

# Achieving Consistent Low Latency with All-Flash Arrays

Brian McKean

Netapp


E-Series Flash Architect

# Two worlds for Latency

	Dedicated Application Hardware	Shared Storage
Average Latency	<0.5 ms	<1 ms
Latency Sensitivity	Very Sensitive to spikes	Less sensitive to spikes
Feature use	Light	Heavy
Key Metric	\$/IOP	\$/GB Effective

# Latency

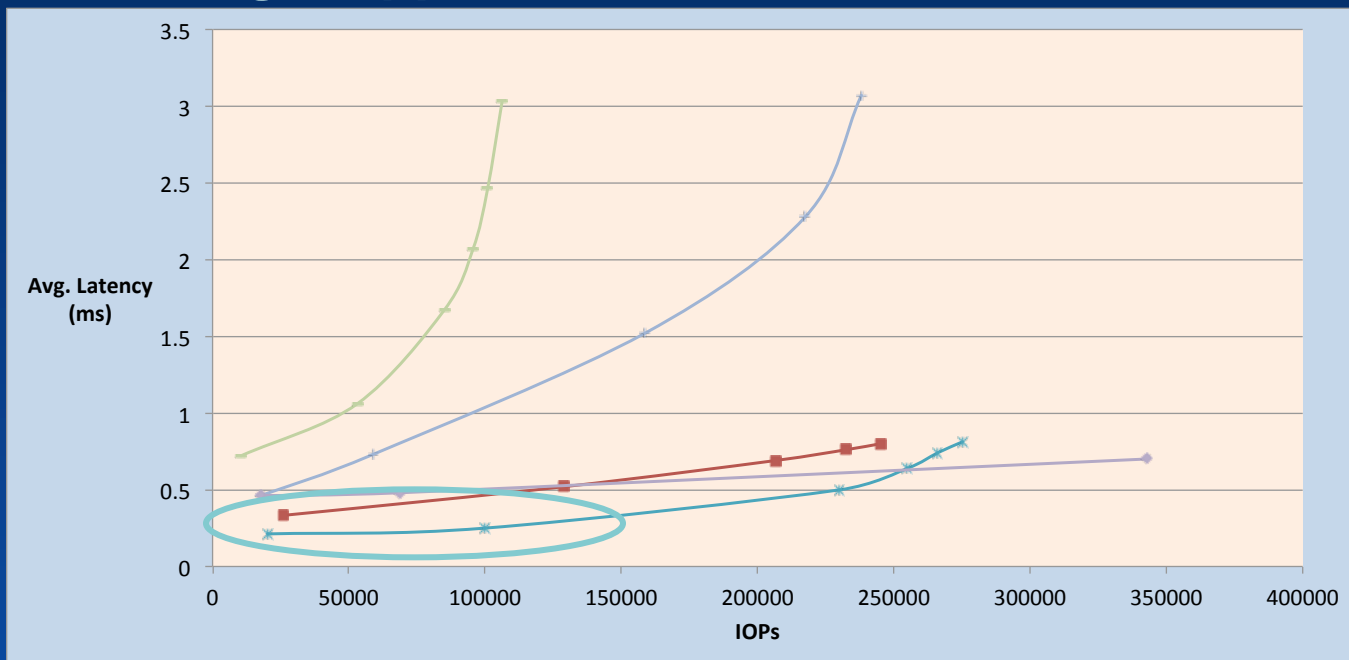
## Share Storage Feature Rich



IOP = 8K Block, 40/60 R/W, 80% Random

# Latency

## Single Application – Dedicated Hardware



IOP = 8K Block, 40/60 R/W, 80% Random

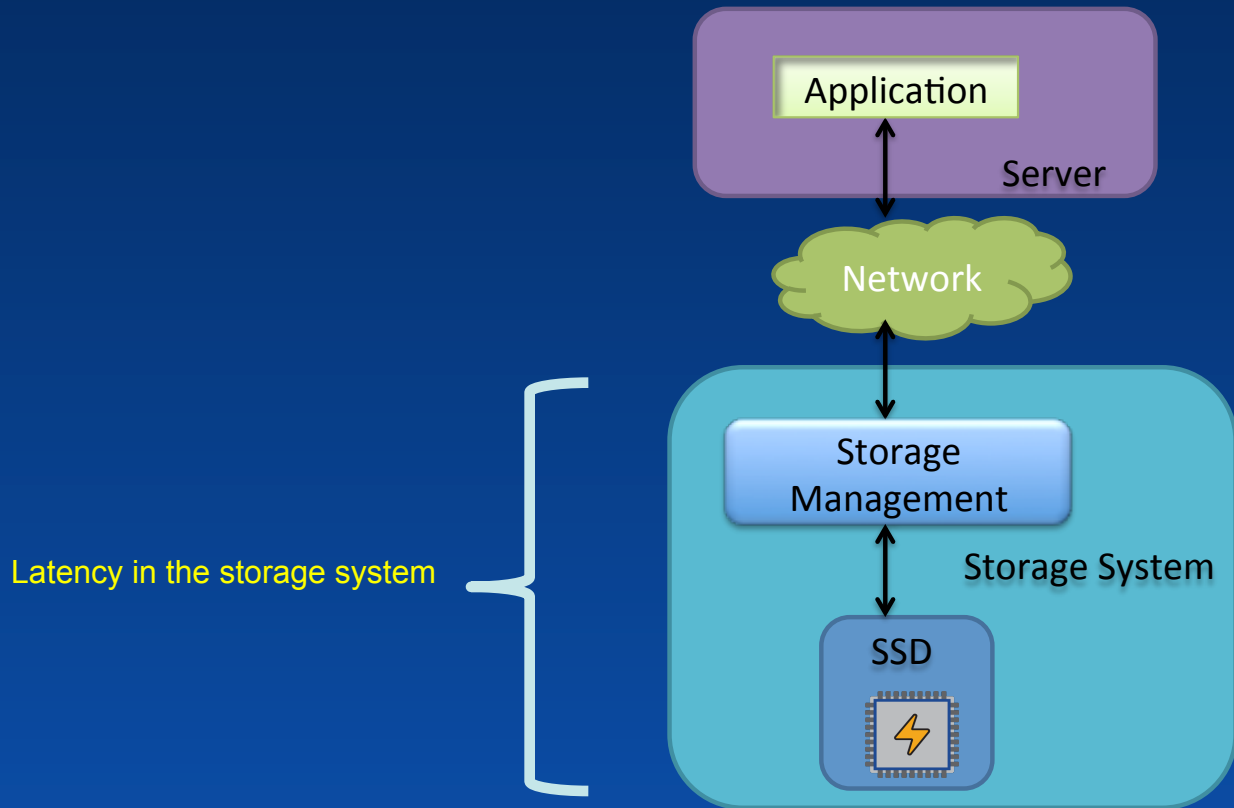
# How do we keep latency low in All Flash Arrays?



Function	Size	Time
Erase	256 KB	2 ms
Write	4 KB	200 usec
Read	4 KB	50 usec


We want to keep read and write to near flash speeds

# Storage Stack

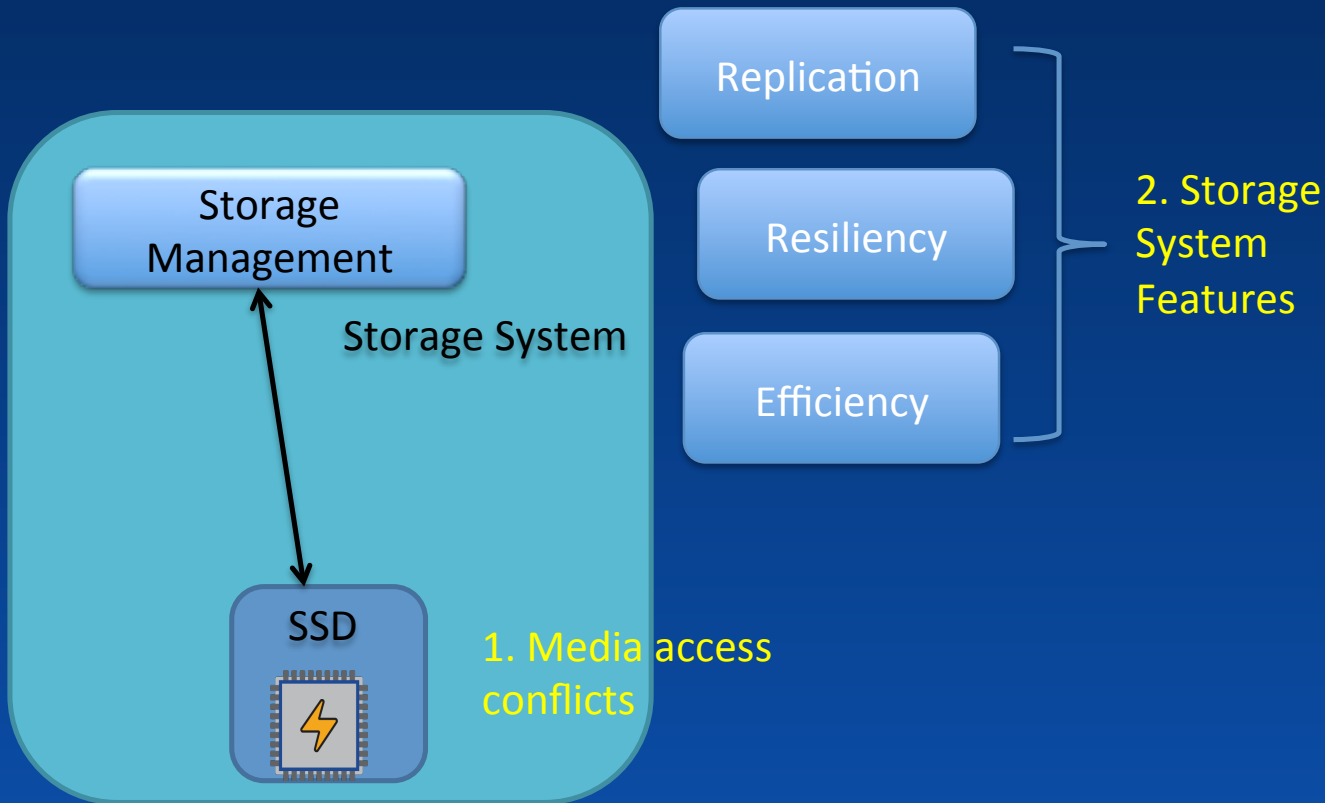


# What gets in the way of low latency?

1. Conflict with Flash media operations
  - Read/write waiting for for a 2 ms erase
2. Complexity of Storage Management Features
  - Layers of indirection & mapping
  - Multiple media accesses

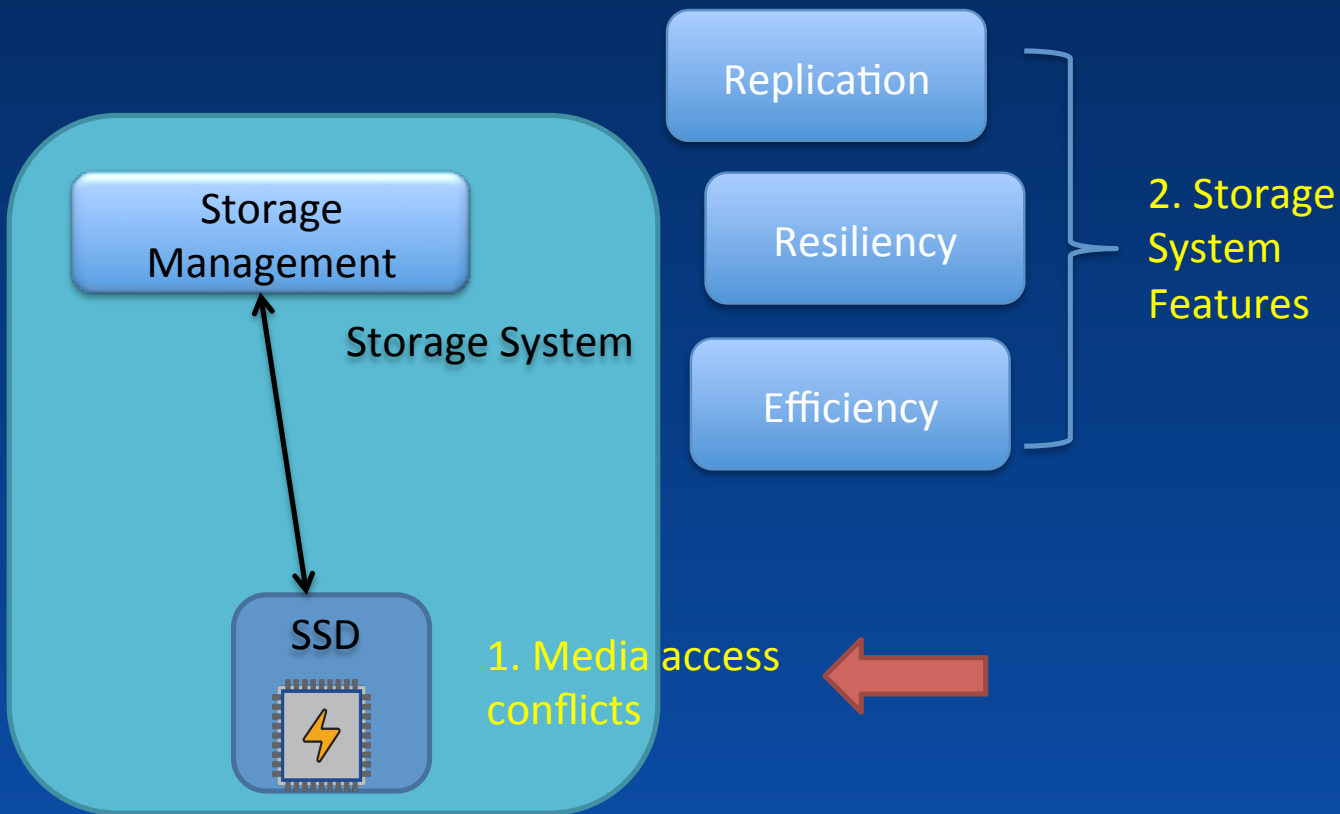


# What gets in the way of low latency?

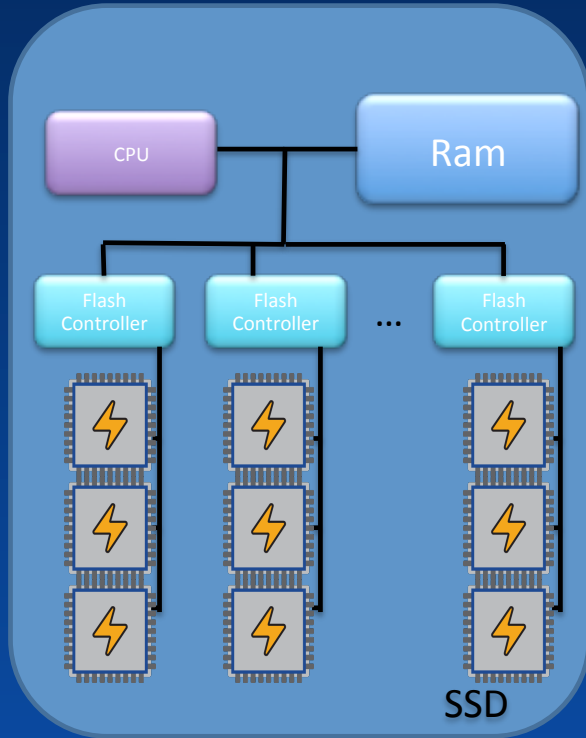




# What gets in the way of low latency?



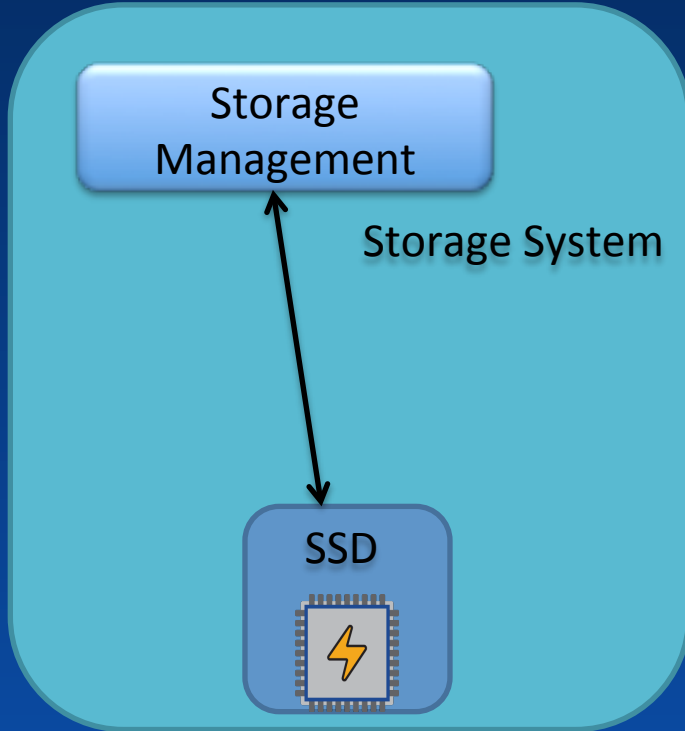
# Handling Media Access Conflicts



Within Device

1. Flash lanes and flash controllers
2. CPU / RAM
3. Over provisioning flash capacity

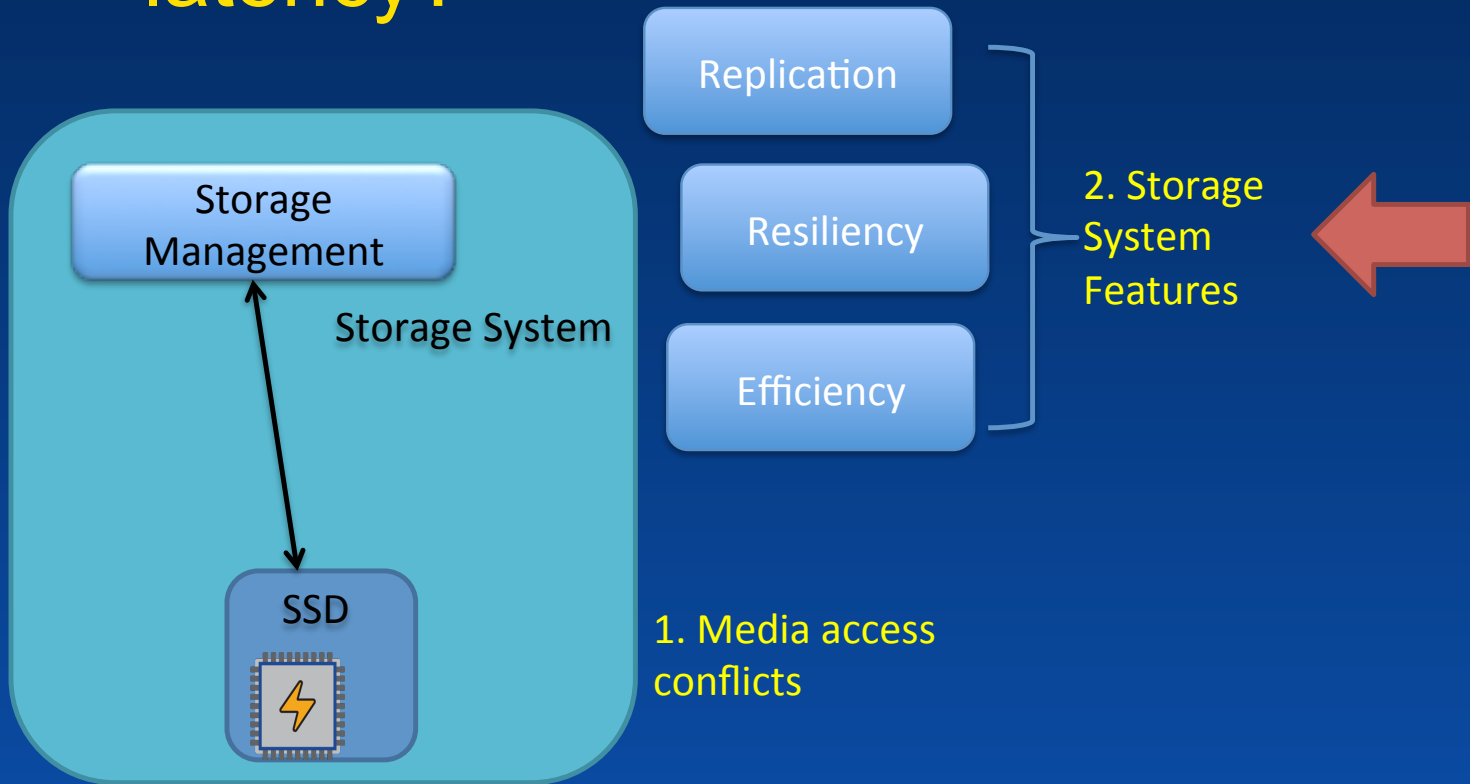
# Handling Media Access Conflicts



## SSD & Storage System Together

1. Sequential write patterns
2. Access other devices
3. SCSI Stream commands

# What gets in the way of low latency?



# Handling Latency from Storage System features

- Replication
  - Snapshots
  - Clones
- Latency mitigation
  - Built into architecture
  - Hardware Assist
    - More cores
    - More DRAM
  - Higher performing SSD



# Handling Latency from Storage System features

- Resiliency


- Erasure codes / RAID
- Replication



- Latency mitigation
  - Hardware Assist
    - More cores
    - Specialized HW
  - NV Mirror
  - Higher performing SSD
  - RAID Level choice
  - Full stripe writes

# Handling Latency from Storage System features

- Efficiency
  - Thin provisioning
  - Compression
  - De-duplication
- Latency mitigation
  - Built into architecture
  - Hardware Assist
    - More cores
    - More RAM
    - Specialized HW
  - Higher performance SSD
  - Fewer IO to SSD from efficiency



# NetApp EF-Series Latency Optimization

## Storage System Features

### Resiliency

- SANtricity OS
- Custom drivers
- Streamlined IO Path
- Direct mapped
- Mirrored NVRAM
- Hardware assist for erasure coding



## Media Access Conflicts

### SSD Selection & Qualification

- Over provisioning
- Flash lanes / device
- CPU/RAM

Top 5 SPC-1 Price-Performance  
#1 SPC-1 LRT



# All Flash FAS Latency Optimization


Storage System  
Features

Replication

Resiliency

Efficiency

- Data ONTAP OS
- Features Built into the Architecture
- Cores & DRAM, Mirrored NVRAM
- Flash Optimized Read IO Path
- Efficient Sequential Write Streams



Media Access Conflicts

- SSD Selection & Qualification
  - Less Over provisioning
  - Flash lanes / device
  - CPU/RAM

Top 5 SPC-1 Performance



# Summary

- Multiple approaches to developing an architecture for low latency
- No single feature is key for low latency
- Look for features and performance to fit your application needs

■ NetApp Booth #511