



Solid State Storage


Changing the world of IT

Priyadarshi Prasad
Sr. Product Manager
HP Storage

Pete Robinson
Manager, Storage Administration
ExactTarget, a salesforce.com company

HP Solid State Storage Strategy

- **Optimize for solid state arrays**
 - Performance, density, form factor
- **Innovate in software for solid state**
 - Rich data services layer
- **Integrate servers and storage**
 - Collaborative flash caching, software defined storage
- **Offer solid state as a tier in servers and storage**
 - Sub-lun tiering, max counts, MLC/SLC
- **Pioneer advanced development**
 - Next generation flash technologies





HP 3PAR StoreServ 7450 Storage


All Flash Array



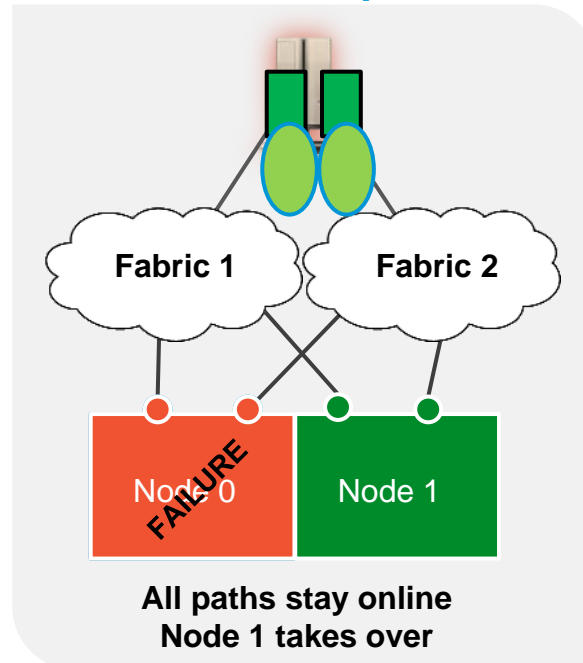
	7450 2 Node	7450 4 Node
Max. IOPS Estimate (4k Random 100% Read)	275,000+ IOPS at under 700 micro-seconds latency	550,000+ IOPS at under 700 micro-seconds latency
Max. Throughput Estimate (Read, RAID 10)	3,200 MB/s	5,200 MB/s
Connectivity Options	FC: 4-12 8Gb/s iSCSI: 0-4 x 10Gb/s FCoE: 0-4 x 10Gb/s	FC: 8-24 8Gb/s iSCSI: 0-8 x 10Gb/s FCoE: 0-8 x 10Gb/s
Cache	64 GB (32GB Data, 32GB Control)	128GB (64GB Data, 64GB Control)
SSDs Supported	100GB SLC, 200GB SLC, 400GB MLC	
Max SSD Drives and Capacity	120 SSDs, 48TB SSD Raw Capacity	240 SSDs, 96TB SSD Raw Capacity

Tier 1 Resilience

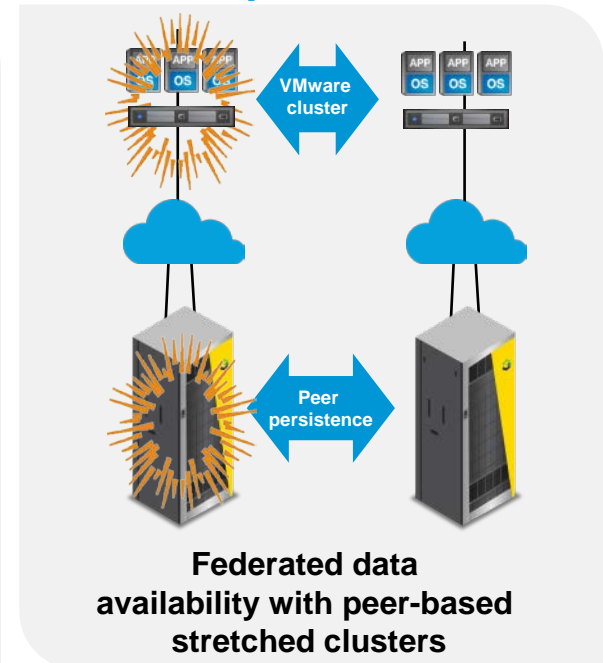
Persistent cache



Persistent ports



Peer persistence



Q. Why is resilience important?

A. Because IOPS from a machine that is down is zero.

Customer Feedback – ExactTarget

- Global marketing SaaS leader
- Based out of Indianapolis
- Offices in US, Europe, Australia
- 3PAR customer since 2006



Customer Feedback – ExactTarget


- **Buying Criteria for all flash array**
 - Performance
 - Application acceleration
 - Business Impact (ROI)
 - Resiliency
 - Architectural maturity
 - Advanced data services
 - Ease of deployment/usage
 - Cost – acquisition and maintenance cost
- **Evaluation results**
- **Use Cases & Deployment scenarios**

Data mobility to manage SLAs at a datacenter level

Hosting provider using federation to non-disruptively manage customers to variable SLAs

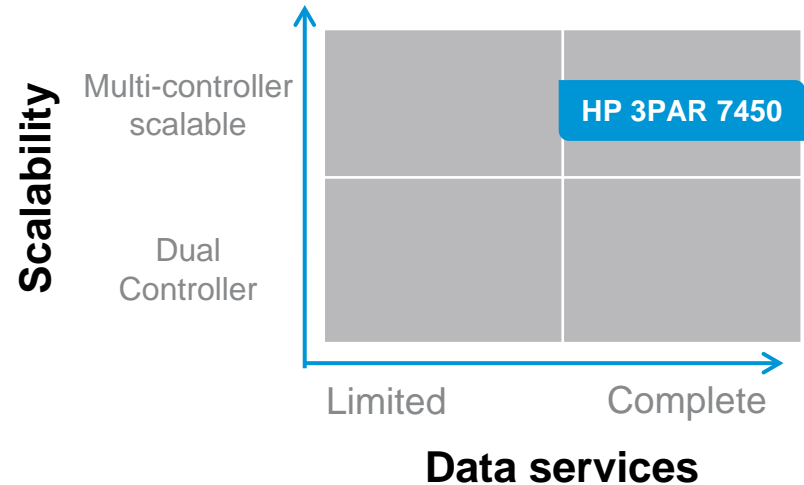
HP Peer Motion + Priority Optimization + Adaptive Optimization

- Move workloads to the right resources
- Manage service levels at a datacenter



Performance Without Compromises

	3PAR 7450
Non-disruptive node/ performance scalability	✓
Online upgrades w/o write-thru	✓
T10 DIF support	✓
Quality of service	✓
Transparent failover (vMSC)	✓
3-site replication	✓
Data-at-rest encryption	✓
Federation support	✓





Thank You!






HP 3PAR StoreServ Delivers Polymorphic Simplicity



- Operating System
- User Interface
- Command Line
- Feature Set


When Value Matters
Starting at \$25K



When Performance Matters
Over 550K IOPS




When Scale Matters
Up to 2.2 PB



OLTP: High frequency trading

Microseconds = millions of dollars

High-frequency traders generated about \$21B* in 2008



Business analytics

How fast can I churn the data?

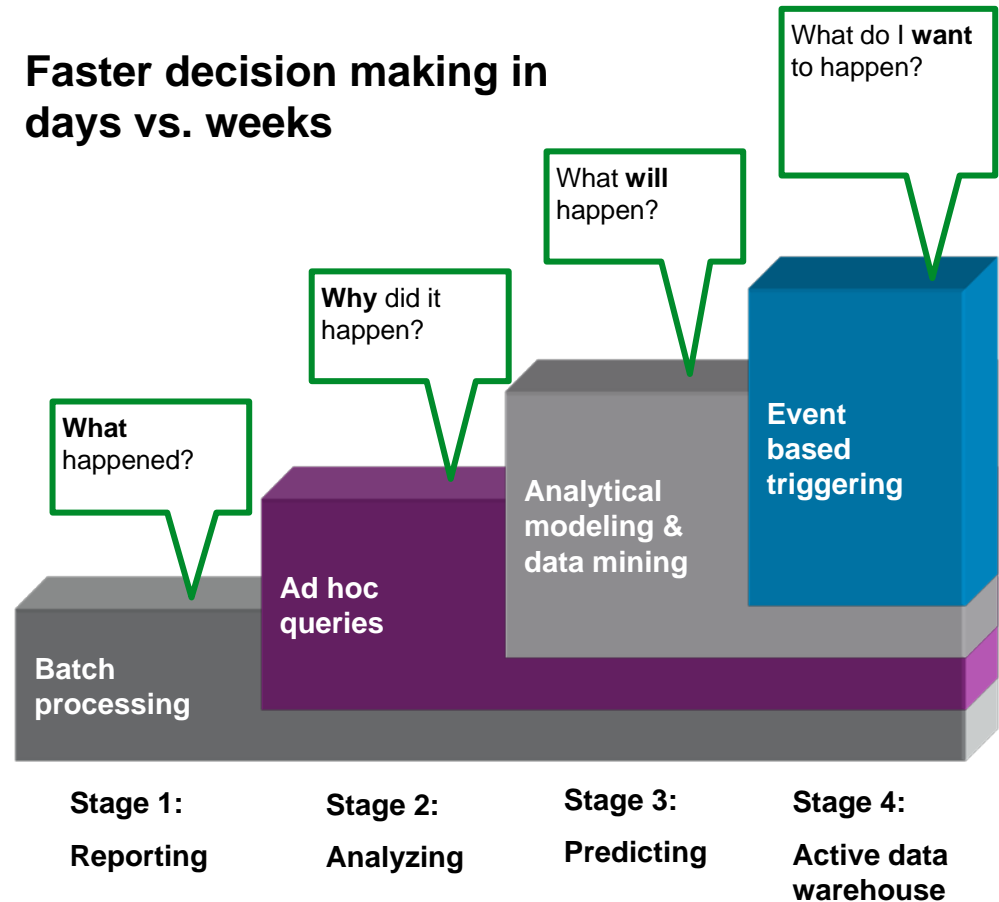
Data warehousing environment

- 10-15 million transactions a day
- Challenge
- Address the bottleneck caused by processing and consolidating huge volumes of transactions into production data mart

Retail example benefits

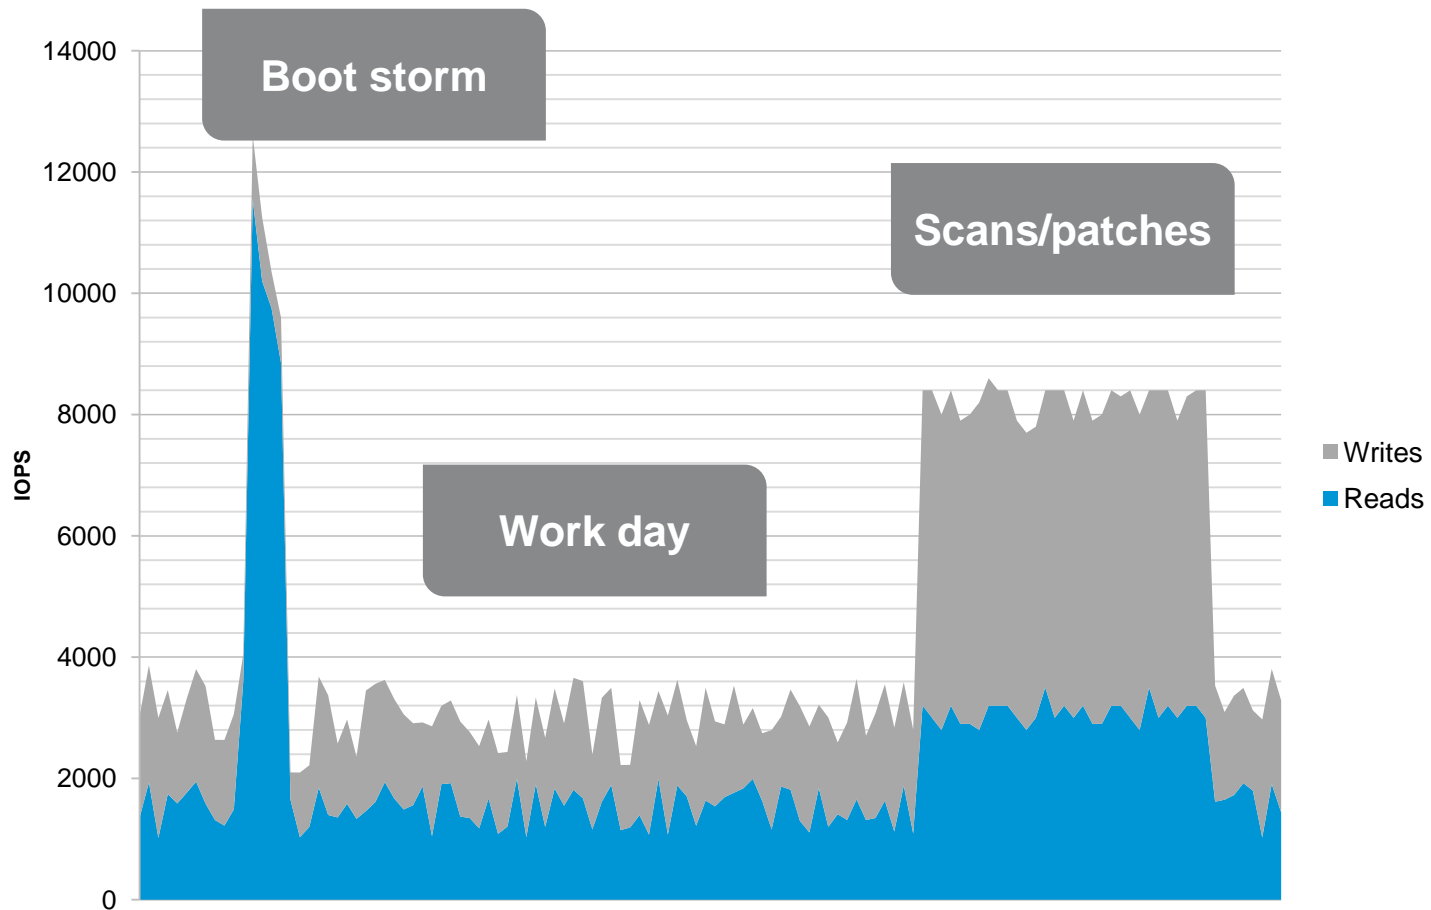
- 3x faster batch processing
- Analyze data and information in real time
- Decision making in days vs. weeks

Faster decision making in days vs. weeks



VDI demands responsive, high-performance storage

VDI storage requirements for hundreds of desktops



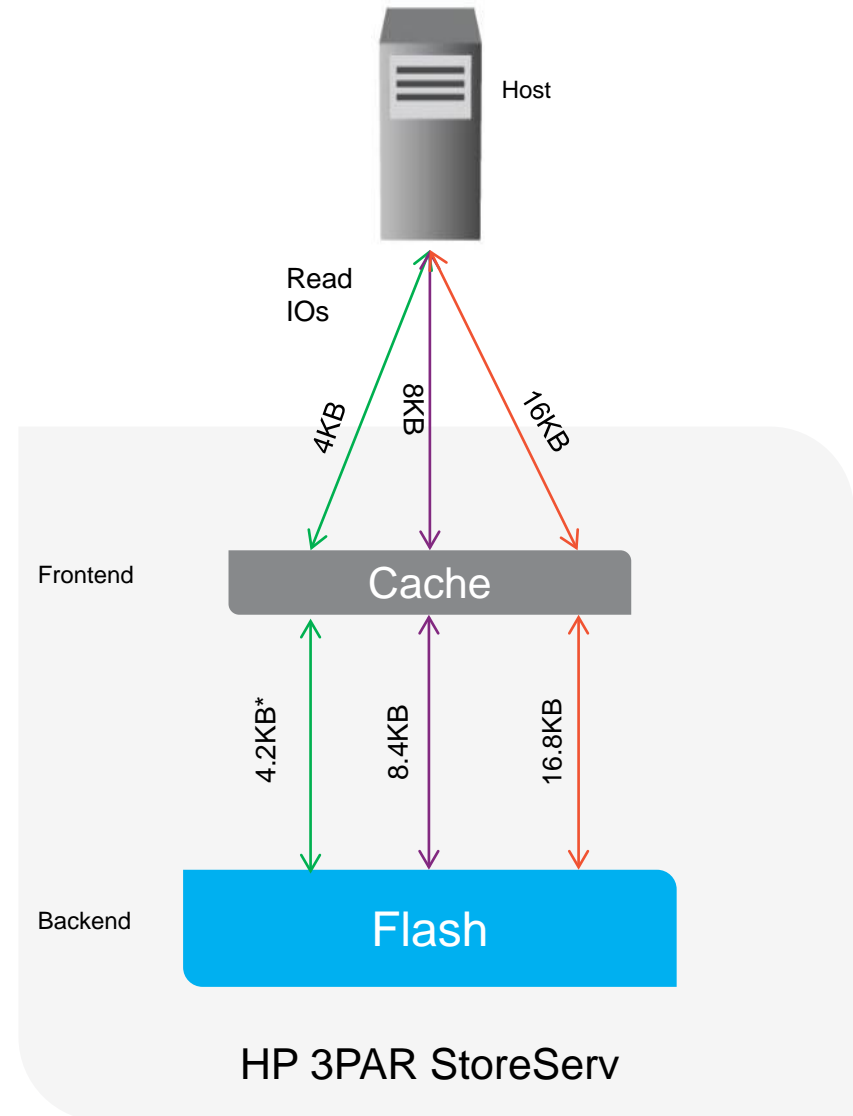
Adaptive read

Read optimization – from flash to cache

- 3PAR architecture adapts its reads from flash media to match host IO sizes

Benefits

- Reduced latency by avoiding unnecessary data reads
- Optimized backend throughput handling



*Extra bytes to account for DIF




Adaptive write

Write optimization to cache

- 3PAR architecture supports a granular cache page size of 16KB
- However, if a sub-16KB write IO comes, 3PAR does a sub-16KB write to cache
- 3PAR keeps a bitmap for each page and only the dirty part of the page

Benefits

- Reduces latency and backend throughput as well as extends flash life by avoiding unnecessary data writes
- For R1 volumes, adapting writes to match IOs avoids latency penalties associated with read-modify-write sequences



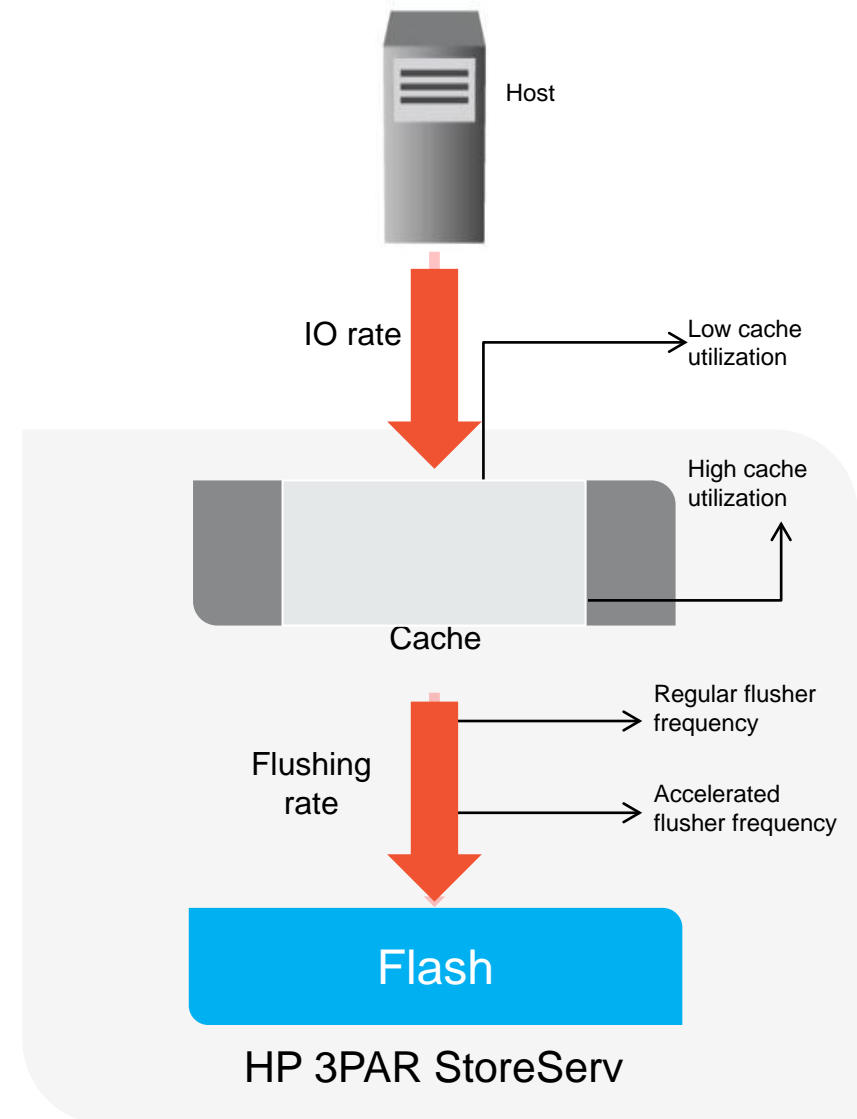
Autonomic cache offload

Workload-driven cache algorithm

- 3PAR architecture autonomically changes cache-to-media offload frequency based on utilization
- 3PAR also keeps track of read cache hits and keeps hot data in cache itself
 - This increases cache-hit rate for frequently access data, lowering latency

Benefits

- Allows servicing hundreds of thousands of IOPS at lower latencies



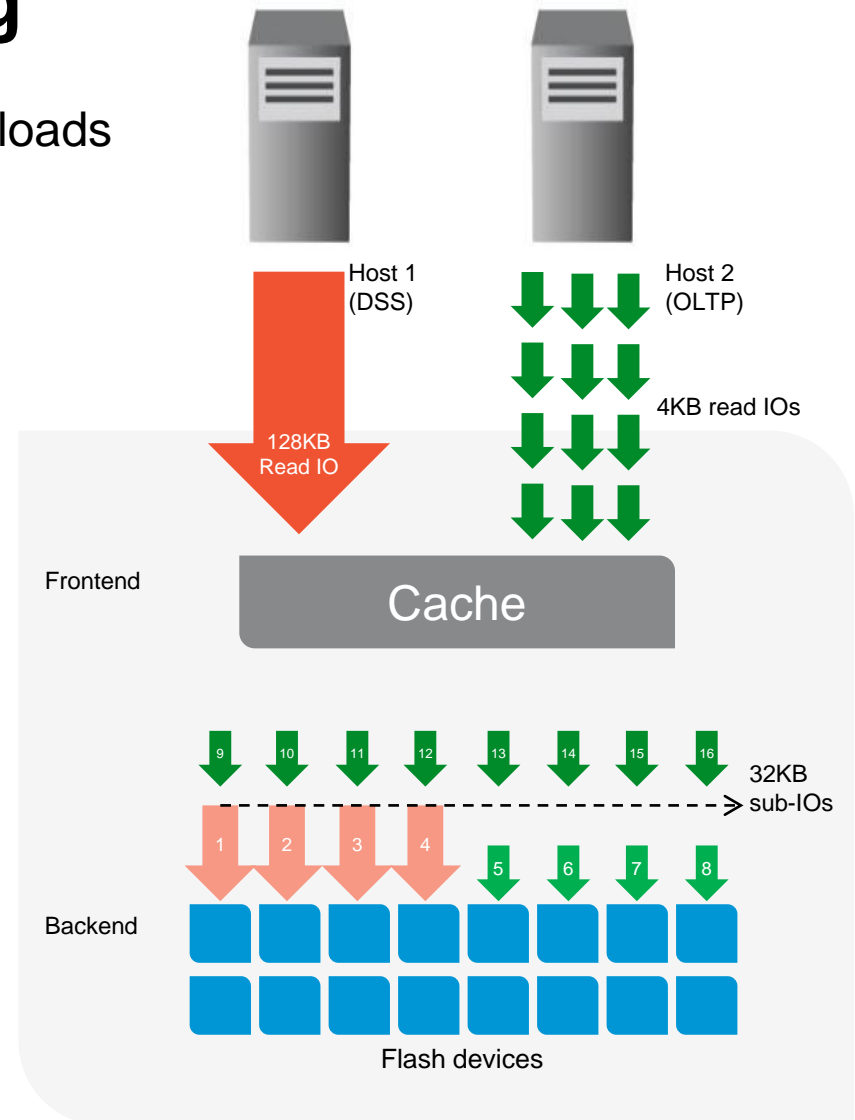
Multi-tenant I/O processing

Maintaining service levels under mixed workloads

- Frontend:
 - Cache to media write process is multi-threaded, allowing for each IO to start its own thread in parallel without waiting for threads to get free
- Backend:
 - 3PAR architecture splits large R/W IOs into 32KB sub-IOs before sending them to flash media
- Ensures that smaller read IOs do not suffer from higher response times

Benefits

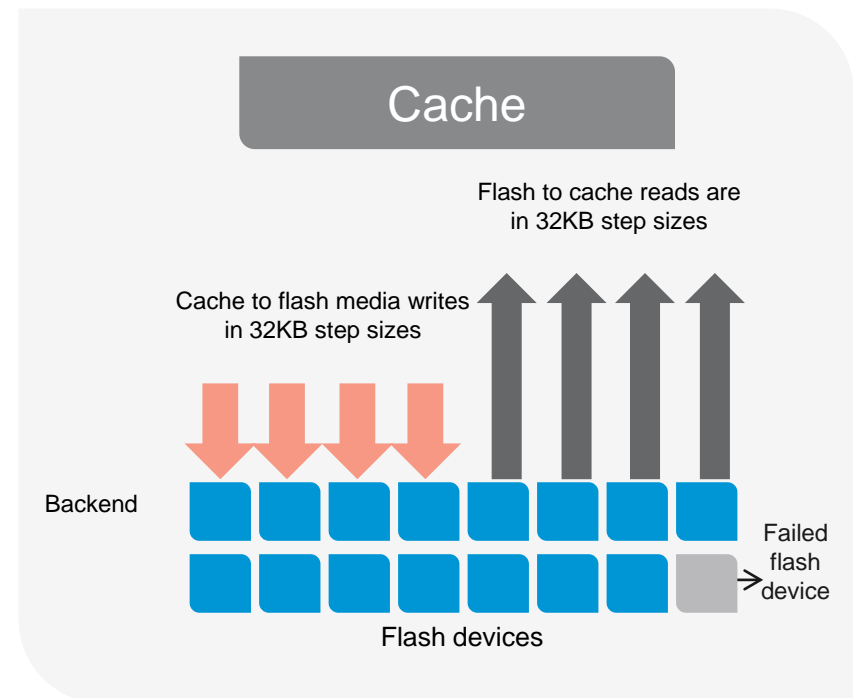
- Allows 3PAR to serve sequential workloads without paying a latency penalty on OLTP workloads



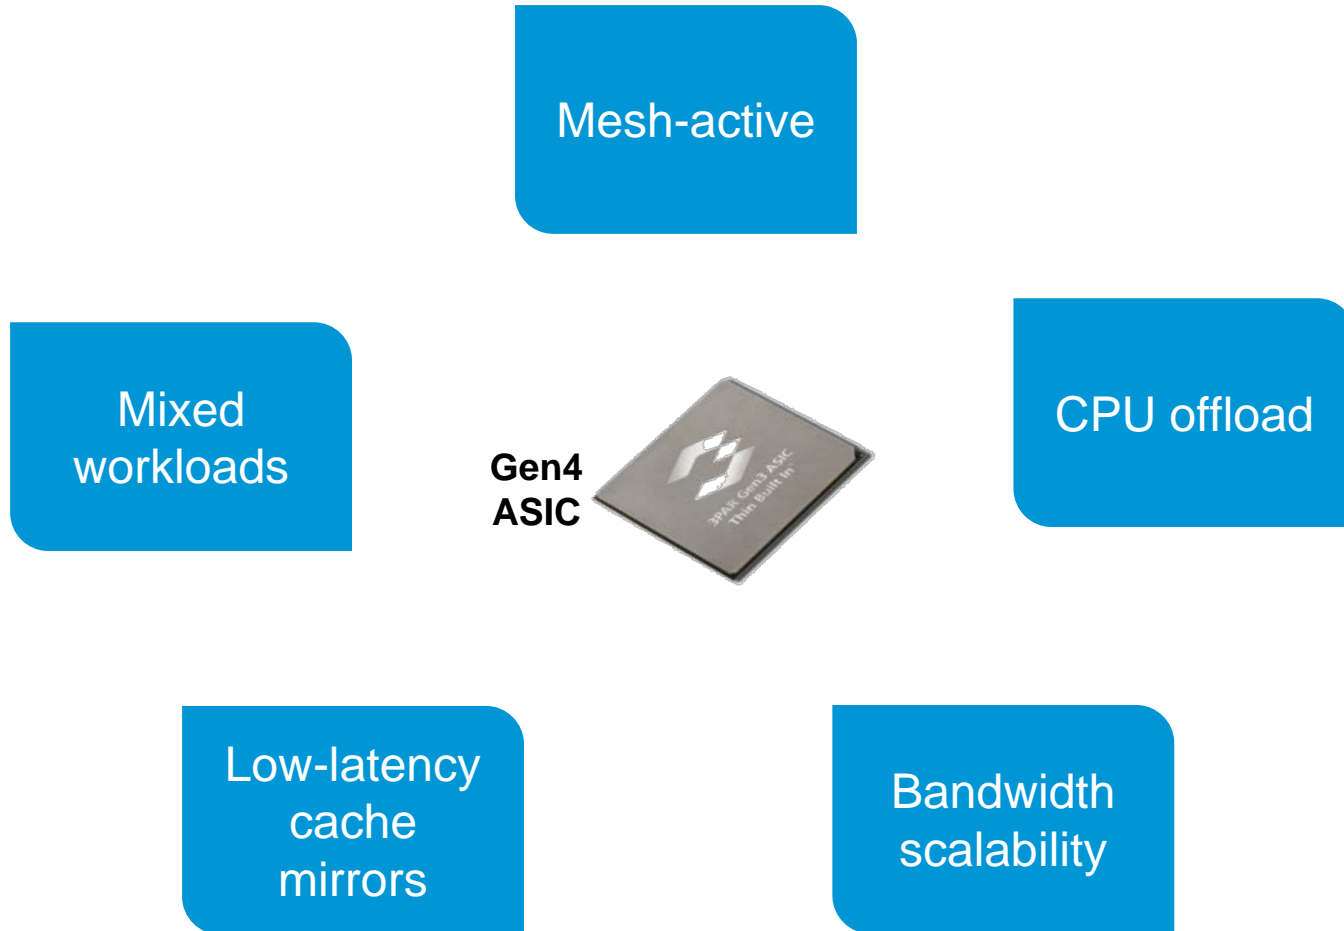
Multi-tenant I/O processing

Maintain performance during media failures

- Flash media has poor performance with large sequential writes
- In the event of a failure of a flash device, data is read from other flash devices and is written into spare chunklets
 - This mimics a sequential workload, that if not handled appropriately, can result into poor performance
- 3PAR architecture writes data to flash media in 32KB step sizes (instead of 128KB), thereby protecting the flash media performance

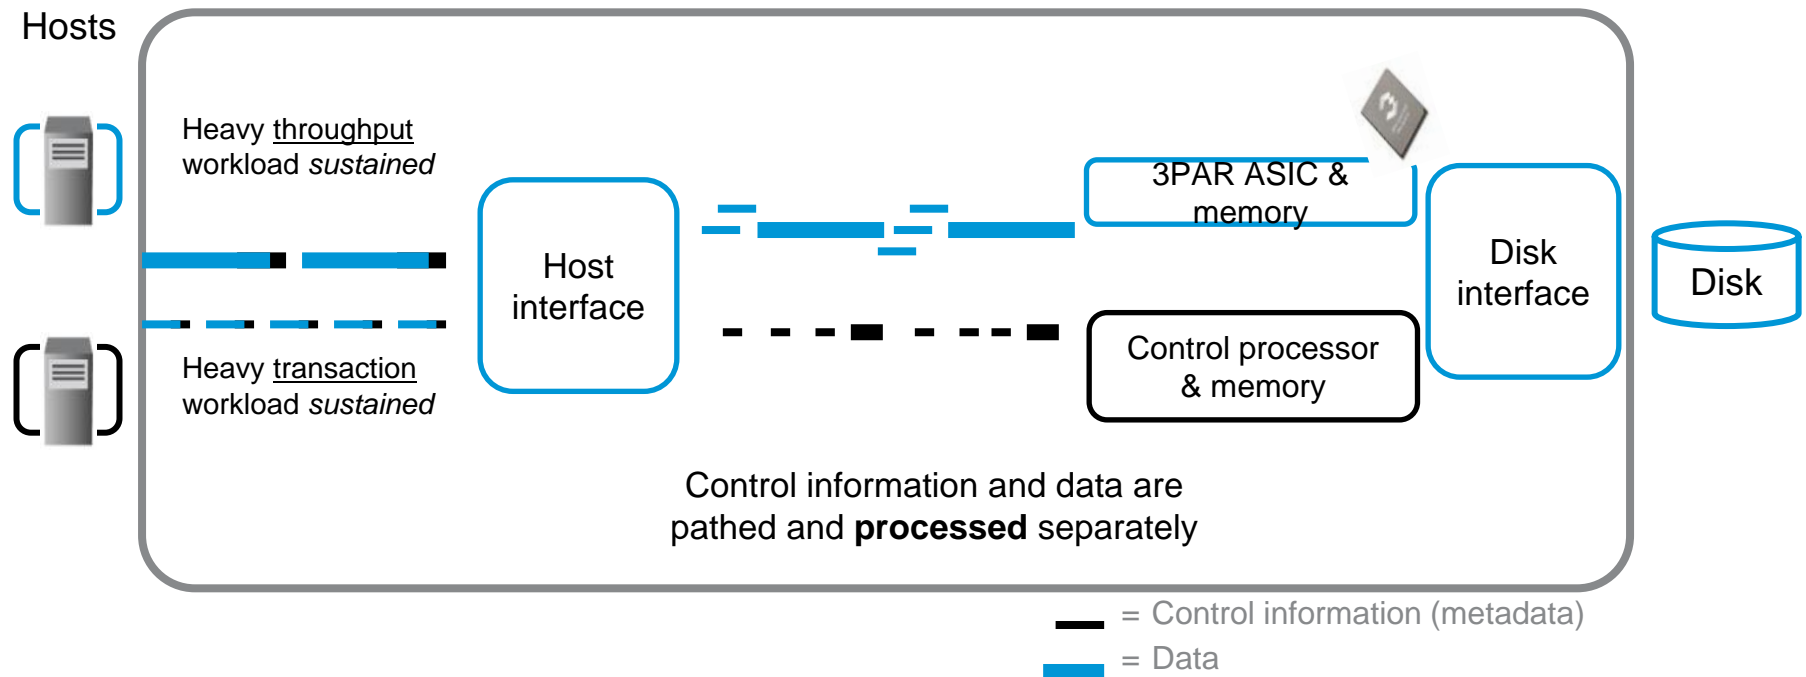


Offloading CPU for greater IOPS scalability



Specialized ASIC: For mixed workload support

I/O processing: 3PAR controller node



Benefits

- Specialized ASIC offloads data processing load from CPU and allows 3PAR to deliver much greater throughput capabilities
- Greater throughput becomes important while serving either a DSS/analytics workload or mixed workloads

Managing performance


Quality of service to enable service level guarantees

- Limit max performance with thresholds
 - At an application or tenant level
- Features
 - Max threshold for front-end IOPS and bandwidth
 - Configurable thresholds by:
 - Volume Set (proxy for application)
 - Hierarchical with over-provisioning in each hierarchy
- Enforcement in real-time
 - Seconds or sub-seconds

Benefit

- Service level certainty and predictability for app/tenants


Application 1



 Bandwidth control

 IOPS control

Application 2



3PAR Priority Optimization

Discrete performance threshold example

- 2 tenants sharing a single array, one tenant per virtual private array
- Tenant 1
 - Assigned 20,000 IOPS per SLA and “over-commits” at 150% of assigned IOPS
 - Array will cap total IOPS to 20,000 at any point in time
 - Max cap will not interfere nor impact tenant 2
- Tenant 2
 - Assigned 100,000 IOPS per SLA. Chooses not to “over-commit”
 - Array will allocate all assigned IOPS to application
- 3PAR architecture enables QoS & predictability of performance

Single array with 120,000 IOPS

Tenant 1 (20,000 IOPS)	Application 1 (10,000 IOPS)
	Application 2 (10,000 IOPS)
	Application 3 (10,000 IOPS)
Over-commits at 150% of assigned IOPS	
Tenant 2 (100,000 IOPS)	Application 4 (100,000 IOPS)

