



# What Flash Vendors Won't Openly Disclose


Patrick Allaire  
Hitachi Data Systems

# The Battle For Attention



## Flash Focus: Get The Headline

- Flash delivers 1,000,000+ IOPS
- Flash 25, 50, 100, 4000x faster vs HDD
- 100K, 200K+ IOPS per device
- >95% less power per TB
- 40K, 50K, 60K, 100K P/E cycle
- Lower cost than traditional storage



***Such claims aren't exactly FALSE...***

***...but also aren't very HELPFUL.***




# Performance vs. Feature Value

Key Cost Factors For Storage

## Sources of Savings:

- Waste Reduction
- Reclamation
- Data Migration
- Outage time reduce
- Mgmt Labor Effort
- Maintenance Fees
- Environmental
- Misc Ops Efficiency

**HARDWARE COST IS ~20% OF TCO**





# Claim: "Just Add Flash"

*Full Disclosure: A Chain Is Only As Strong As The Weakest Link*

## FLASH ECOSYSTEM

### Commodity

- 1,2,3 bits per cell

### Specialty

- Enhanced Reliability
- Enhanced Endurance

**NAND  
Technology**

**Application  
Workload Demands**

**System  
Hardware and  
Software**

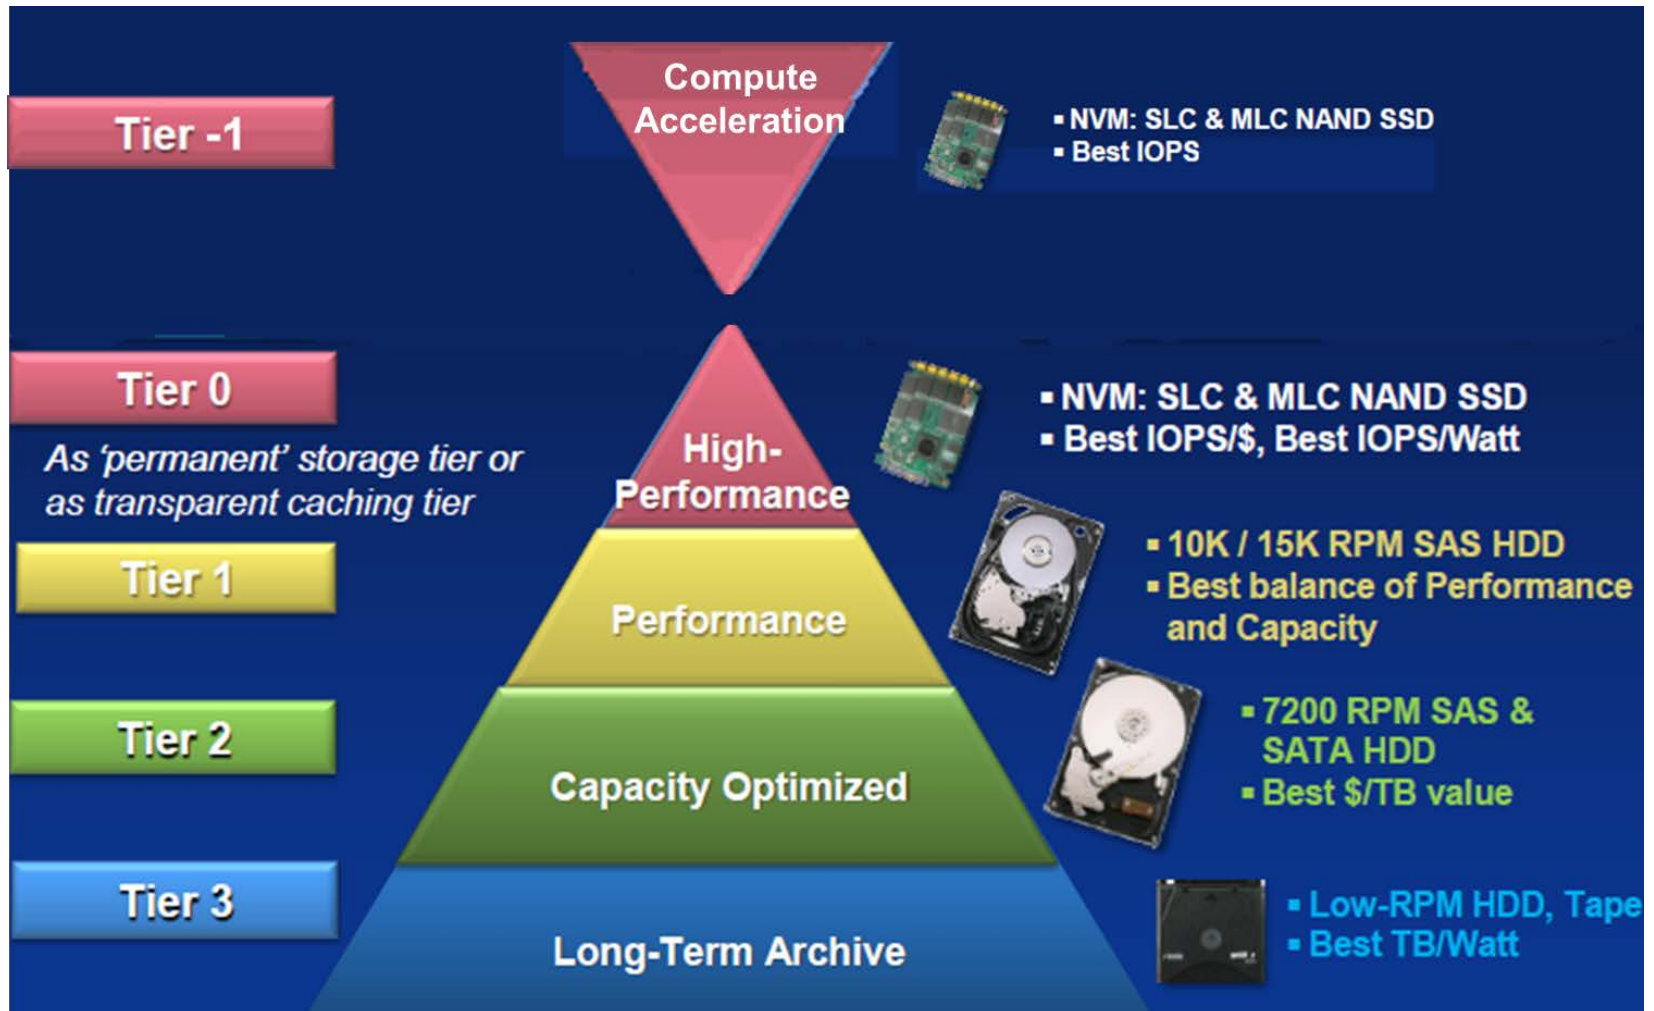
**Embedded  
Controller and  
Firmware**

- Thin provisioning
- Copy services
- Data Mobility
- Auto Tiering
- Caching
- Compression
- De-duplication
- Data protection
- Encryption

- Multicore
- Wear leveling
- Data refresh
- Performance manager
- Error correction
- Read retry
- Endurance manager
- Data protection

# Claim: "All You Need Is Flash"

*Full Disclosure: One Size Does NOT Fit All*



# Getting To Full Disclosure On Performance Claims

- Peak vs sustained performance
  - >40% difference between 4K peak vs. sustained reads measurement
- Response Time
  - Where was the measurement taken: host, cache or device level?
- Read I/O vs. Write I/O
  - Specs based on 100% read I/O; write I/O can be as low as 10% of reads; size requirements at proper read/write %
- I/O block size
  - Vendors often promote 4K block performance; 8K throughput is often 40-60% less; select size to your app's spec



# Getting To Full Disclosure On Performance Claims

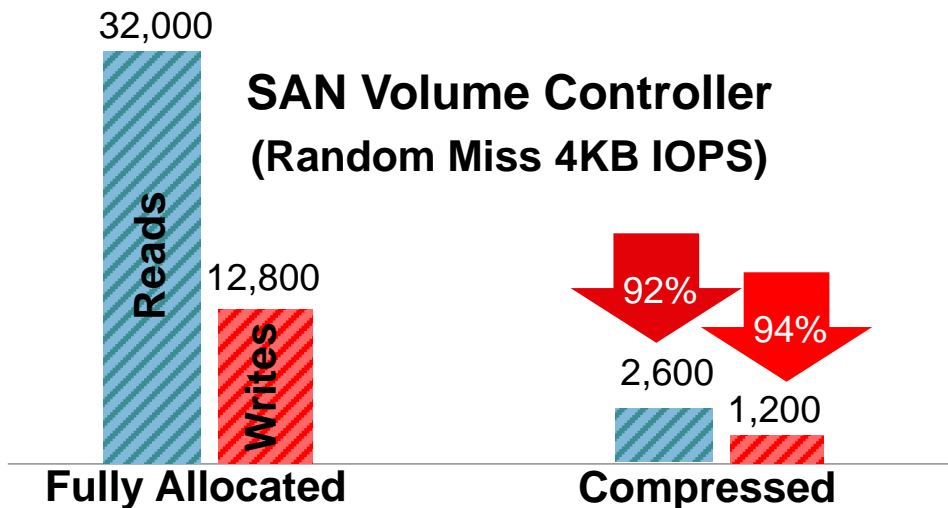
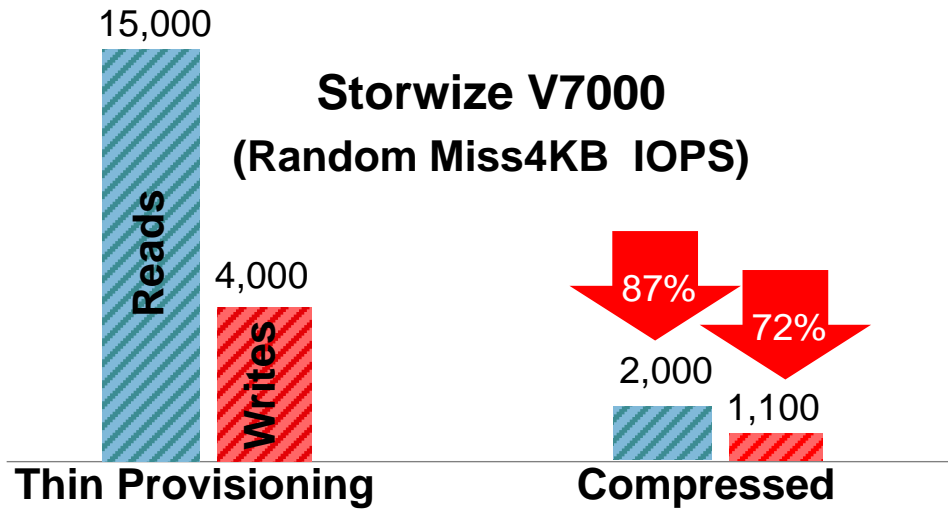
- Four corners workload performance profile
  - Random 4K and 8K block size, cache hit vs. cache miss, response times with 100% write
- Multi-threading capabilities and configuration
  - Request configuration details behind performance numbers
  - Will your hosts be able to drive similar performance?
- Feature overhead: What was turned on or off?
  - Which features were running during performance measurements?
  - Features like dedupe and compression can significantly degrade performance





# Avoid Overload: Understand impact of Software features

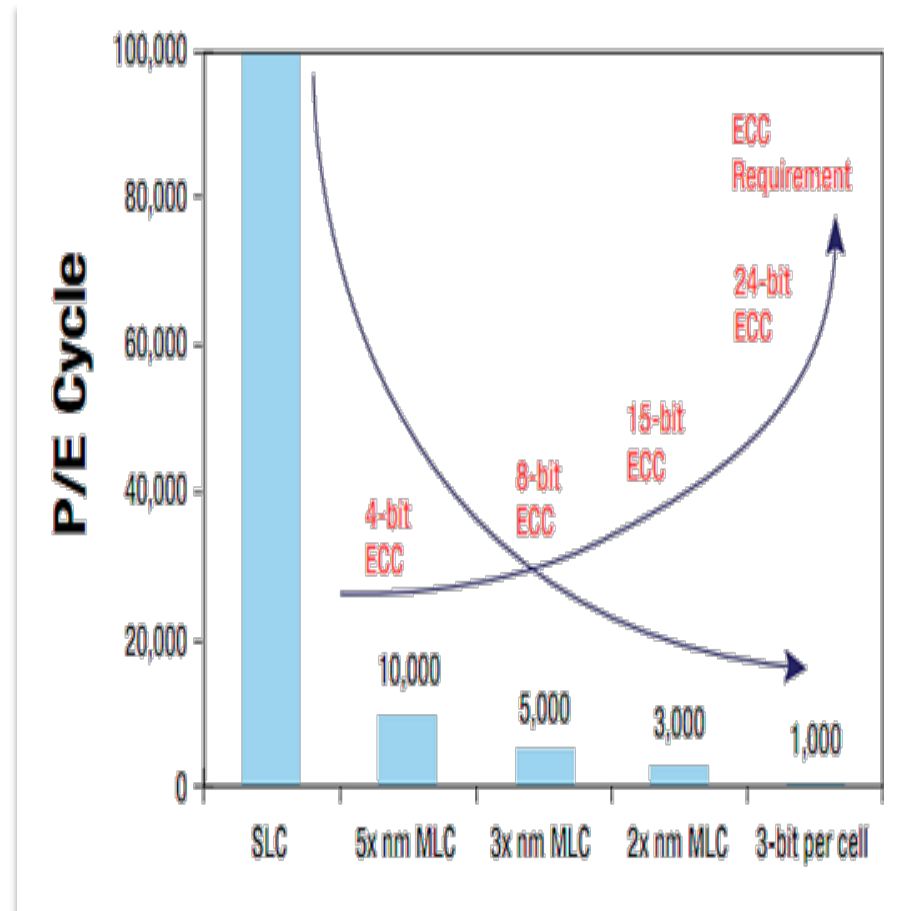
*What is the performance/Capacity Tradeoff Are Acceptable*



# Claim: “Flash Endurance Levels Are All The Same”


*Push Your Vendors For More Specificity*

- Real world endurance
  - Flash will wear out at different rate based on the technology and geometry
  - What special technology leveraged to extend useful life?
- Enterprise support is critical
  - Vendor must ensure technology provides advertised benefits
  - Will the vendor replace flash capacity under warranty or support if the flash reaches the end of its useful life?





# Let Your Workloads define Your Flash Investment Strategy



## Making the Best Decision Means Getting To Full Disclosure

- Establish your flash strategy early
  - Prioritize your flash investment based on application workload
- Assess economic impact of dedicated vs. shared flash
  - How will economics change with data growth
- Size your flash performance requirement based on middle of the road read/write ratio
- Use Proof of Concepts when concerned about new vendors' claims





**THANK YOU**