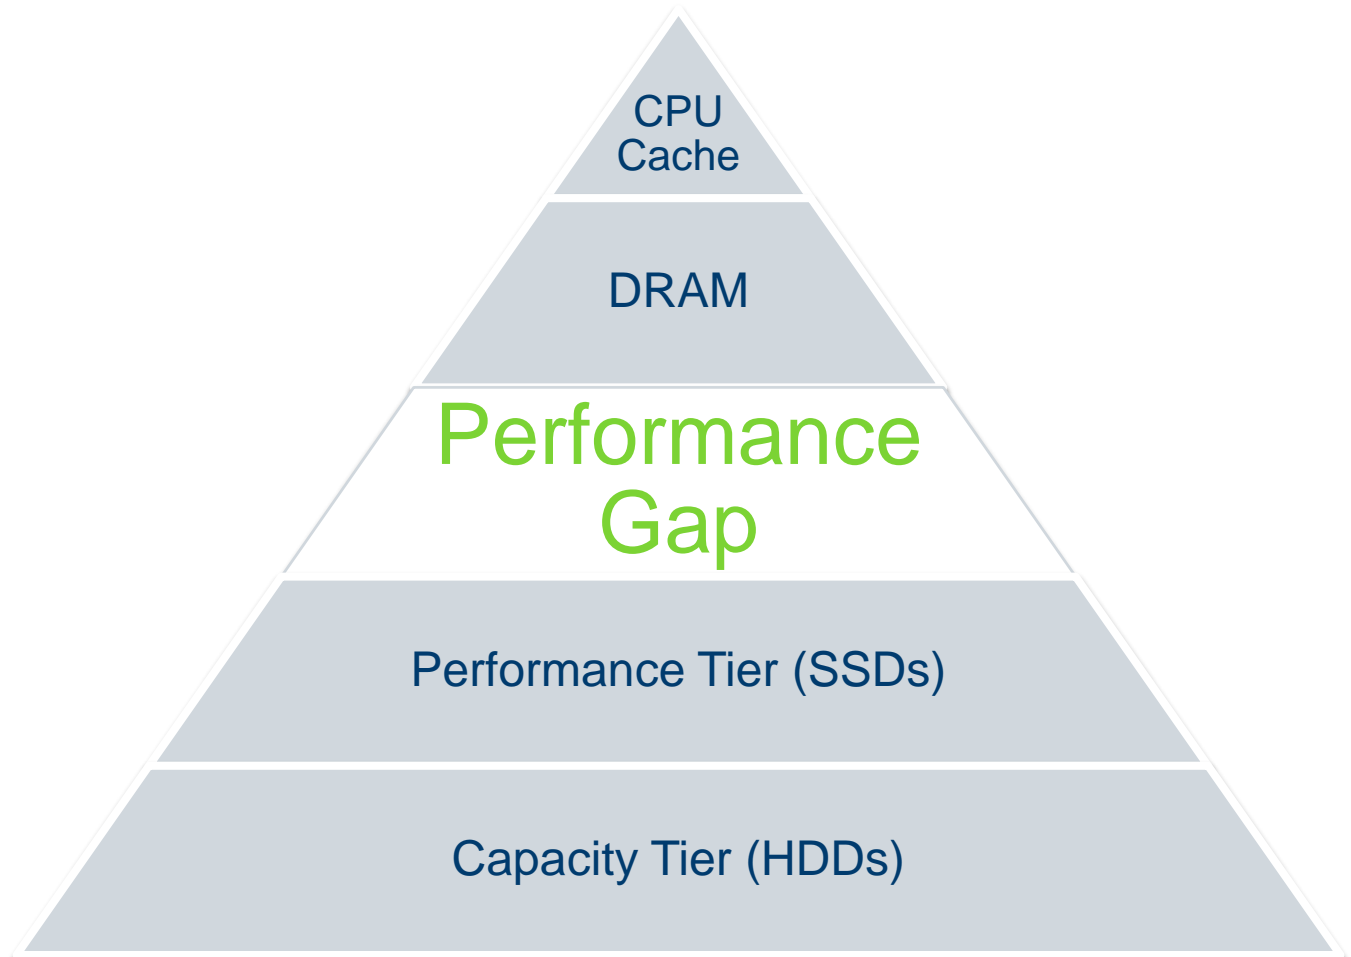





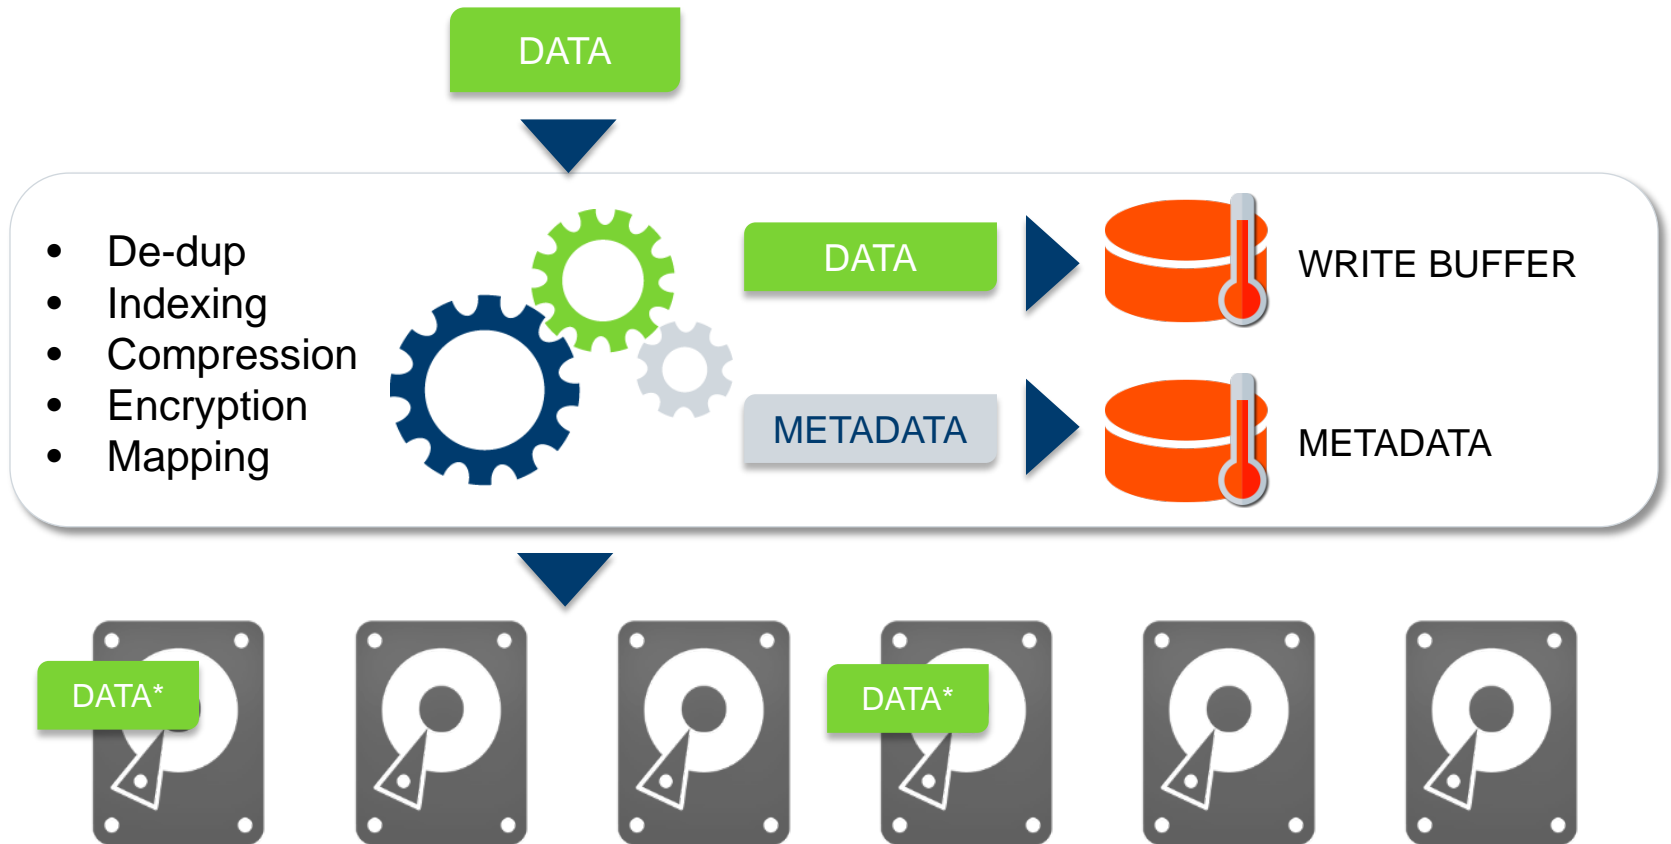
# PCIe Storage Beyond SSDs

Fabian Trumper  
NVM Solutions Group  
PMC-Sierra

# Classic Memory / Storage Hierarchy

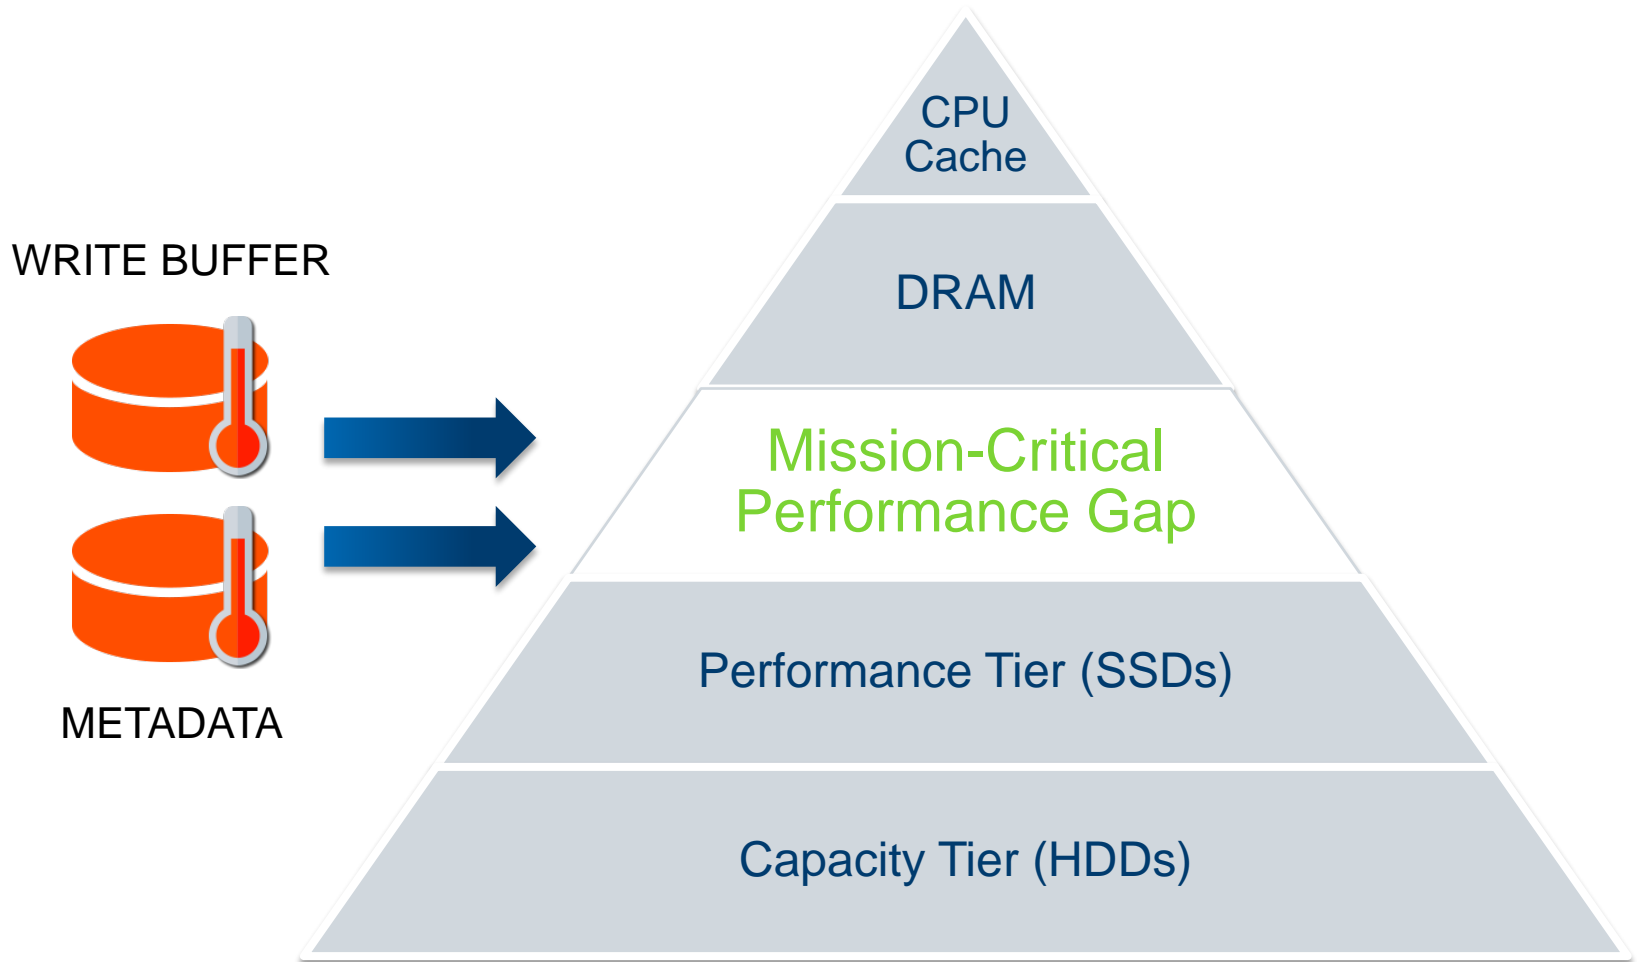


# Some system hot spots don't fit well within that hierarchy

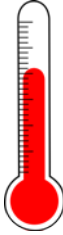
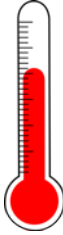


Hot Spot I/Os  $\geq \sum$  (all I/O in the System)


# Both fall into the Performance Gap



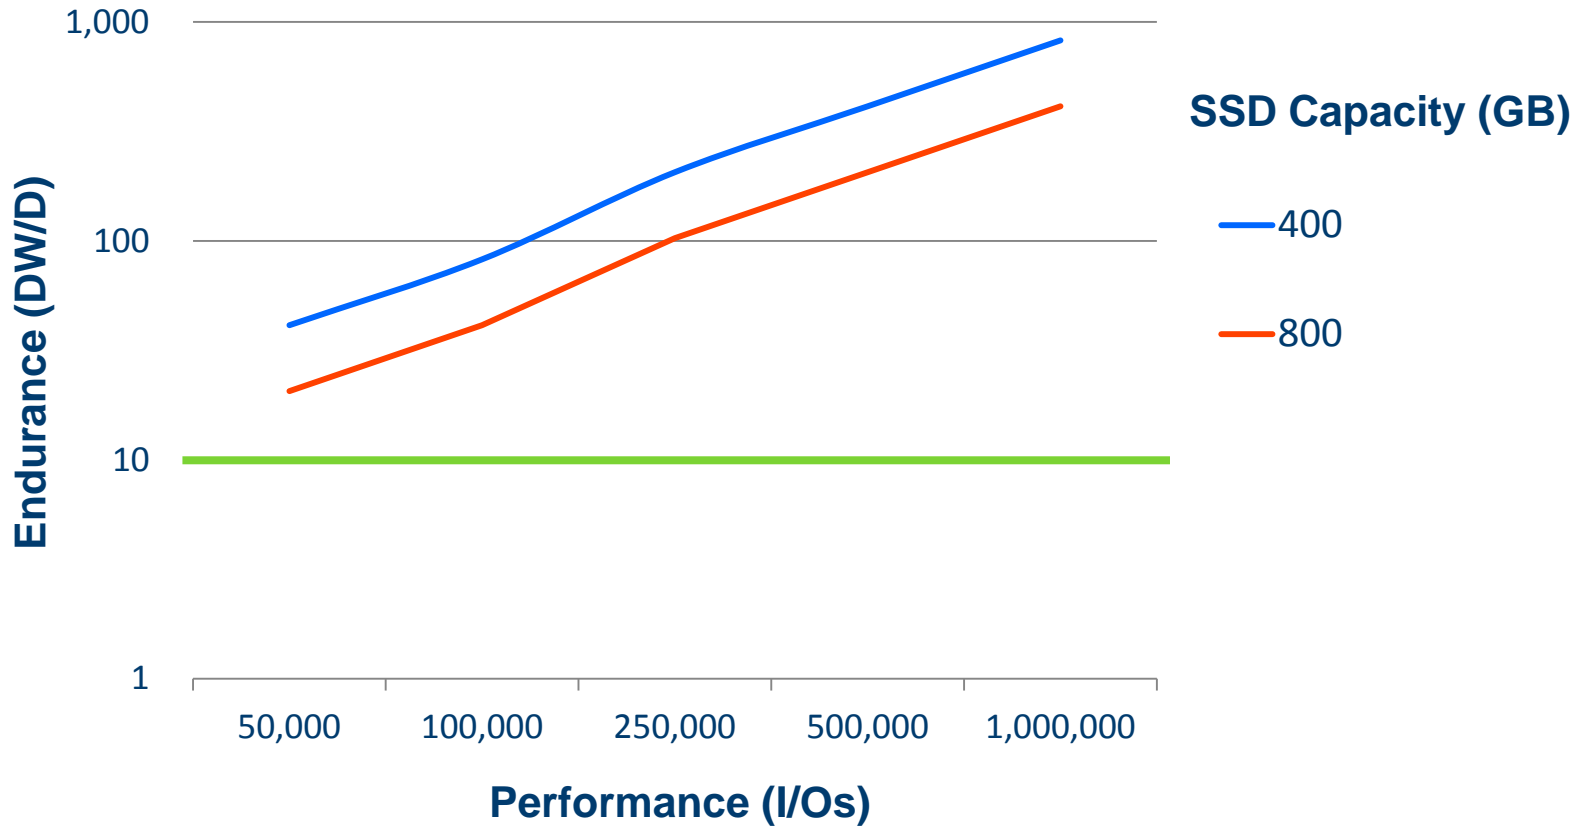
# Two Mission Critical Data Sets

Requirements	WRITE BUFFER 	METADATA 
Transaction Size	Variable / Small	Fixed
Bandwidth	Low/Moderate	Very High
Response Time	Very Low	Very Low
Non Volatile	Yes	Yes
Write Endurance	Very High	Very High
Multi-Port Access	Preferably	Preferably
Mission Critical (Fail Safe)	Yes	Yes

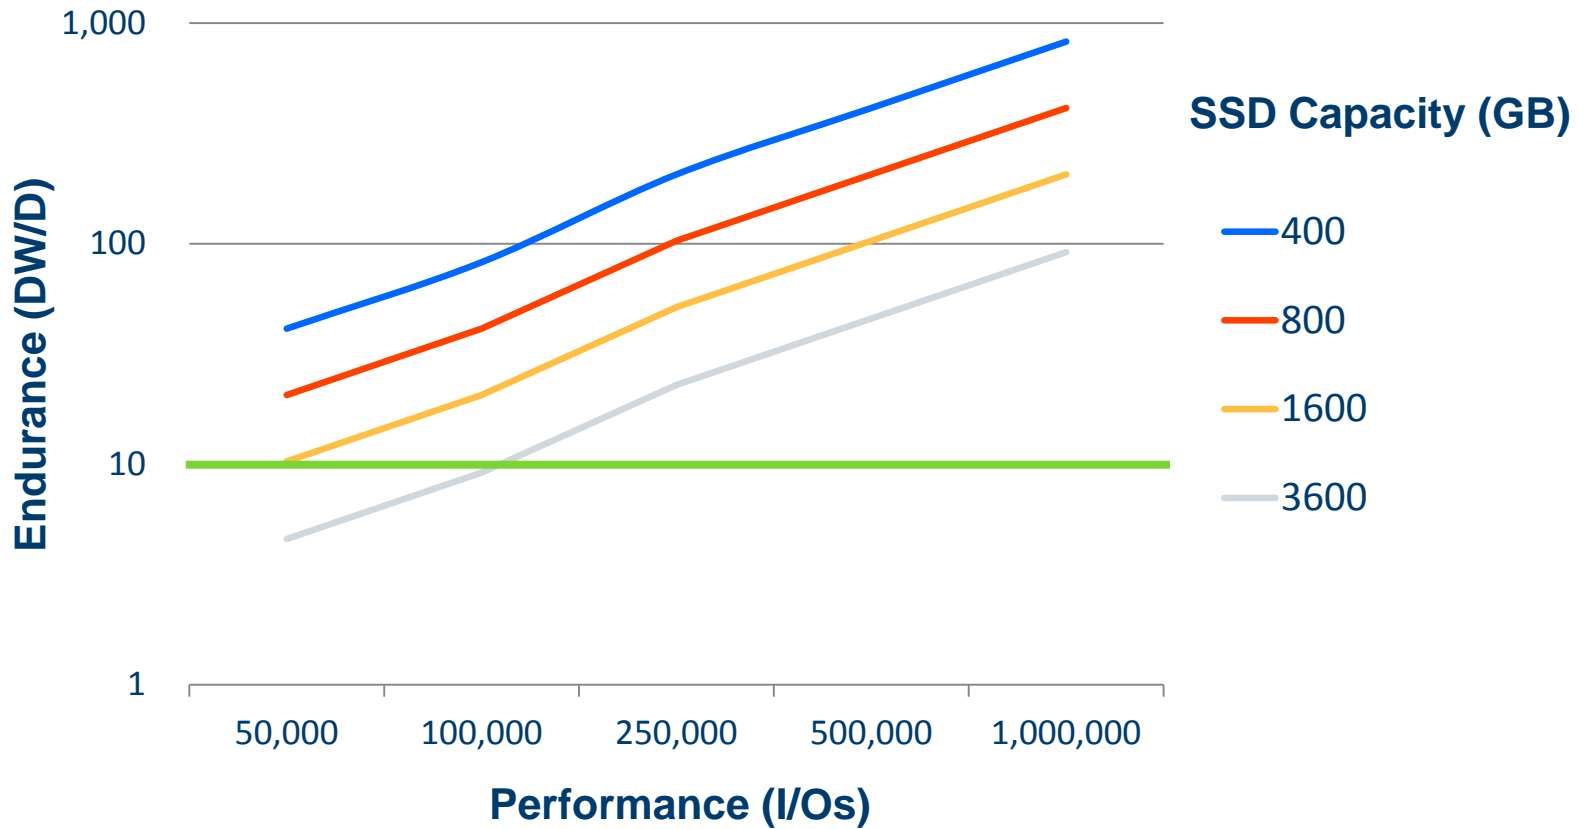
# Can we use SSDs for that?



# What if we doubled the capacity?




# Quadrupled? More?





# Can't solve this with Flash...




## How about DRAM then?

- DRAM is susceptible to **power failures**
- To protect mission critical information,

either -

- Use UPS or BBU ▶ maintenance nightmare
- Log everything to Storage ▶ Complexity, Performance and Endurance trade-offs



*Closed  
due to  
Power Outage!*

What about PCM? RRAM? MRAM?  
We've been waiting for years...

A photograph of an elderly woman with white hair, wearing a white jacket and a dark skirt, walking from left to right. She is using a wooden cane. She is walking past a large, dark, textured wall that has the words "Coming Soon..." written on it in large, white, sans-serif font. The wall appears to be made of dark panels or a rough surface. The woman is looking towards the sign. The overall scene is set outdoors on a sidewalk.

Coming Soon...



שנה ברטי!  
הנבחרת הנורמה הנורמה הנורמה  
הנבחרת הנורמה הנורמה הנורמה  
David luz

[davidluz2000.deviantart.com](http://davidluz2000.deviantart.com)

# Back to the drawing board..



# Can we combine both?

Requirements	 <b>DRAM</b> <small>POWERED BY</small> <b>DDR</b> <small>MEMORY™</small>	<b>Solid State Drive</b>  <b>PCI EXPRESS®</b>	<b>PCI-Express</b> <b>Memory Mapped</b> <b>Storage</b>
DMA Initiator	No	No	Yes
Memory Addressable	Yes	No	Yes
Variable Size Objects	Yes	No	Yes
Write Endurance	Yes	No	Yes
Multi-Port Access	No	Yes	Yes
Persistent	No	Yes	Yes



# Mt Ramon Project – Hybrid Memory Card


- So we went ahead and built it 😊
- 16GB NV-RAM
- NVMe and Direct Memory Interface Access Modes

4GB/8GB/16GB DRAM  
Unlimited Endurance

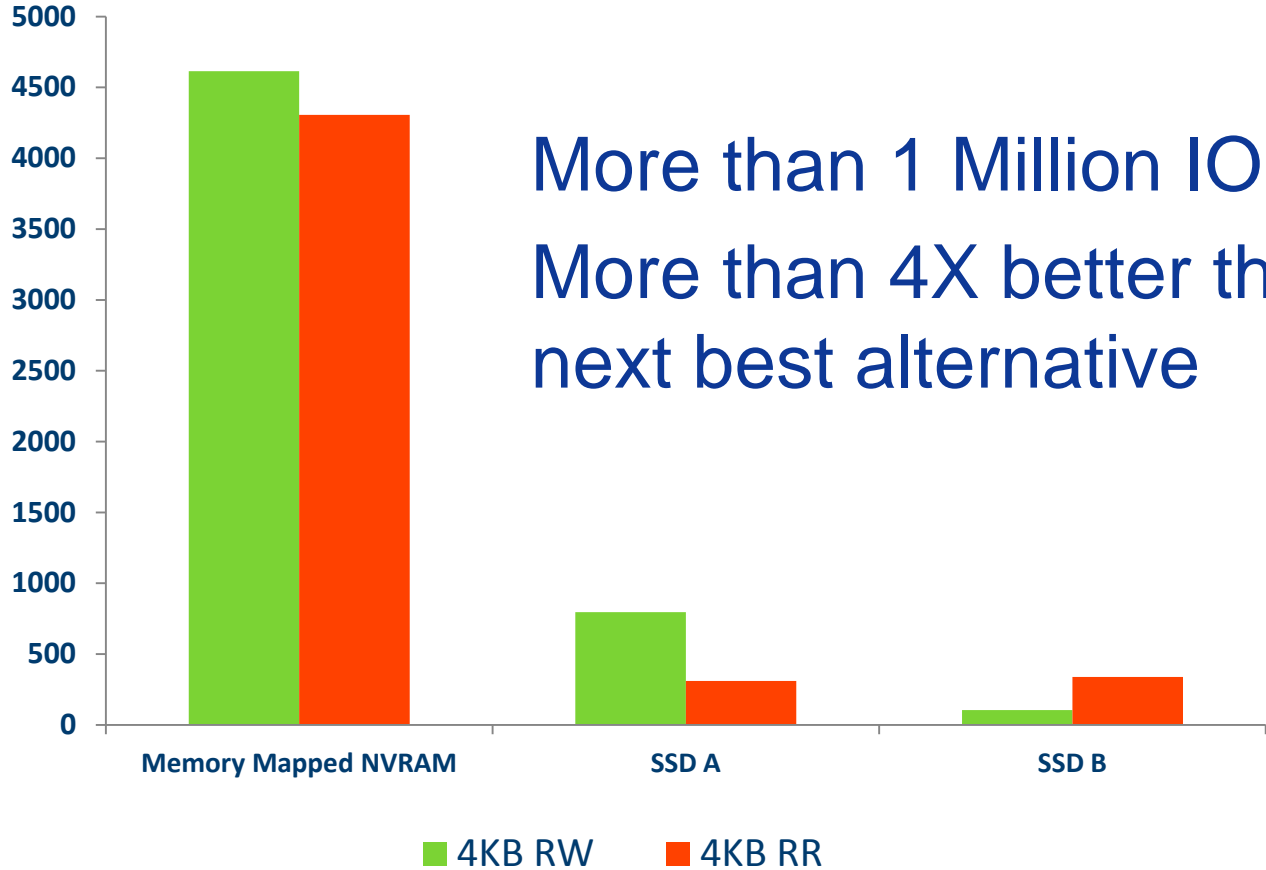
Flash Backup Module

PCIe 3.0 Performance  
Low Latency Access

Backup Power Module




# What did we get?



More than 1 Million IOPS  
More than 4X better than  
next best alternative

Sustained and Symmetrical Performance

# And what about CPU utilization?

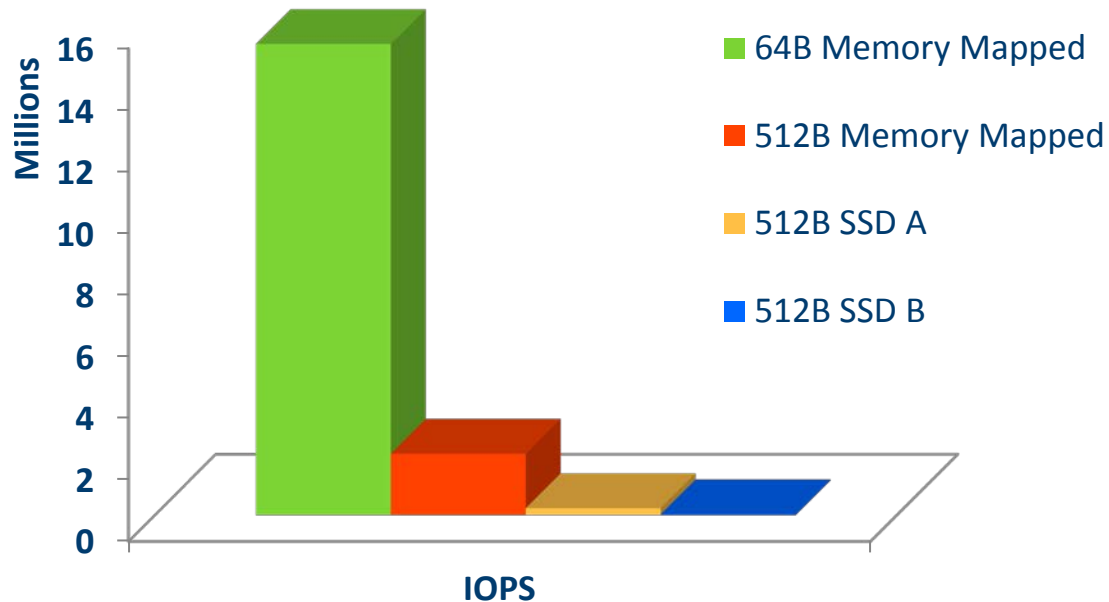


HW	Details
CPU	I5-3470
DDR	1333MHz
Frq	3.2Ghz
PCI	X8 Gen3
# of cores	4



# Memory Mapped Write IOPS

## 15 Million Random Write IOPS for 64B Transactions



- Memory Mapped Interfaces provides higher IOPS for small write transactions
- Ideal for in-place updates of metadata and database records

## What next?

- Enable the Industry to support Memory-Mapped Storage by working through standards bodies:
  - NVMeexpress
  - SNIA NVM Programming TWG
- Enable Applications and Operating Systems ISV
  - Adopt methods and APIs to recognize and take advantage of memory-mapped devices



# Thank You!



See us in booth # 416 for a live demo

[www.pmcs.com](http://www.pmcs.com)

[blog.pmcs.com](http://blog.pmcs.com)    [@pmcsierra](https://twitter.com/pmcsierra)