Panel: Persistent Memory Adoption in Operating Systems (-ish)

Moderator: Jeff Chang | AgigA Tech, NVDIMM SIG Co-Chair

Meet The Experts

Tom Talpey
Architect

Microsoft

Tom Coughlan
Senior Engineering Manager

redhat

Amit Golander
CTO

PLEXISTOR

Pratap Subrahmanyam
Fellow

vmware

NVDIMM

Non-Volatile Dual Inline Memory Module:

[Hardware] A dual inline memory module that operates as standard RAM while also having persistence across power cycles.

- pg 190, 2016 SNIA Dictionary

The Three Legged Stool

The Three Legged Stool

The Three Legged Stool

HARDWARE BUILDING BLOCKS

The Three Legged Stool

HARDWARE BUILDING BLOCKS

SOFTWARE INFRASTRUCTURE

The Three Legged Stool

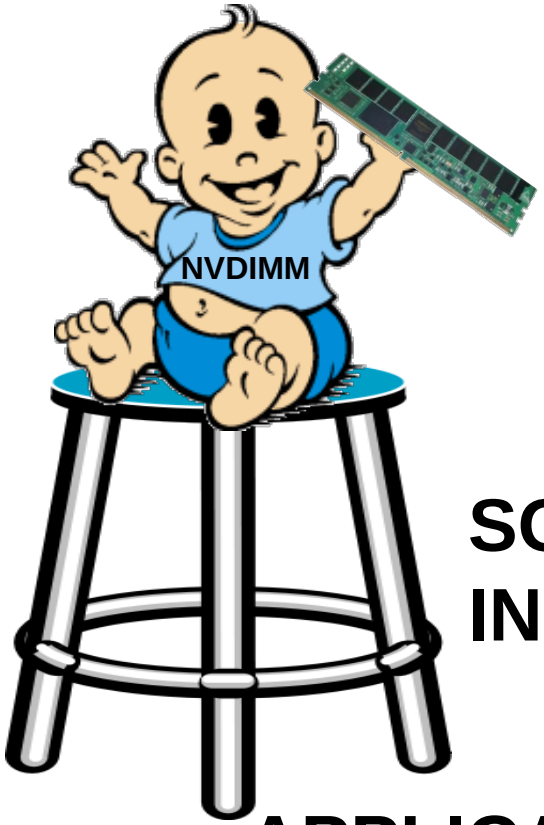

HARDWARE BUILDING BLOCKS

SOFTWARE INFRASTRUCTURE

APPLICATIONS

The Three Legged Stool

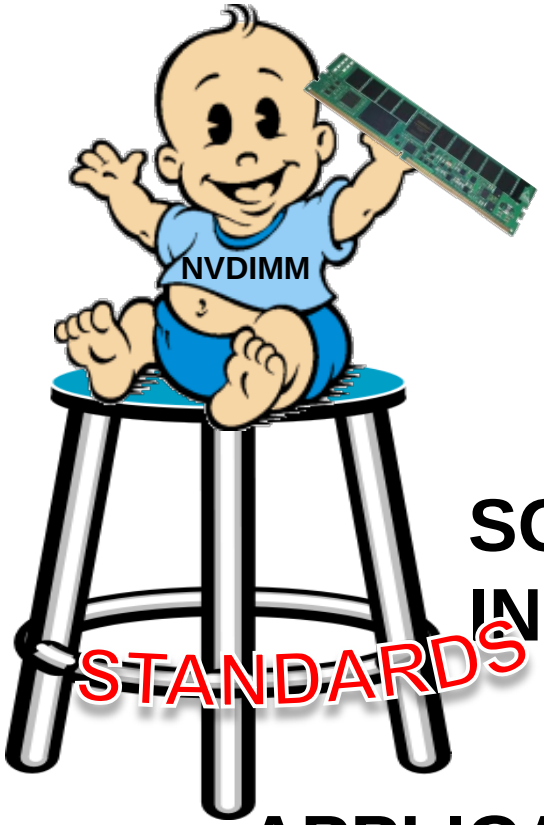

HARDWARE BUILDING BLOCKS

SOFTWARE INFRASTRUCTURE

STANDARDS

APPLICATIONS

The Three Legged Stool

VALIDATION PROGRAM

HARDWARE BUILDING BLOCKS

SOFTWARE INFRASTRUCTURE

APPLICATIONS

The Three Legged Stool

VALIDATION PROGRAM

HARDWARE BUILDING BLOCKS

SOFTWARE INFRASTRUCTURE

APPLICATIONS

First Question

- Amit, what does Plexistor do?

PM Adoption

Amit Golander, PhD
Plexistor, CTO

PLEXISTOR

Where does Plexistor fit in?

Enterprise-grade platform to make PM trivially consumable by applications →

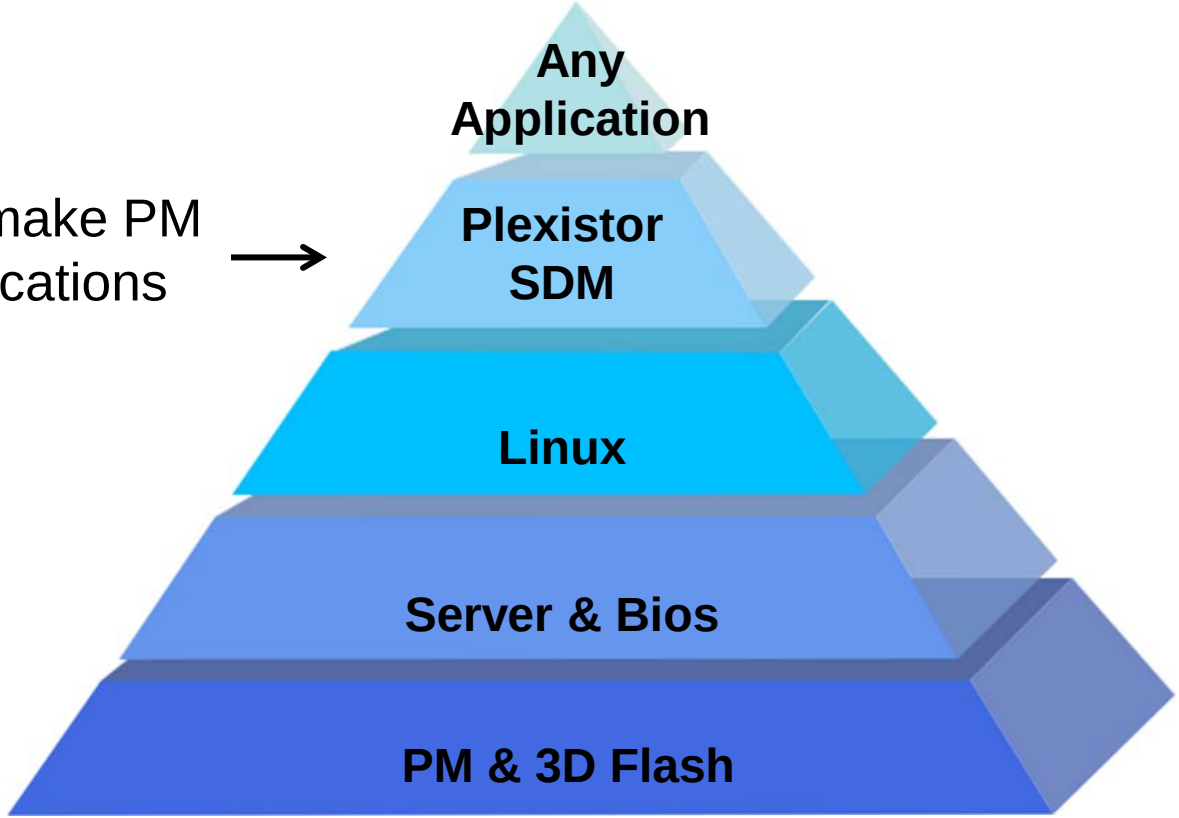

From Promise to Delivery in 3-4 years

Promise → Deliver

	PM Summit Focus	Plexistor Status	
Jan. 2013	What are the problems?		
Jan. 2014	The **Future of Computing** – The convergence of Memory and Storage	Concept	Prototype
Jan. 2015	Realizing the Benefits of the Convergence…	Alpha	AutoTiering
Jan. 2016	The Status of the Convergence…	Public Beta	DAX, Stability
Jan. 2017	**Delivering** on the Convergence…	GA	Clones, PMoF

Use cases →

State of the Nation (Jan 2017)

Legacy and PM aware

SDM v2.1
as SW or pre-installed SKUs (partners)

Kernel 4.4+
Ubuntu 16.04
RHEL 7.3

HPE Proliant Gen9
Supermicro X10

16GB NVDIMM-N
JEDEC FW

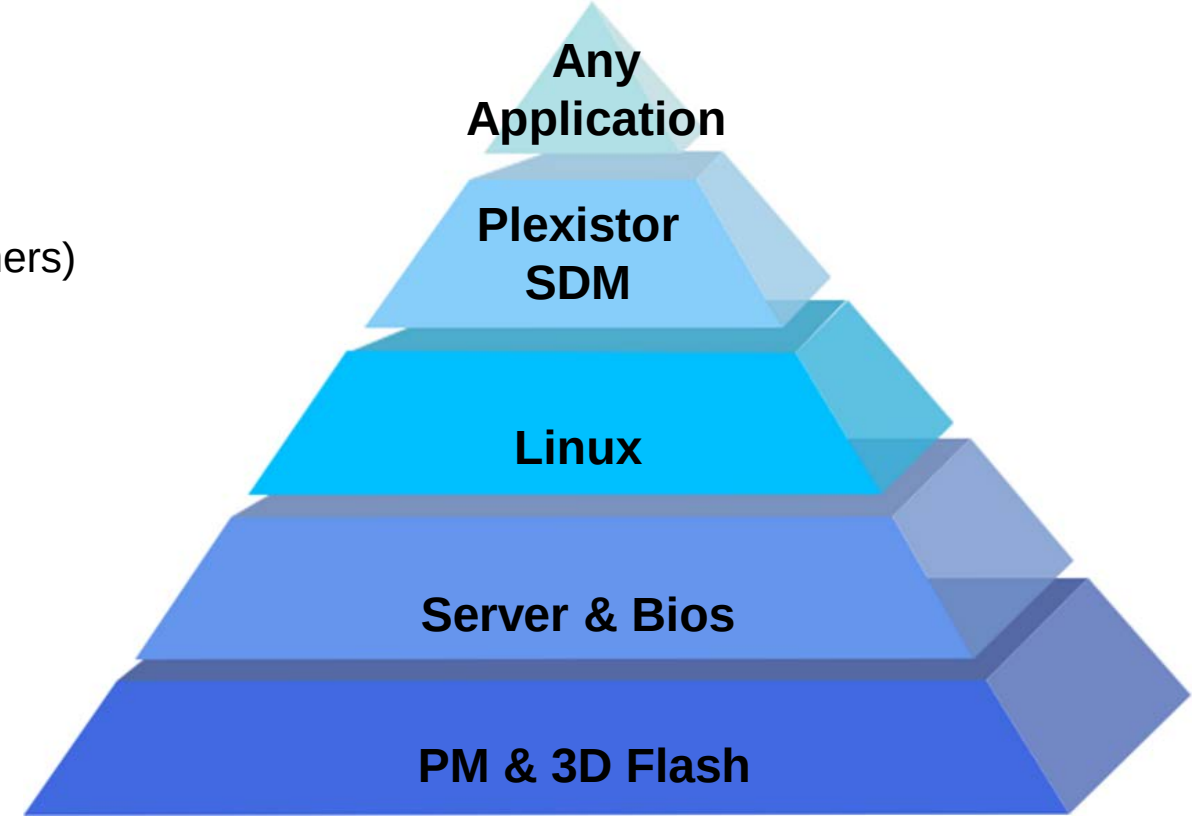

The Plexistor SDM Solution

- Installed on top of Linux w/ NVDIMM driver www.plexistor.com/download/
- Runs any application on bare metal, container or virtualization
- Enterprise-grade features:
  - Storage and Memory semantics concurrently
  - Auto-tiering between NVDIMMs and Flash SSDs
  - Light-weight Clones
  - Mirroring to PMoF
  - NUMA optimized, Dashboarding, Proactive support, Forward migration…

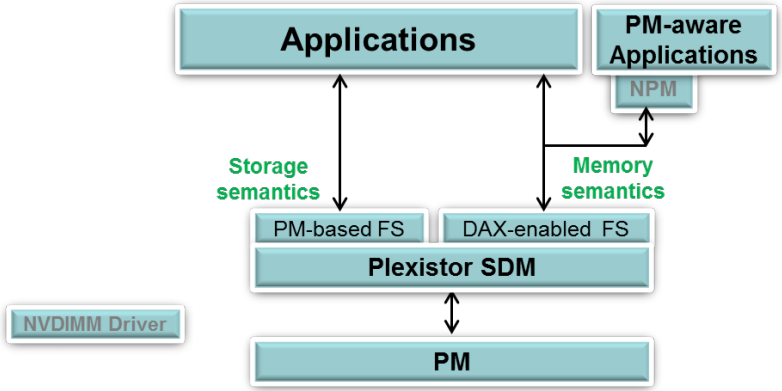

Performance

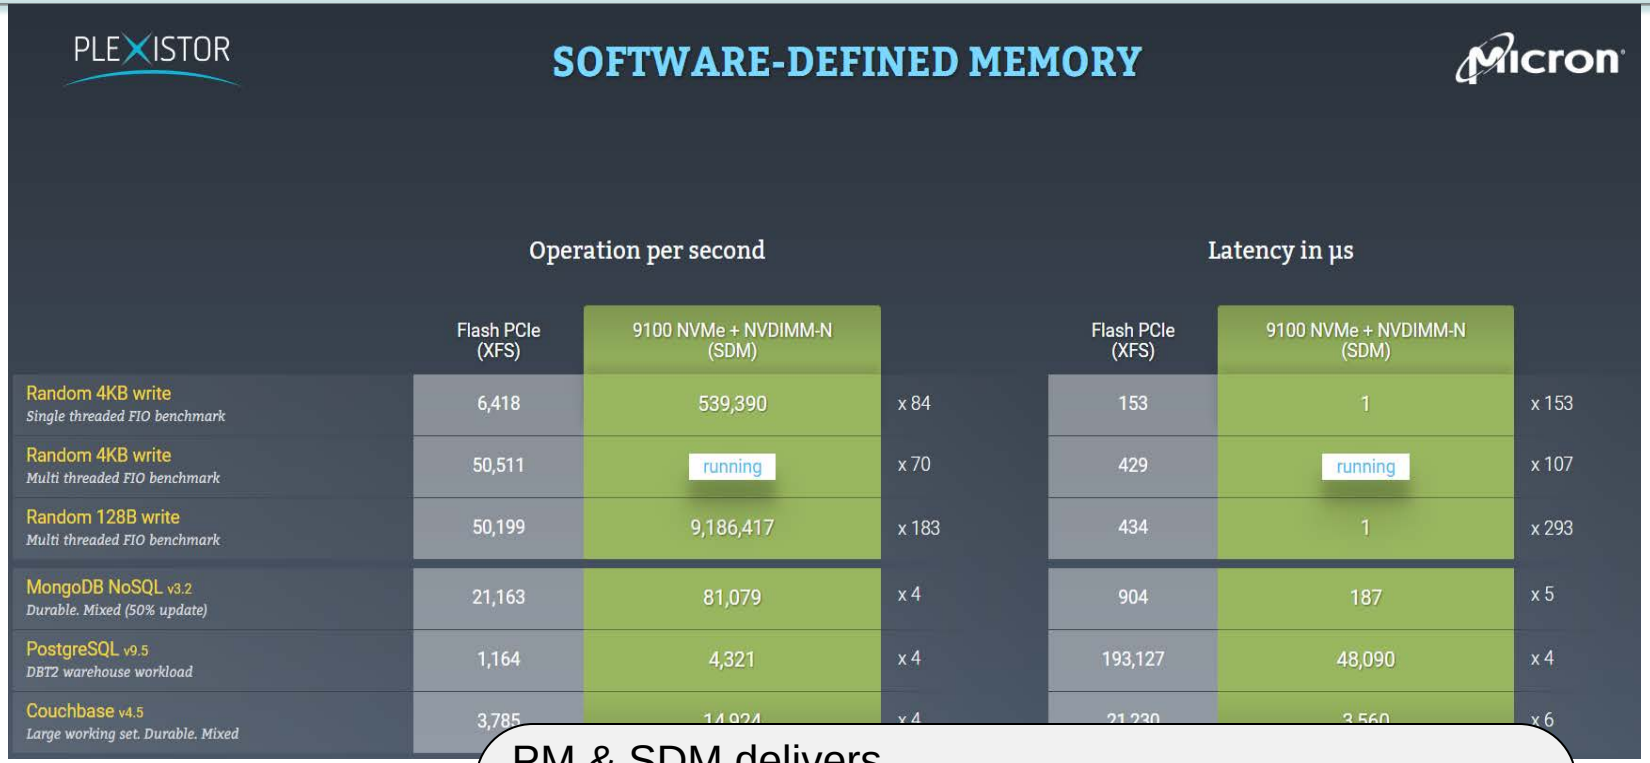

	Operation per second			Latency in µs		
	Flash PCIe (XFS)	9100 NVMe + NVDIMM-N (SDM)		Flash PCIe (XFS)	9100 NVMe + NVDIMM-N (SDM)	
Random 4KB write *Single threaded FIO benchmark*	6,418	539,390	x 84	153	1	x 153
Random 4KB write *Multi threaded FIO benchmark*	50,511	running	x 70	429	running	x 107
Random 128B write *Multi threaded FIO benchmark*	50,199	9,186,417	x 183	434	1	x 293
MongoDB NoSQL v3.2 *Durable. Mixed (50% update)*	21,163	81,079	x 4	904	187	x 5
PostgreSQL v9.5 *DBT2 warehouse workload*	1,164	4,321	x 4	193,127	48,090	x 4
Couchbase v4.5 *Large working set. Durable. Mixed*	3,785	14,924	x 4	21,230	3,560	x 6

PM & SDM delivers,
provides data services
and is production ready

PMoF Sneak Peek

- SDM v3.0 www.youtube.com/watch?v=geZSYJxPPfQ&feature=youtu.be

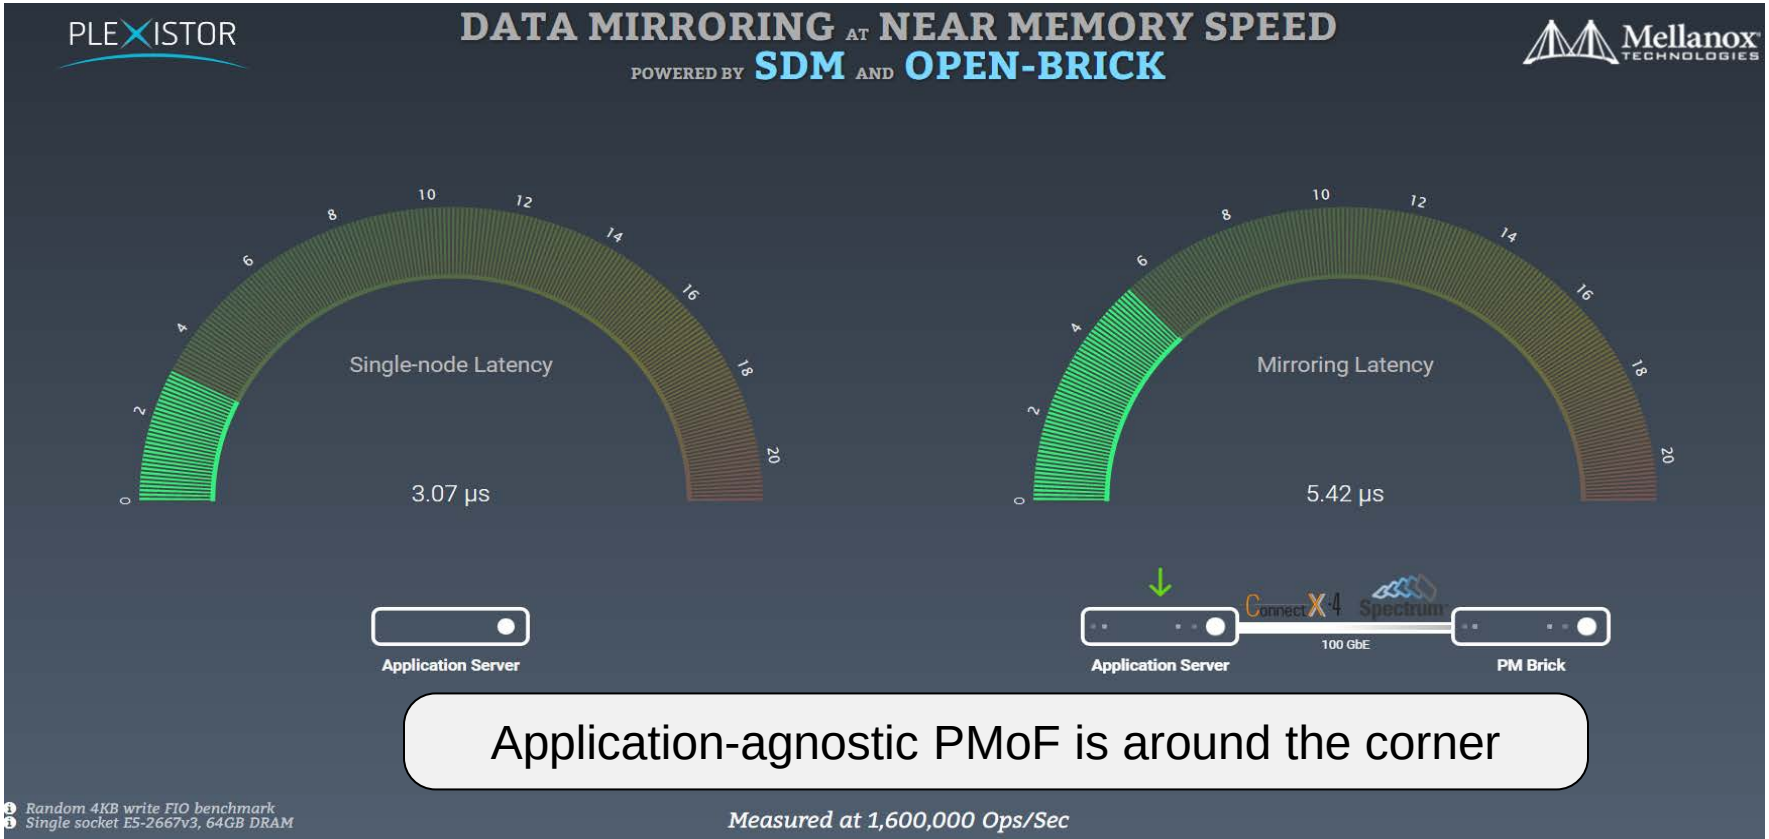

Example: Oracle Database for OLTP

Baseline

Exadata X6-2 Full Rack
Running Oracle RAC 12C

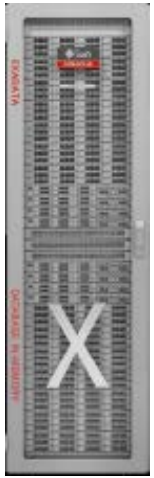

Max IOPS	5.6M
Throughput	21GB
Latency	250us
#Cores	72

License cost **$2.1M**

Plexistor Solution

1 HPE 2x2699 server
Running Oracle 12C

Max IOPS	6.5M
Throughput	30GB
Latency	2us
#Cores	24

License cost **$540K**

Swiss Federal Government

Run the exact same application
Faster and at **75% lower cost**

Now To The Panel…

Questions????

Tom Talpey
Architect

Microsoft

Accelerating SQL 16 with PM

	Row Updates / Second	Avg. Time / Txn (ms)
NVMe SSD	63,246	0.379
Dax Mode NVDIMM-N	124,917	0.192

From Neal Christiansen @ FMS'16

Tom Coughlan
Sr Engineering Mgr

redhat.

Modified Kernel Subsystems

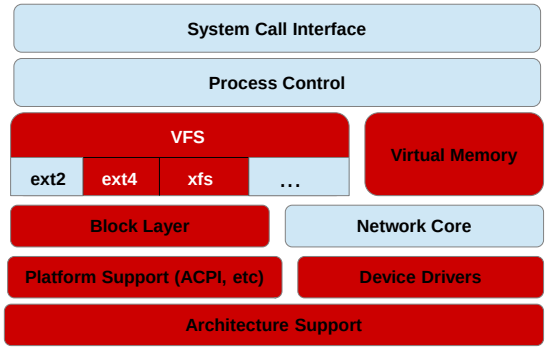

Amit Golander
CTO

PLEXISTOR

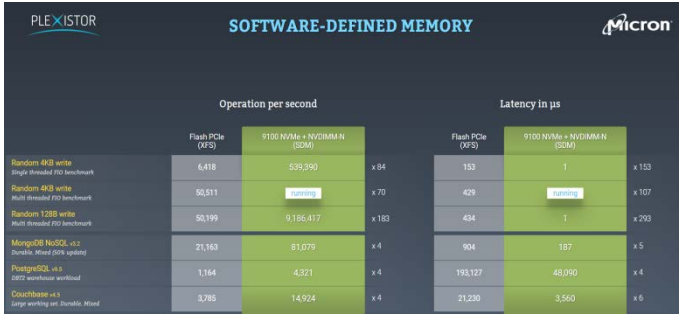

Pratap Subrahmanyam
Fellow

vmware

vSphere-based NVDIMM Emulation Vehicle

- **Available Now**
- **Emulates all of the capabilities of NVDIMMs from different vendors**
- **Works with off-the-shelf commercial servers**