



# Persistent Memory (PM) for Big Data & Analytic

Kumar Prabhat, Sr Director Product,  
Plexistor

The logo for the Flash Memory Summit features a stylized yellow sunburst with multiple rays. Below the sunburst, the word "Flash" is written in red, "Memory" in black, and "SUMMIT" in white on a blue rectangular background.

# Flash Memory Summit Agenda

- Enterprise Applications/IT Environment (2006 & Beyond)
- Emergence of New Generation of Applications
- The In Memory Challenges
- Opportunity for Flash and PM
- A New Approach to Storage
- Benefits
- Summary



# Enterprise Applications/IT Environment (2006 and Beyond)



Mobility



Social  
business



Big Data  
Analytics



Cloud

- Volume
- Velocity
- Variety



- Big Data
- Fast Data



# Emergence of In Memory Applications and No SQL: 2010 & Beyond



- Millions of nodes deployed
- Massive Adoption of open source
- In-memory compute gave up on **Storage** and moved the working set to **Memory**
  - tremendous performance gains, but challenging



# The In-Memory Challenge

- DRAM is limited and expensive
- Puts Data at Risk
- Suffers from slow recovery time during power failures
- Over provisioning to achieve QoS for dynamic workload
- Multiple compute & data silo's optimized for different applications

**BIG DATA → BIG DATA CENTER**



## Good News...

- Emergence of persistent memory
  - different cost/performance characteristics
  - NVDIMM, 3DxPoint, ReRAM
- Massive Adoption of Flash in Data Center
  - SAS, SATA, PCIe, NVMe
- Availability of Network Speed (40GB -> 100 GB)



## Opportunity

- Bring the best of both worlds
  - leverage persistent memory and Flash
- Memory like speed to address application latency requirements
- High capacity & low cost of traditional storage



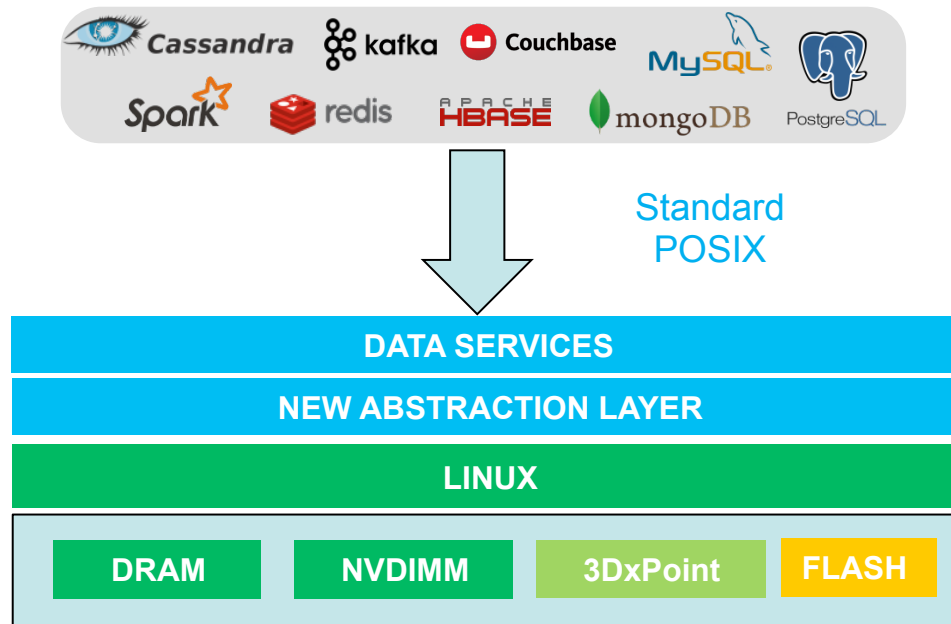
## A New Approach to Storage

- New Software Stack for Persistent Memory
- Convergence of memory and storage
- Unlock the power of byte addressability of memory
- Tier Architecture with SCM and Flash
- Memory Oriented Data Services
- Support for standard API
  - POSIX, NVML





# A New Approach





# Benchmark running in Amazon EC2

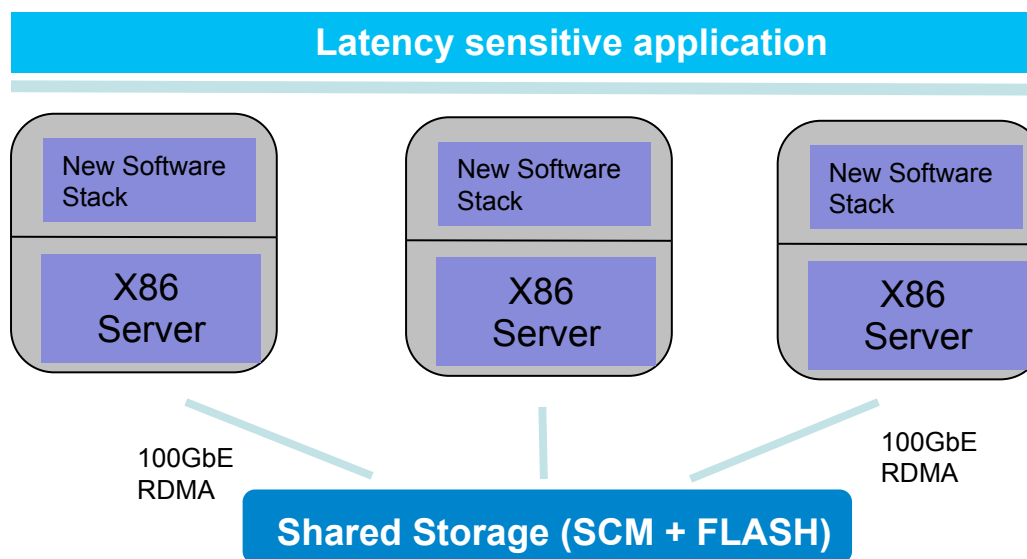
		Operation per second			Latency in $\mu$ s		
		ZFS	XFS	PLEXISTOR SDM	ZFS	XFS	PLEXISTOR SDM
<b>Random 4KB write</b>	<i>FIO benchmark</i>	2,146	26,068	5,057,700	8,313	1,452	3
				x 2357			x 2771
<b>Random 128B write</b>	<i>FIO benchmark</i>	2,177	27,944	7,080,780	8,263	1,358	2
				x 3253			x 4051
<b>Cassandra v3.0.4</b>	<i>Large working set. Mixed(50% update)</i>	4,395	6,398	25,535	8	8	7
				x 6			x 1
<b>MongoDB v3.2</b>	<i>Durable. Mixed</i>	30,194	31,168	90,532	1,447	1,399	346
				x 3			x 4
<b>Couchbase v4.5</b>	<i>Large working set. Durable. Mixed</i>	5,047	4,514	18,656	20,692	22,082	3,396
				x 4			x 6
<b>PostgreSQL v9.5</b>	<i>DBT2 warehouse workload</i>	870	1,050	5,052	595	535	22
				x 6			x 27

\* To see live dashboard, press the running button. To reset data, click [here](#)

ⓘ E5-2650 v3 CPU, 32GB DRAM, 32GB NVDIMM, CloudSpeed SSD



# Performance Benefit – Shared Storage

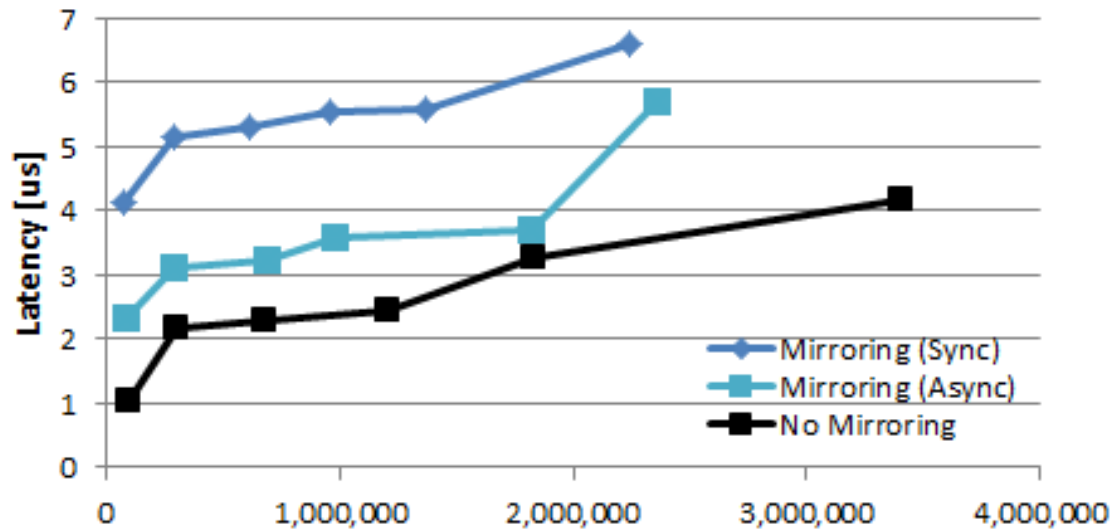


Millions of writes utilizing Mellanox 100GbE/RDMA with few  $\mu$ s of latency, an order of

magnitude better than Rack Optimized Flash Storage solutions such as DSSD (~100  $\mu$ s)



# Very Fast Data Replication



Using Mellanox 100GbE with RDMA Rnd 4KB wr/s

**At 1 million write ops/s:**  
**(40Gb)**

- Async < 4us
- Sync < 6us



# Benefits



## Predictable End User Experience

- Address dynamic work loads
- Eliminate storage spikes



## Single Storage Platform

- Run all NoSQL DB's - I/O & Latency sensitive
- Persistent



## Easy to Deploy & Manage

- No overprovisioning, no dedicated hardware
- Faster Recovery
- No Application Modification



## Summary

- Memory and storage convergence is the future!
- Persistent Memory is great technology, but also complex and expensive than Flash!

The right approach will be to mix FLASH with PM  
and new software stack