Storage Acceleration, Driven by Autonomic Software

Phillip Clark / JJ Kane, Sr. Software Engineer


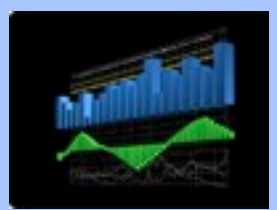
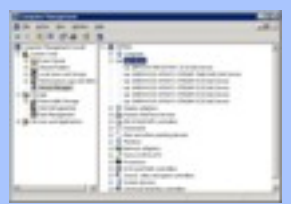
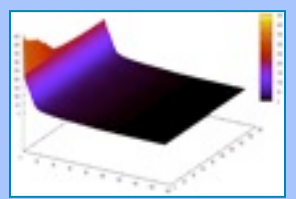

The Storage Dilemma

Increasing Performance Gap between Servers and Storage

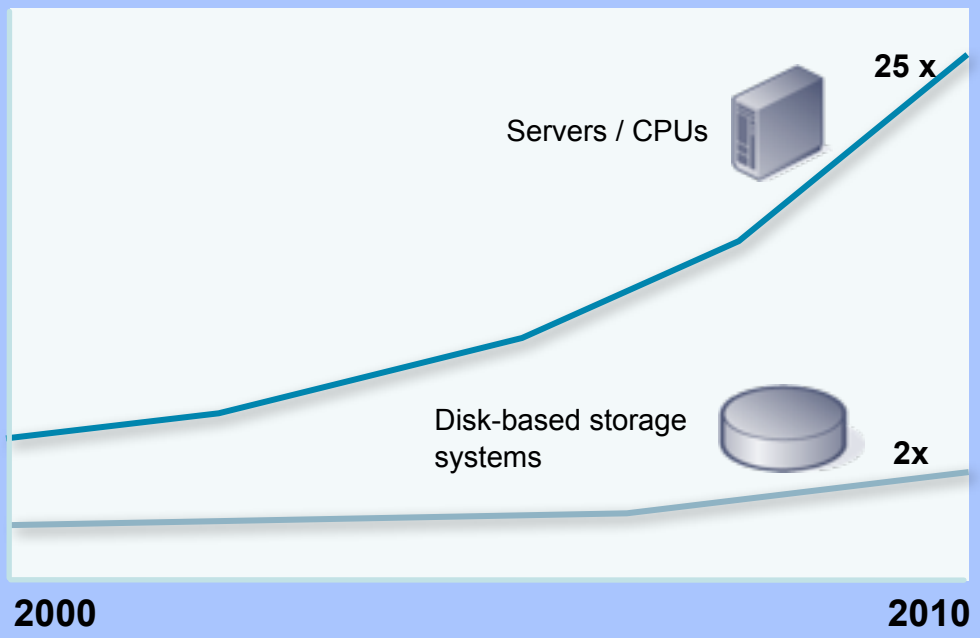

Increasing server performance

Traditional disk performance

The Storage Dilemma

Increasing Performance Gap between Servers and Storage

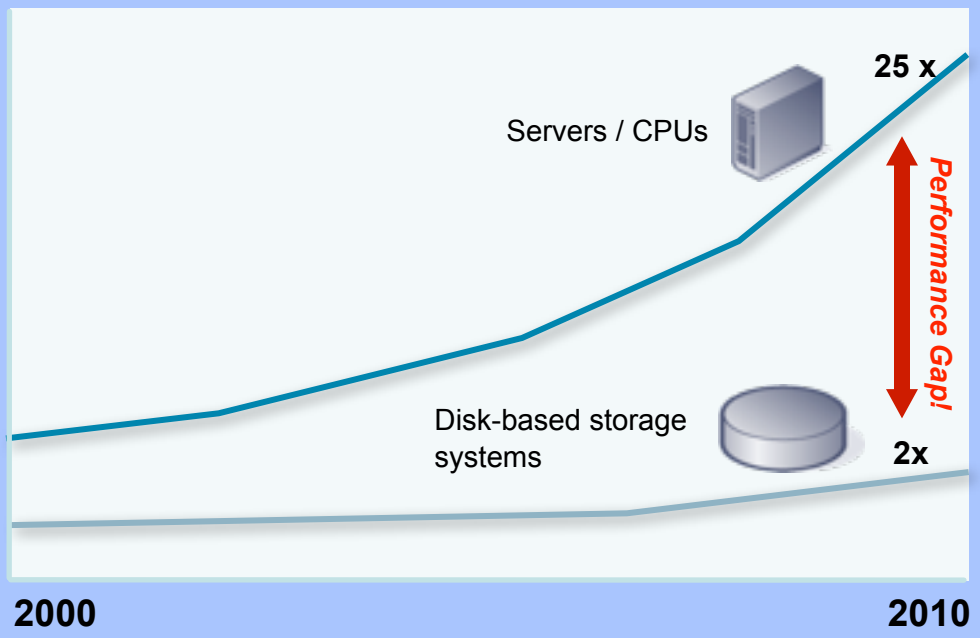

Increasing server performance

Traditional disk performance

The Storage Dilemma

Increasing Performance Gap between Servers and Storage

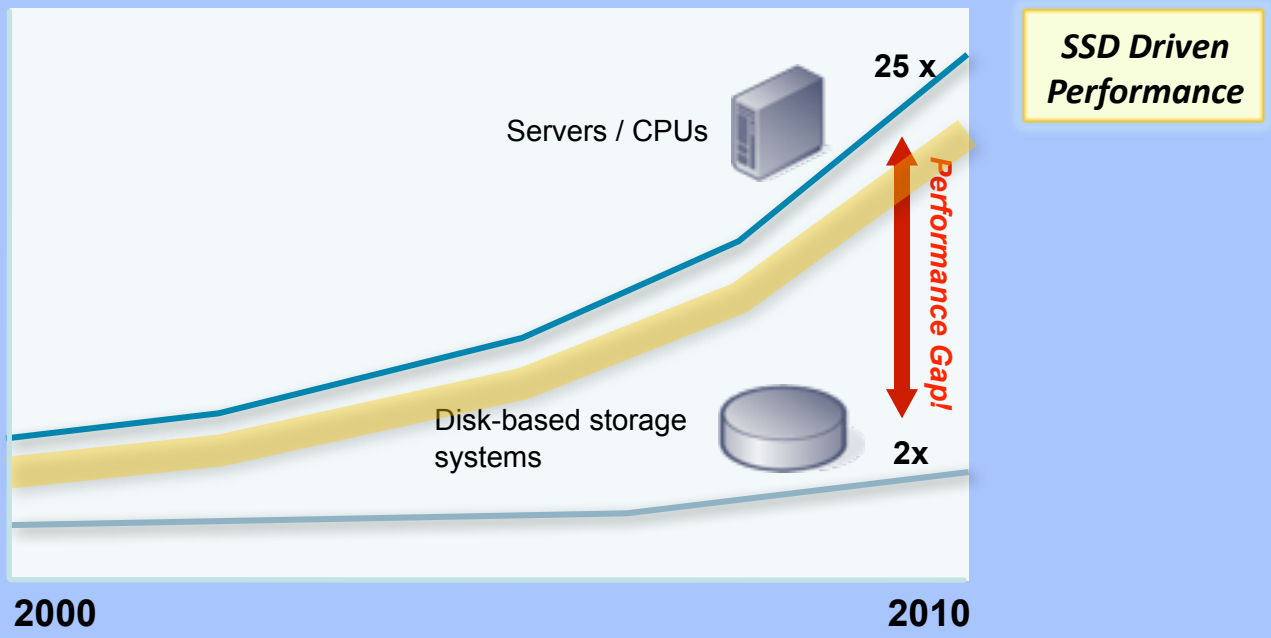

Tiered Storage Architecture

Current Limitations

Tiered Storage Architecture
Current Limitations

- Lack Knowledge to Size Storage Tiers
  - Can cause over-provisioning of costly SSDs
  - Unable to predict/show performance gains
  - No metrics to measure improvement

Tiered Storage Architecture
Current Limitations

- Lack Knowledge to Size Storage Tiers
  - Can cause over-provisioning of costly SSDs
  - Unable to predict/show performance gains
  - No metrics to measure improvement
- Need Scalable High Capacity/Density Arrays
  - Not bandwidth matched to scale capacity
  - Does not leverage HDD=Capacity, SSD=Access
  - Limited support for large scale, semi-random workloads
  - Cannot span from cacheable to pure SSD random access

Tiered Storage Architecture
Current Limitations

- Lack Knowledge to Size Storage Tiers
  - Can cause over-provisioning of costly SSDs
  - Unable to predict/show performance gains
  - No metrics to measure improvement
- Need Scalable High Capacity/Density Arrays
  - Not bandwidth matched to scale capacity
  - Does not leverage HDD=Capacity, SSD=Access
  - Limited support for large scale, semi-random workloads
  - Cannot span from cacheable to pure SSD random access
- Inefficient Management
  - Not adaptive to changing access patterns
  - Requires IT time and resources
  - Inability to scale effectively

New Data Center Storage
ApplicationSmart Self-Optimization

Storage Cache and HSM Limitations: Cache is limited in scale/scope, HSM is slowly activated

Traditional Storage Arrays

Too Small – GB, need TB
- Locality and set association limited
- Limited memory scaling

Time Linear Cache

Imprecise, Not Adaptive
- RAM cost/GB high
- Cache hit rate efficiency with LRU/MRU low

Bandwidth Mismatch w/ Cache & HDD

Missed Penalty, Too High
- Over-provisioning of cache
- Too much rack space for HDD
- HSM migration with SSD is slow and steals bandwidth

ApplicationSmart Provides Data Access Acceleration:
Manages cacheable data in real-time

Data Storage for Random Data Access

Identifies Cacheability by Application and Computes Speed-up

Profile Monitors Access Changes

If Data is Completely Random, Need Spindles or Random Access Storage

Terascale

SSD

SAID

Petascale

Tier 0 Replication (TME)

Autonomic Storage Tiering

Customer Benefit	Description
Sizes SSD requirements	• Analyzes and recommends amount of SSD *prior* to purchase • Only what is needed for applications, based on profile • No over buying or over provisioning
No added management	• Enables autonomic data tiering, no policies to set • Anticipates SSD needs based on data access patterns
Eliminates overhead	• Data is replicated but remains resident on HDDs • Avoids migration to and from HDD and SSD • No unnecessary IO, all tiering is opportunistic

Autonomic Storage Tiering

Hybrid VLUN spans SSDs and HDDs

Customer Benefit	Description
Sizes SSD requirements	• Analyzes and recommends amount of SSD *prior* to purchase • Only what is needed for applications, based on profile • No over-buying or over-provisioning
No added management	• Enables autonomic data tiering, no policies to set • Anticipates SSD needs based on data access patterns
Eliminates overhead	• Data is replicated but remains resident on HDDs • Avoids migration to and from HDD and SSD • No unnecessary IO, all tiering is opportunistic

The Bottom Line - Hybrid Storage Delivers the Flexibility to Solve Problems

Fundamental Storage Customer Requirements		
Performance	**Capacity Scalability**	**Cost**

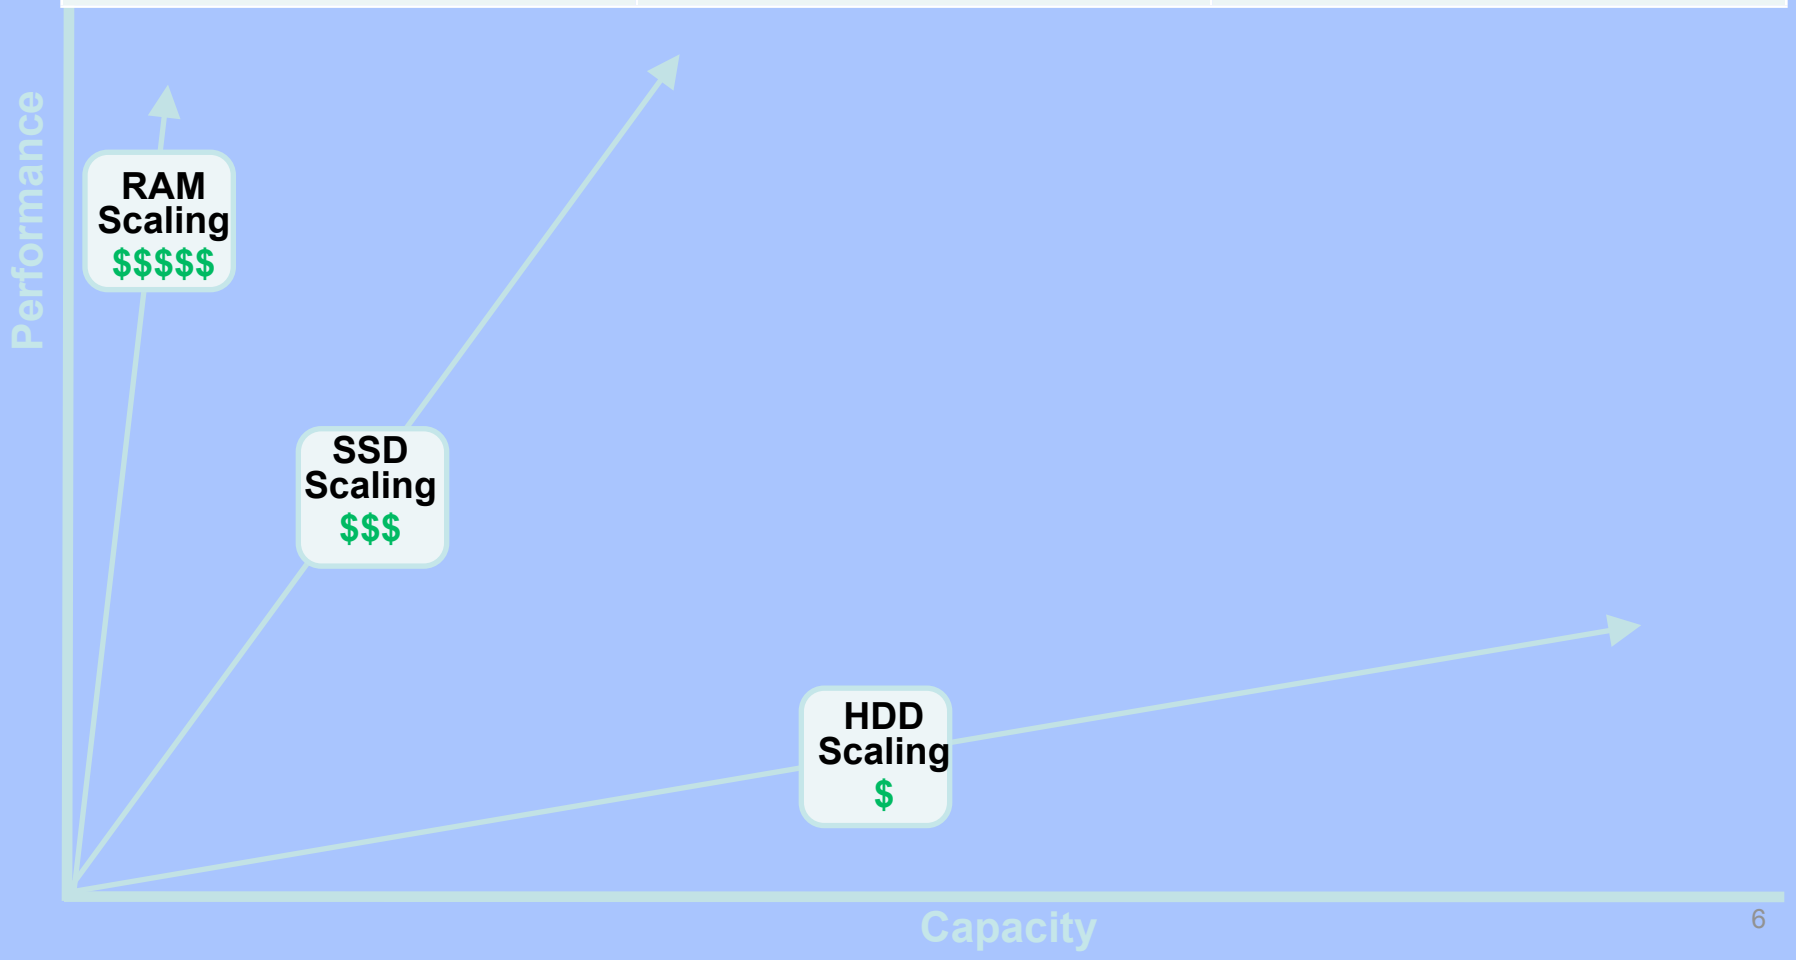

The Bottom Line - Hybrid Storage Delivers the Flexibility to Solve Problems

Fundamental Storage Customer Requirements		
Performance	**Capacity Scalability**	**Cost**

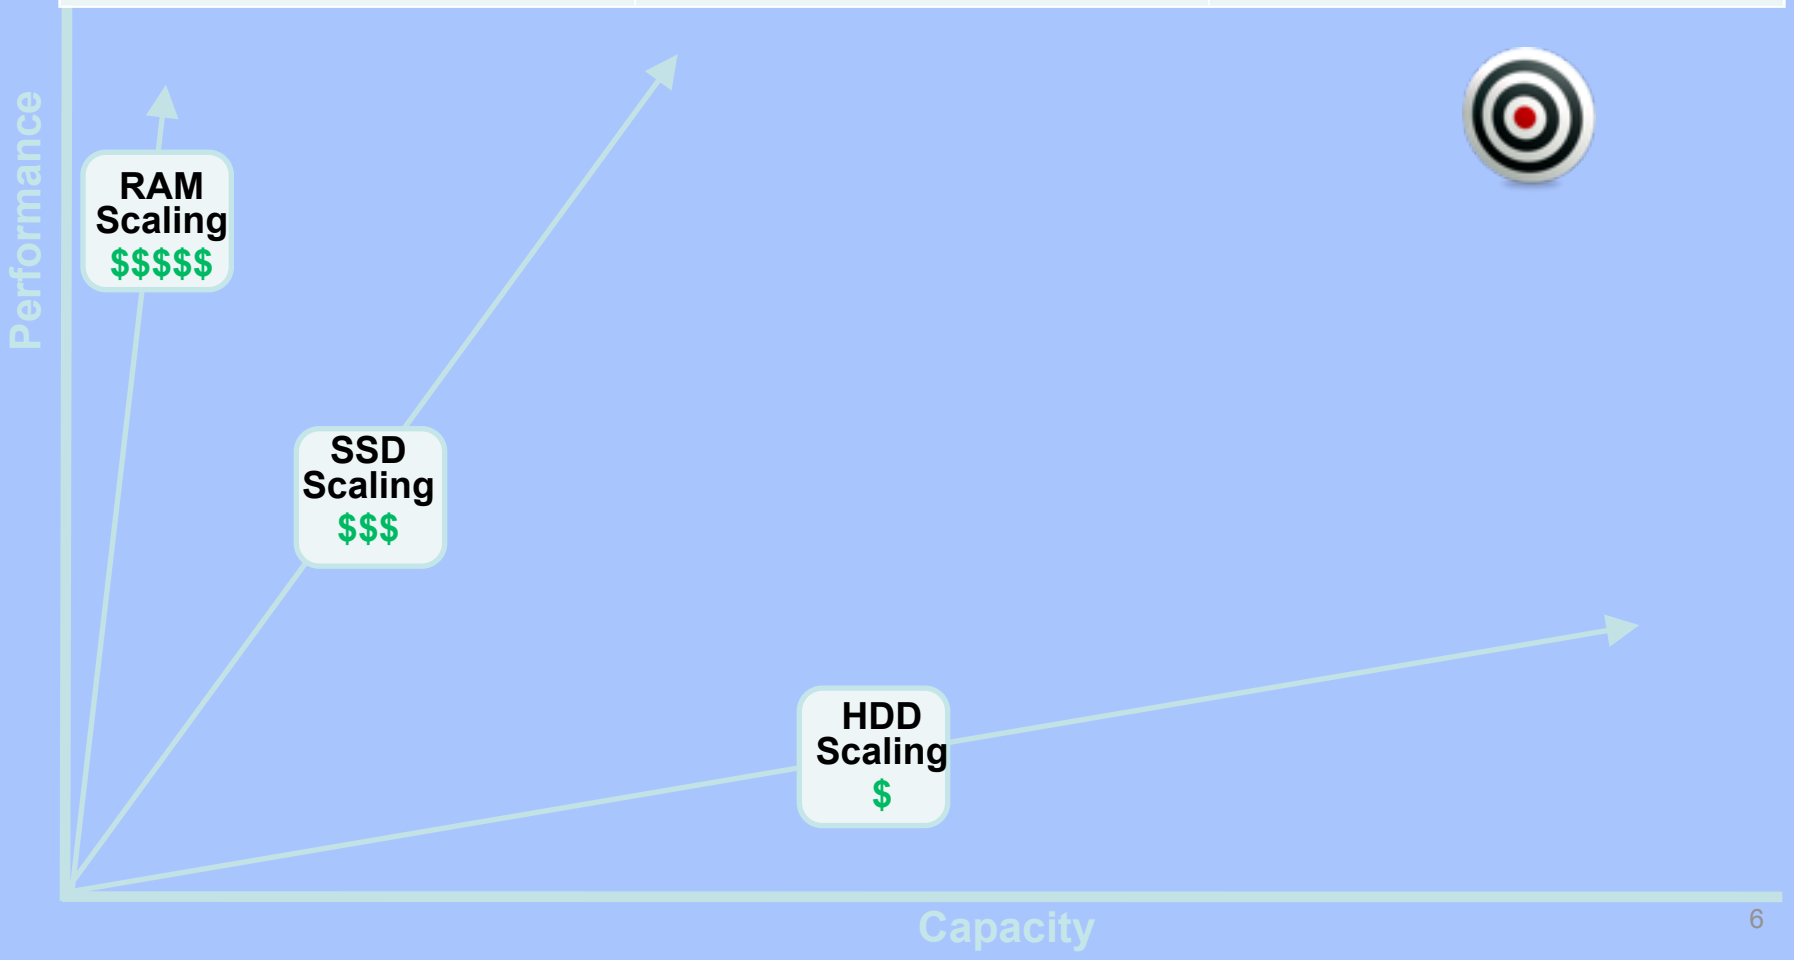

The Bottom Line - Hybrid Storage Delivers the Flexibility to Solve Problems

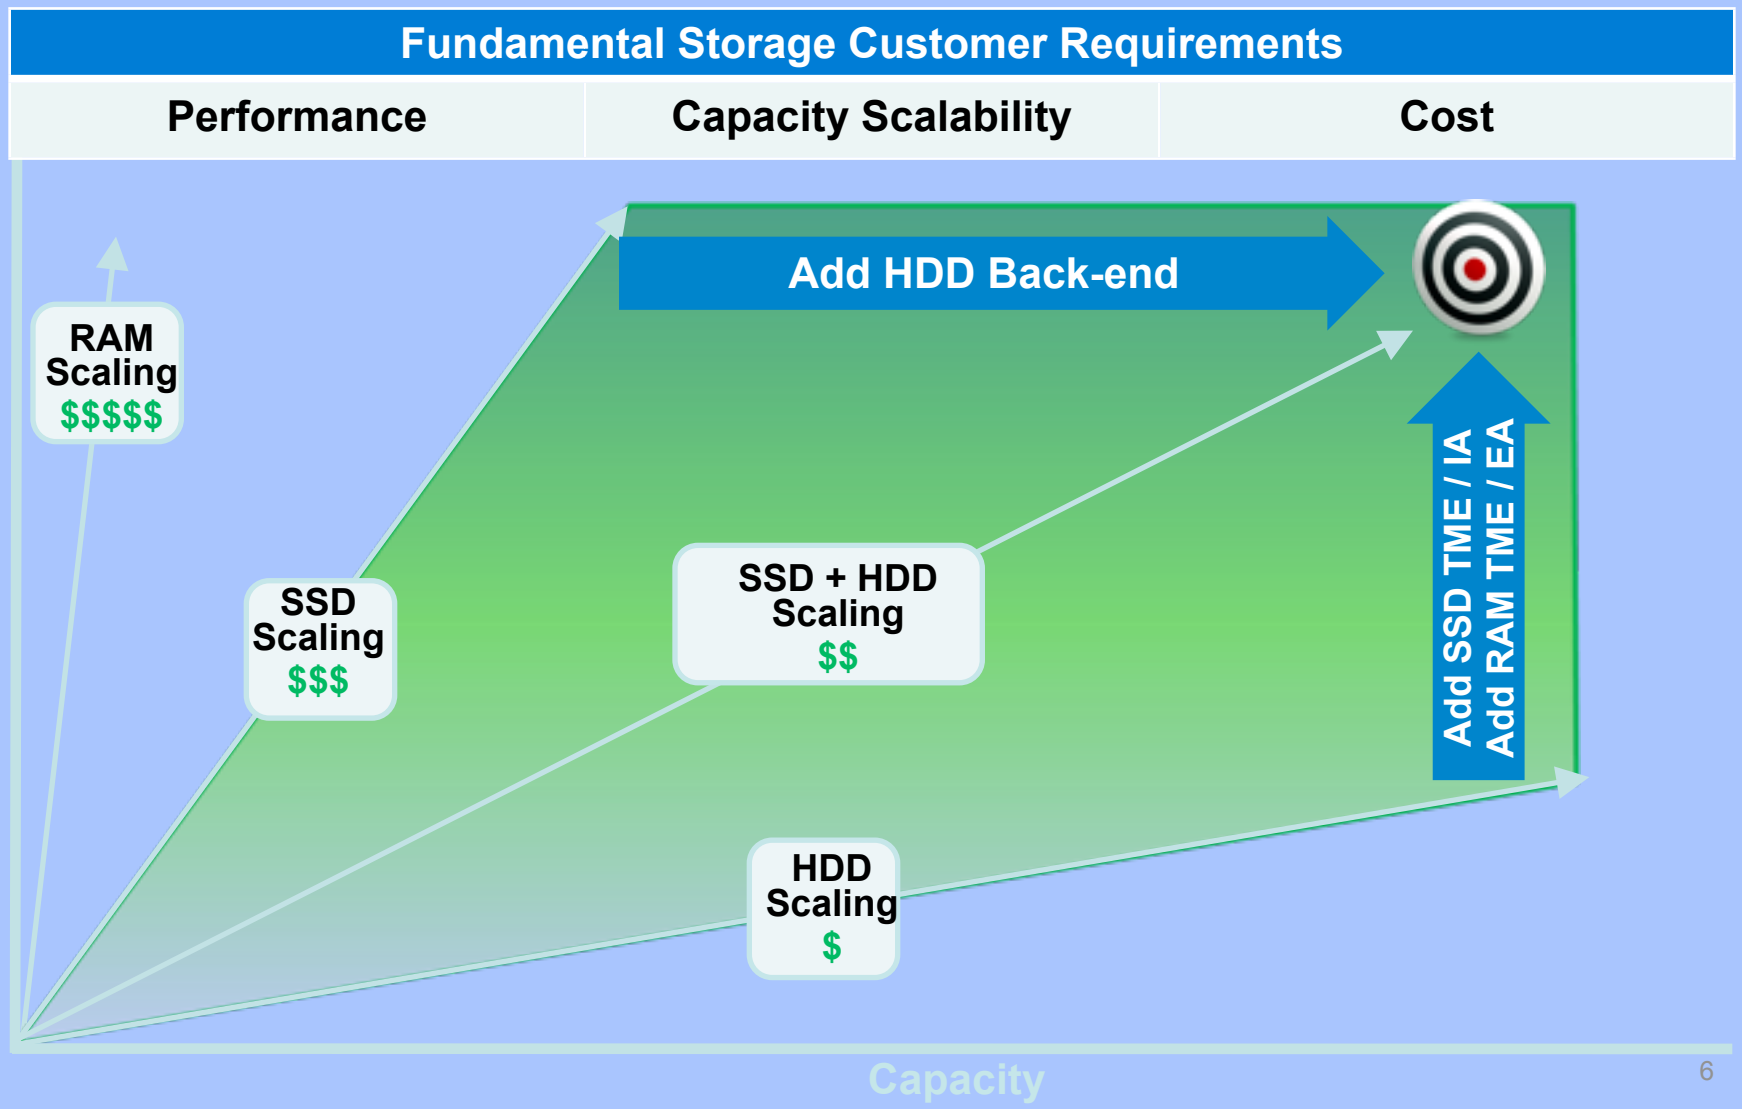

Performance Increase with SSDs

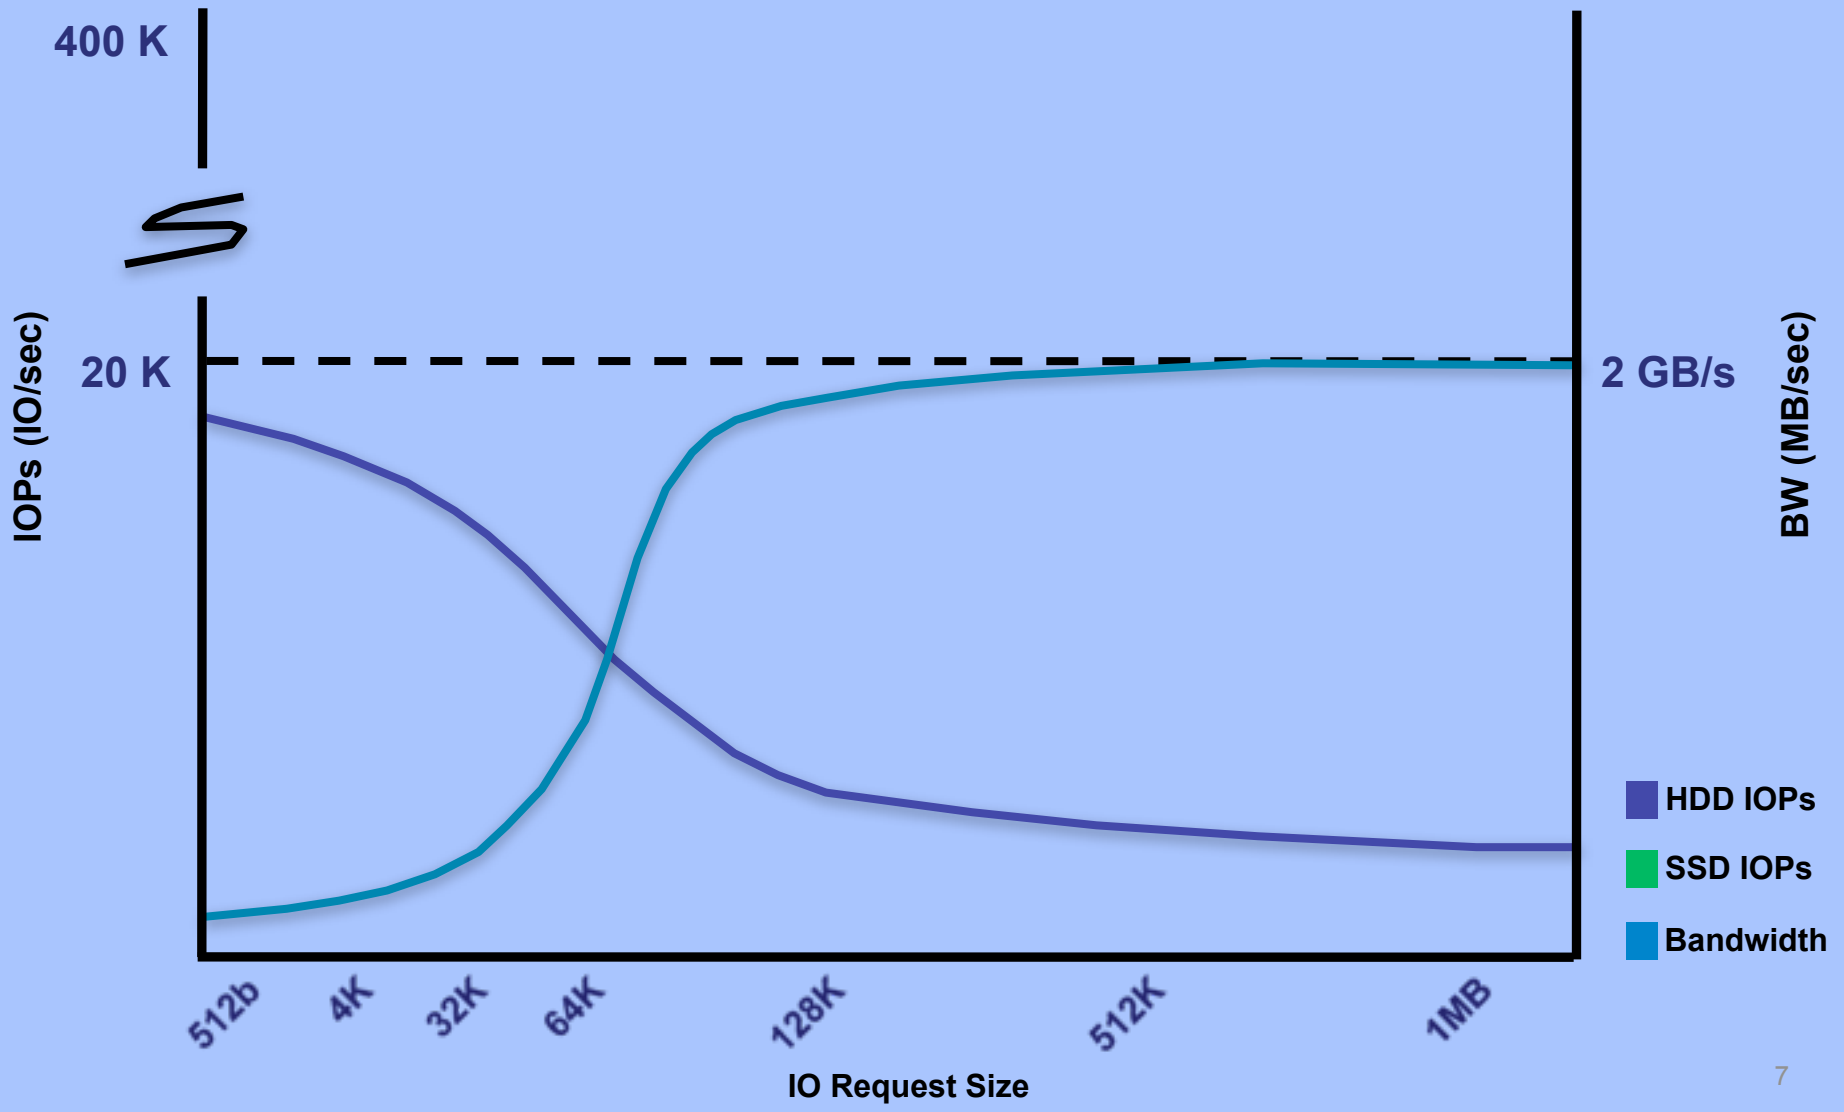

Performance Increase with SSDs

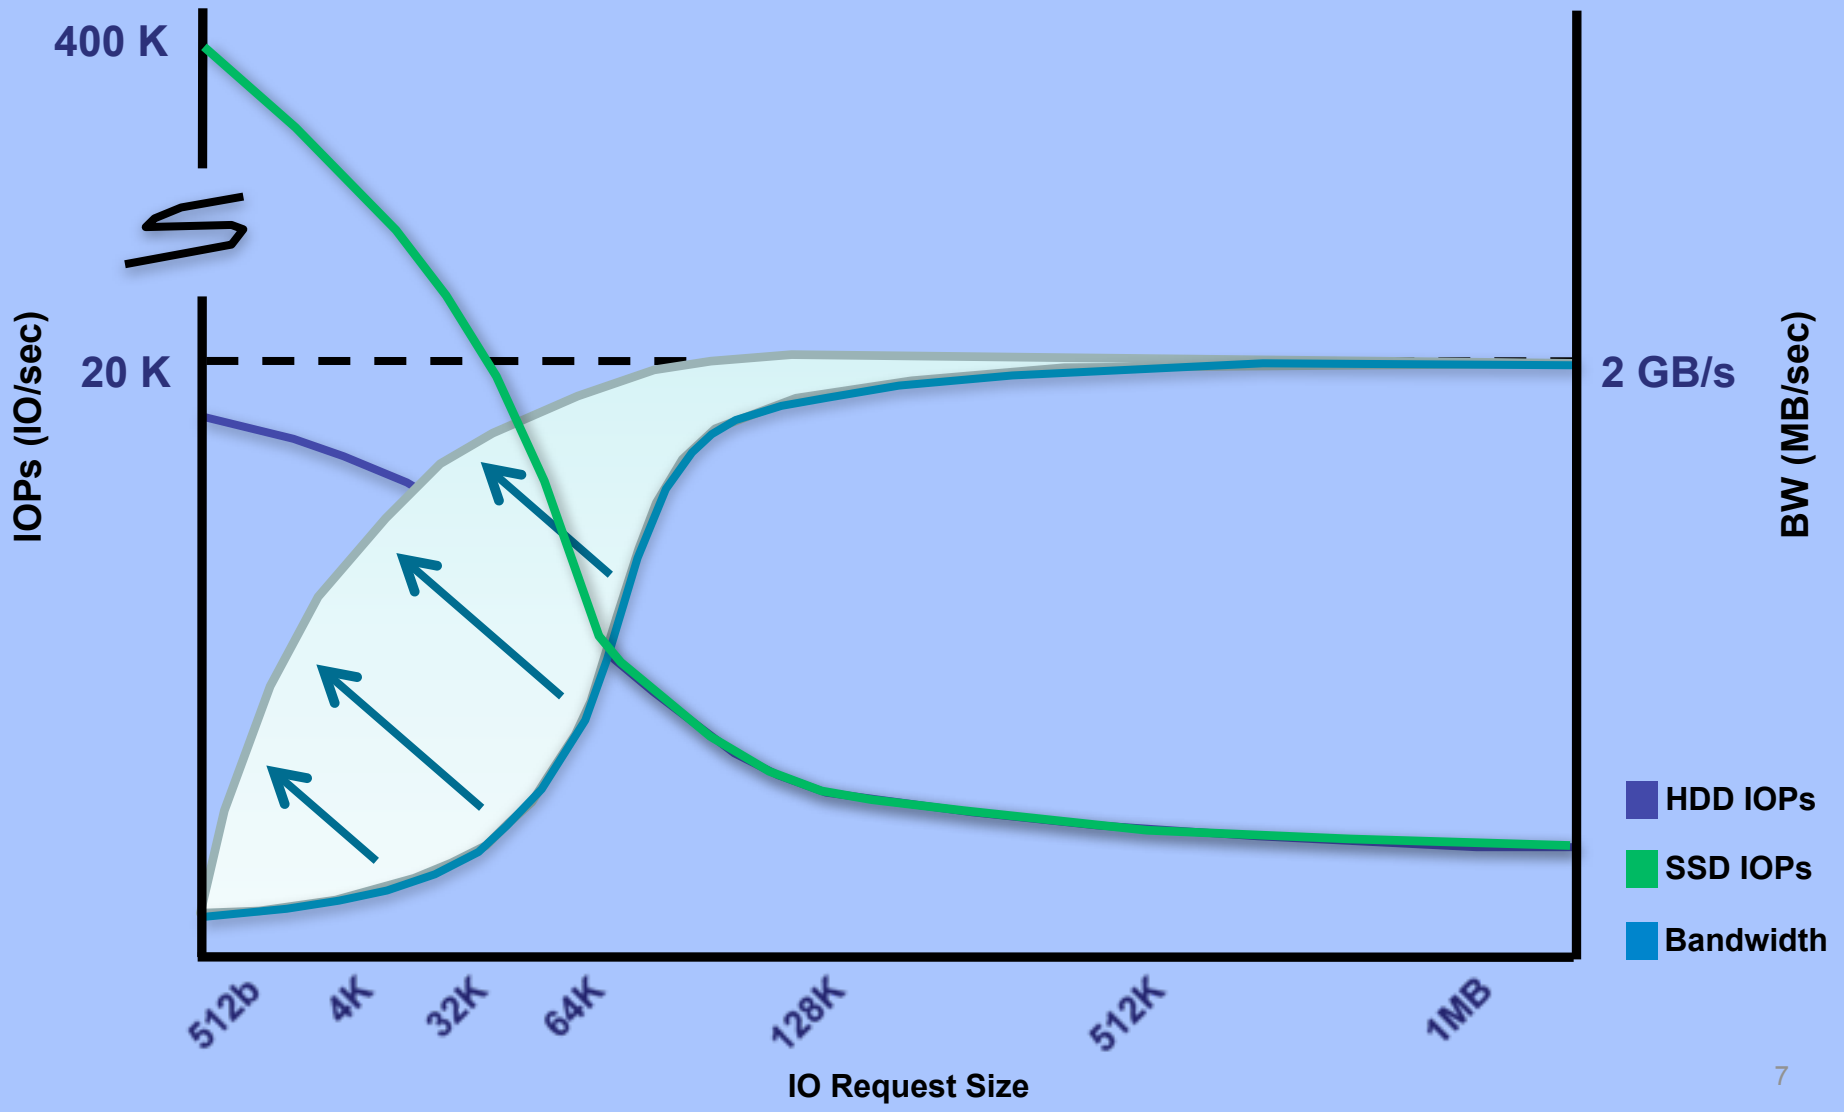

Multi-Tiered Management Software

ApplicationSmart™

Access Profiler	• Adaptive histogram, highly compressed, scales to PB • Drives TME to accelerate IO for high access content
TME (Tiered Management Engine)	• Dynamic block replication with access pattern changes • Optimal FBR (or plug-in heuristic) set replacement • Mapped to LUNs or pools of LUNs
Ingest Accelerator	• Tuned for RAID access (FIFO, back-end IO reforming) • Lower latency, higher throughput with log-structured FIFO
Egress Accelerator	• Detector for sequential/random initiator streams • Read-ahead cache with auto enable/disable
SLM (SSD LUN Manager)	• Full AVS VLUN creation and management • SSD storage pool, data lifetime protection options

Data Access Profiler

Provides real-time application storage access patterns

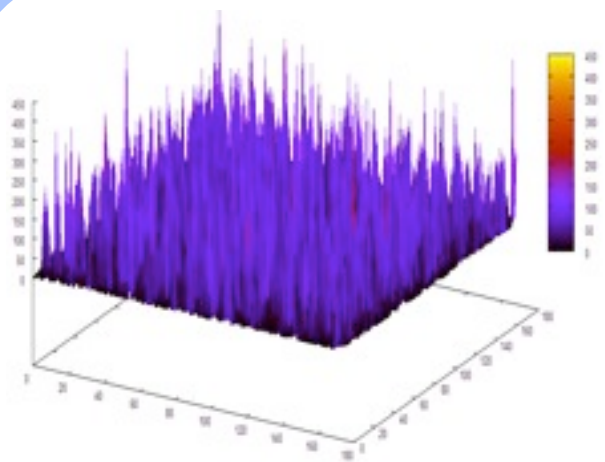

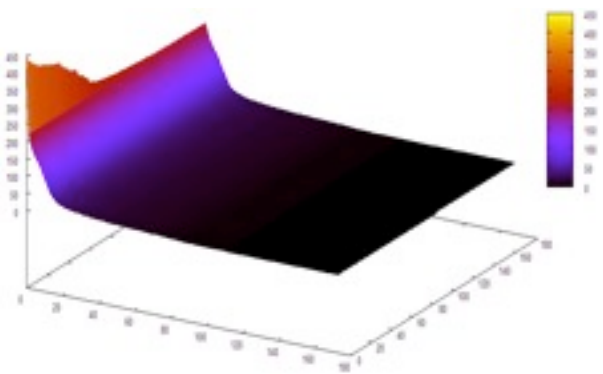

Histogram Analysis

- Identifies access hot-spots
- Notes when access changes are statistically significant
- Mapping integrates with virtualization engine

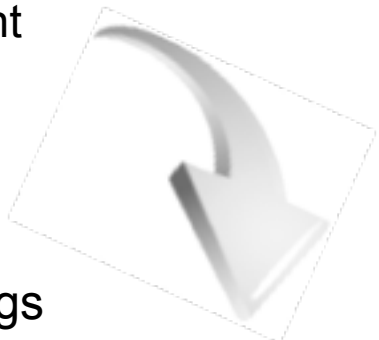

Histogram Groupings

- Drives TME IO acceleration
- Replicates blocks when statistically significant
- Provides continuous opportunistic updates
- Uses access visualization

Works w/ Wide Spectrum of Workloads

Sequential

Egress IO read-ahead

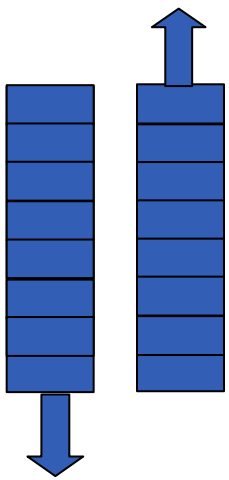

Ingest IO reforming

Fully Predictable
(Solid State FIFOs)

Hot-Spots

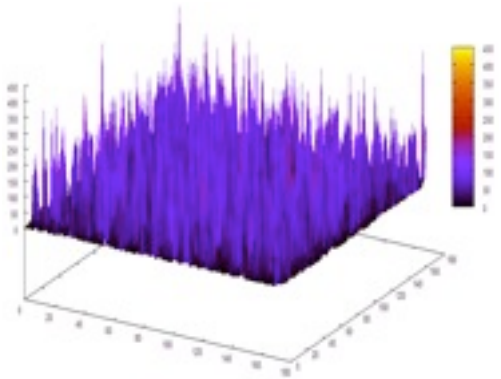

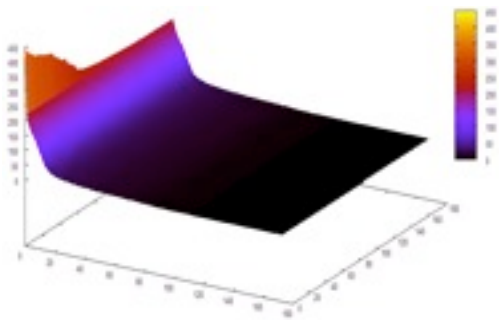

Semi-Predictable
(Scalable Hybrid Flash/Disk)

Random

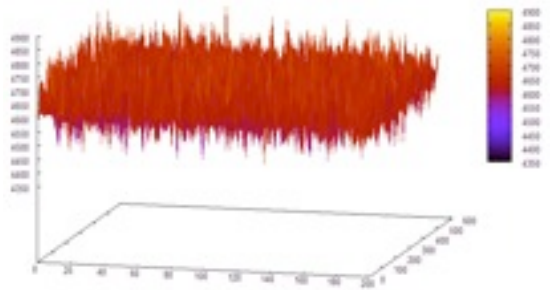

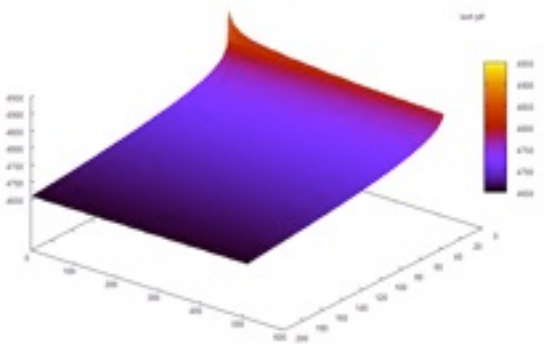

Non-Cacheable
(Solved by Spindle Density or
Random Access Storage Devices)

Tiered Management Engine (TME)

- Uses output from Access Profiler to drive TME
- Dynamic block replication
- As patterns change, new blocks are replicated
- Overwrites less active data

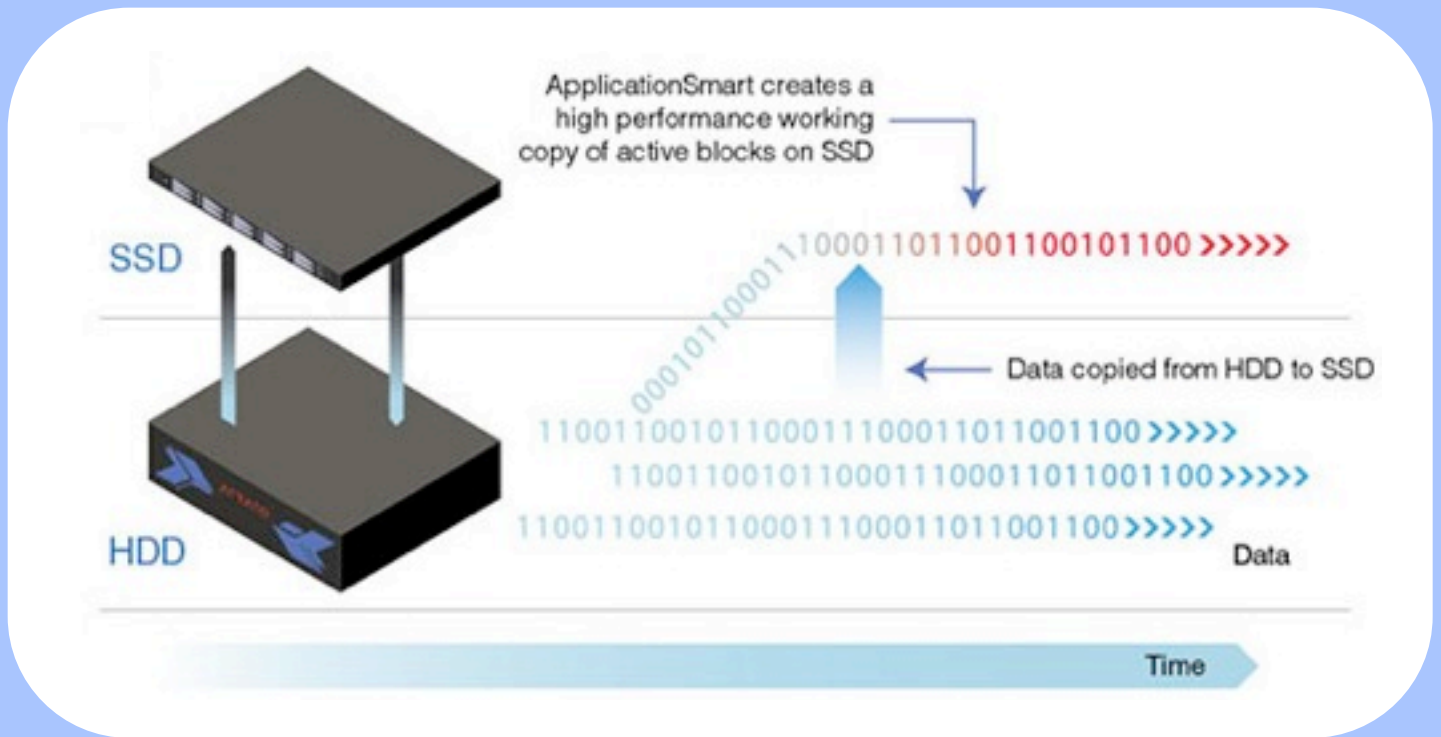

Ingest Acceleration

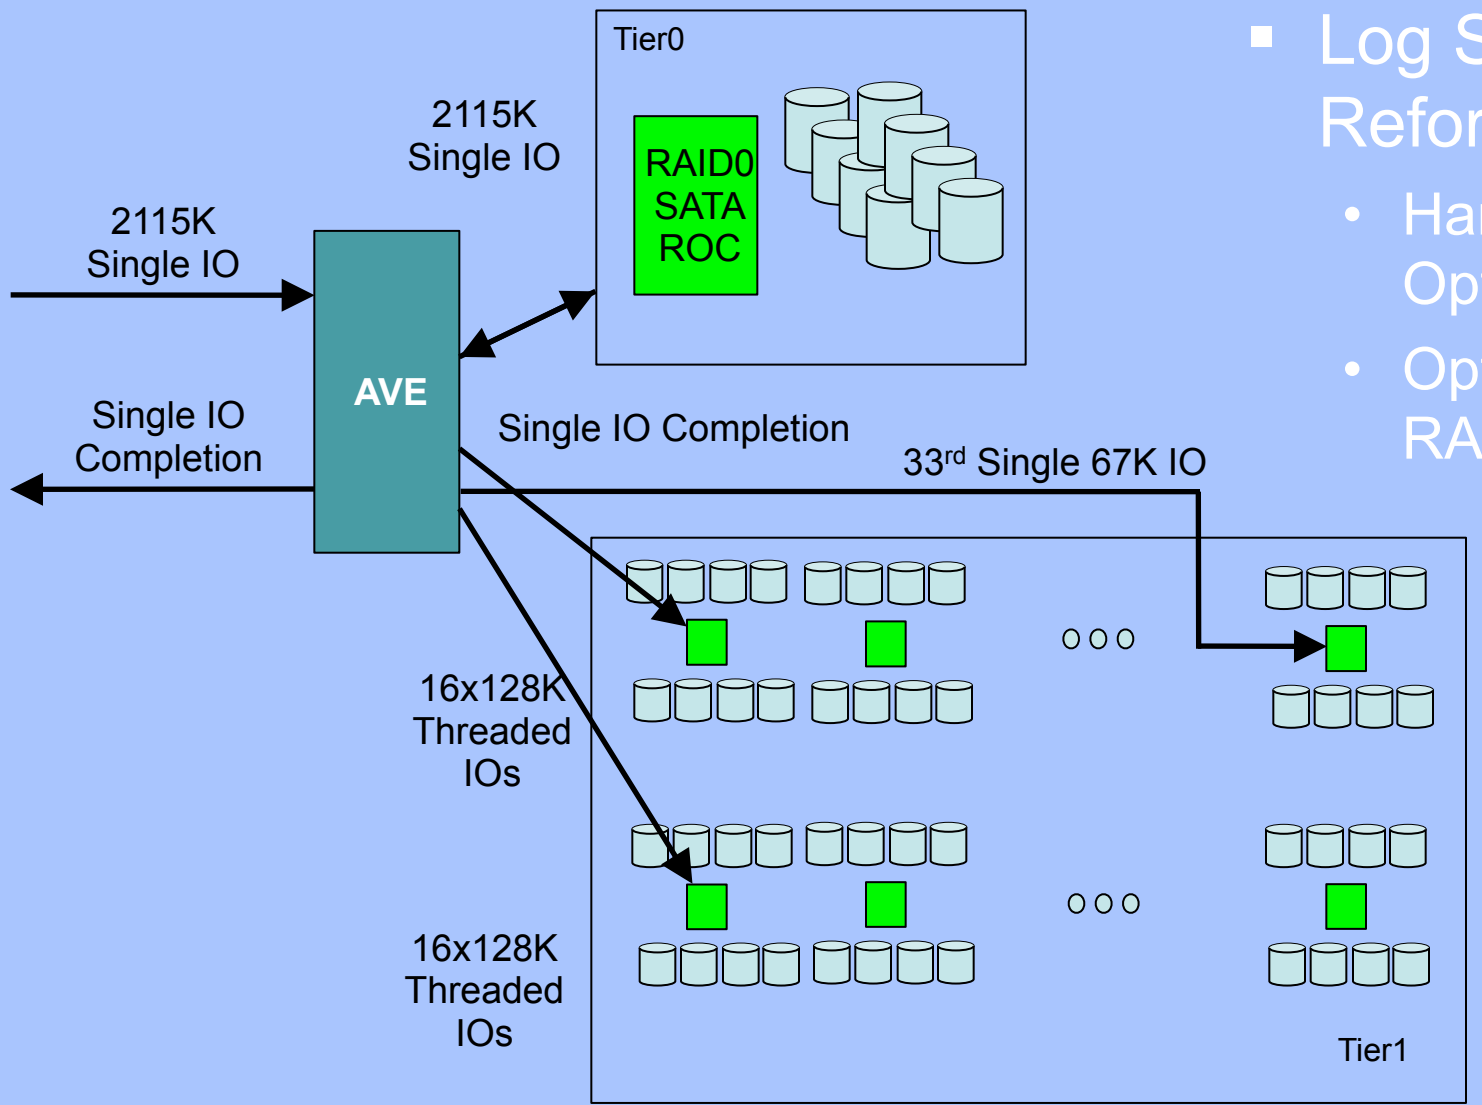

- Log Structured IO Reforming
  - Handles Sub-Optimal Writes
  - Optimizes for RAID Backend

Performance Tiering Checklist

Performance Optimization		**Key Benchmarks:** • *Integrates high velocity storage tiers (Tier 0, Tier 1)* • *Block level movement for increased granularity* • *Works across multiple application workloads*
Dynamic Management	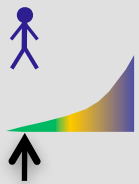	**Key Benchmarks:** • *Understands access patterns and changes* • *Moves data in real-time* • *Supports multiple VLUN configuration options*
Cost Efficiency		**Key Benchmarks:** • *Integrates efficient storage architectures (Tier, 0, Tier1)* • *Recommends SSDs only when needed* • *Fully autonomic, minimizes human intervention*

Flash Memory SUMMIT

Thank You!

Flash Memory SUMMIT

Questions?