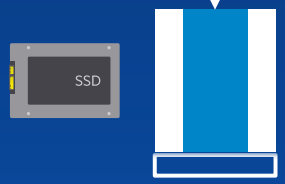

Total: 0.7ms

What does lower latency mean?
User experience is derived directly from latency; high latency means users need to wait a long time for processes to complete while low latency means less waiting time and more productivity

Flash-optimized architecture

Mesh Active Cluster

Active-Active cluster delivers sustainable performance even with high levels of capacity utilization

Fine Grained Virtualization

Three levels of storage virtualization to drive up capacity utilization and accelerate performance

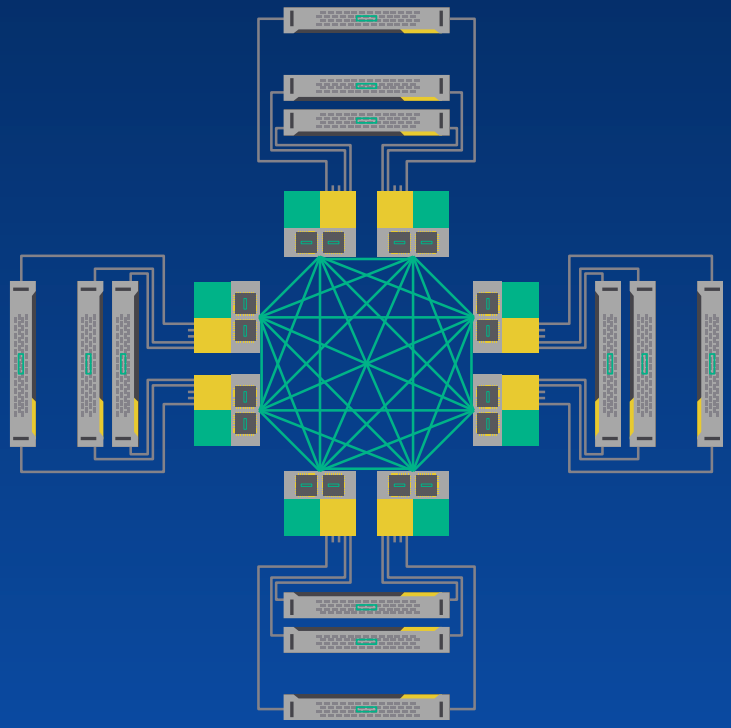

Hardware-acceleration

High-performance engine offloads CPU to boost performance and drives advanced data services

System Wide Striping

Massively parallel striping of data across all internal resources ensures high, predictable service levels

A new approach to sparing

Traditional deployment
High-OP space drive

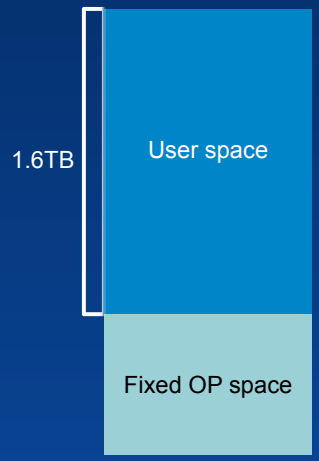

Standard deployment uses drives with large amounts of OP space to provide performance and endurance at high cost

Adaptive Sparing
HPE 3PAR Low-OP space drive

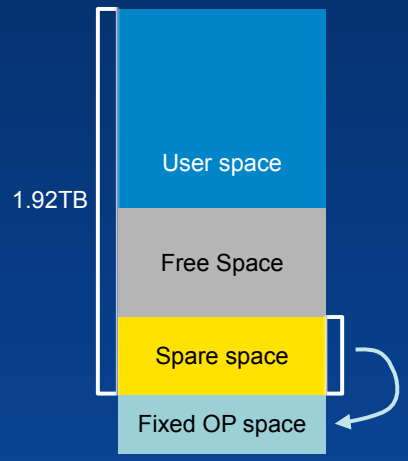

Adaptive Sparing gives the drive access to 3PAR spare space for overprovisioning, increasing endurance and performance

Adaptive Sparing 2.0
HPE 3PAR Low-OP space drive

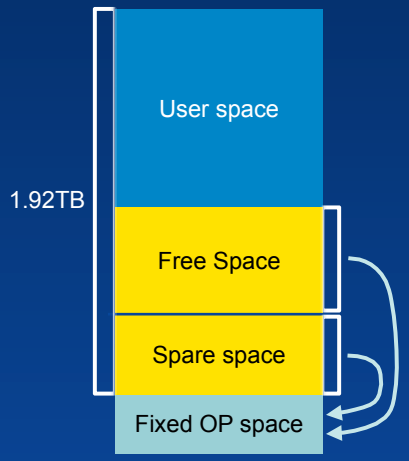

Adaptive Sparing 2.0 works with Adaptive Sparing by also giving the drive access to free system space for overprovisioning

Active/Active end-to-end

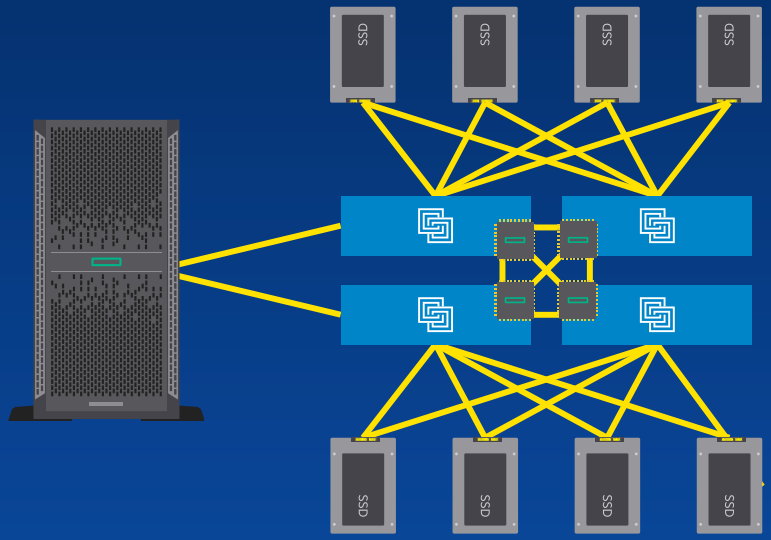

Active/Active host connections
Truly active/active connectivity with no reliance on MPIO DSM or proprietary drivers and plugins

Mesh-active controller architecture
Ensures that all system resources are evenly balanced for optimal performance, drive wear and simplified management

Active/active SSD access
Each node in a node pair has active access to all drives behind the node pair to drive performance and improve efficiency

HPE 3PAR StoreServ Availability

Persistence Cache
Predictable, high performance in case of controller reboot/shutdown

Persistent Ports
High availability in event of switch/cable failure without host side MPIO interaction

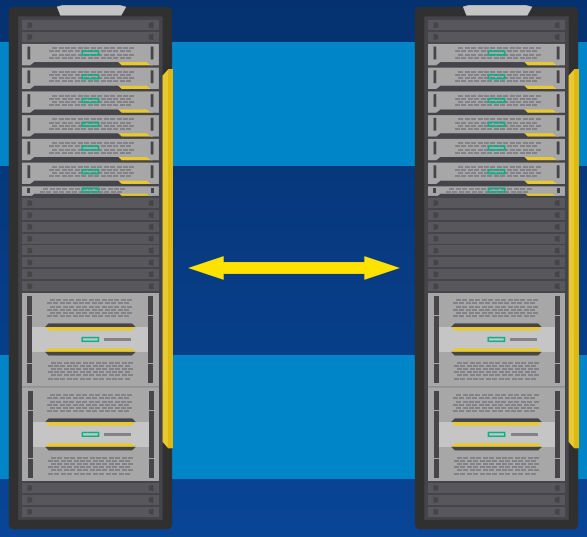

Cage availability
Seamlessly protects against entire drive cage failure or power failure

Priority Optimization
Prevents noisy neighbor and guarantees performance to critical applications

Peer Persistence
Protect against datacenter failure with automated site failover without degradation and zero recovery time

Summary

Solid state and other non-volatile memory technologies have already massively disrupted the industry

Solid state technologies don't change the base requirements for availability, high performance doesn't eliminate resilience

Key differentiators include rich data services (data mobility, protection) and integration into other infrastructure components

Solid state technologies will continue to exist as tiers in most environments including sub-LUN tiering and server-side flash

HPE 3PAR StoreServ

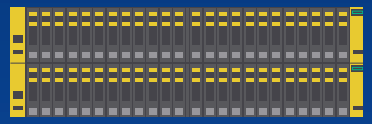

HPE 3PAR 8000

Scales to over 1m IOPS
5.5PB usable capacity
2-4 node scalability
Up to 480 SSDs
24 host ports

1

One OS

One interface

One feature set

Any workload set

Federated mobility

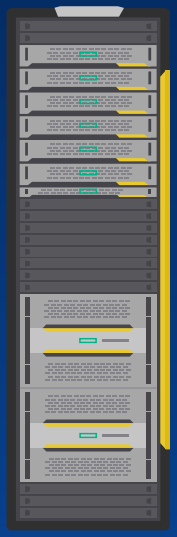

HPE 3PAR 20000

Scales to over 3m IOPS
12PB usable capacity
2-8 node scalability
Up to 1024 SSDs
160 host ports

Enterprise-class data services

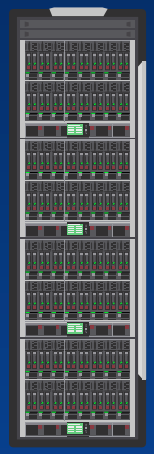

Persistent Checksum
Data consistency

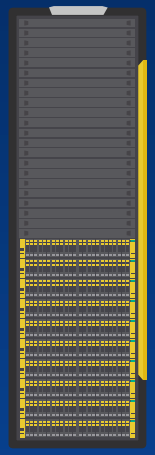

Synchronous
Zero-second RPO

Asynchronous Periodic
5 min RPO

Asynchronous Streaming
RPO in seconds

3-site replication
Zero RPO, 3 data center

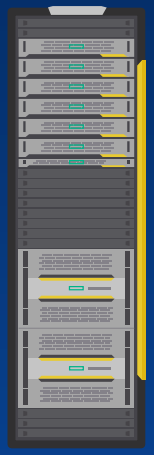

Express Protect
Backup integration

1. End-to-End Data Integrity
2. Complete set of replication options
3. Direct backup to disk-to-disk appliance