



Solving Big Data Delivery Issues

In-Memory Database, Near Data Processing
& Persistent Memory

Gil Russell - Semiscape

IMDB Market Motivation

- Real Time Analytics and Transaction Processing market drivers
- Real time analytics – reason for “your” IT to be on the Cloud
- “Compound Annual Growth Rates” rated between 25% to 43%+
 - ✓ Market leader SAP HANA IMDB achieved 120% growth rate in 2013
- Equipped servers viewed as the birth of “cognitive computing intelligence and predictive analytics” - the next era of computing
 - ✓ i.e. Watson playing Jeopardy in February 2011
 - ✓ Growth rate limited only by availability of application software

SNIA NVM Programming model

The Four Modes

Block Mode Innovation

- Atomics
- Access hints
- NVM-oriented operations

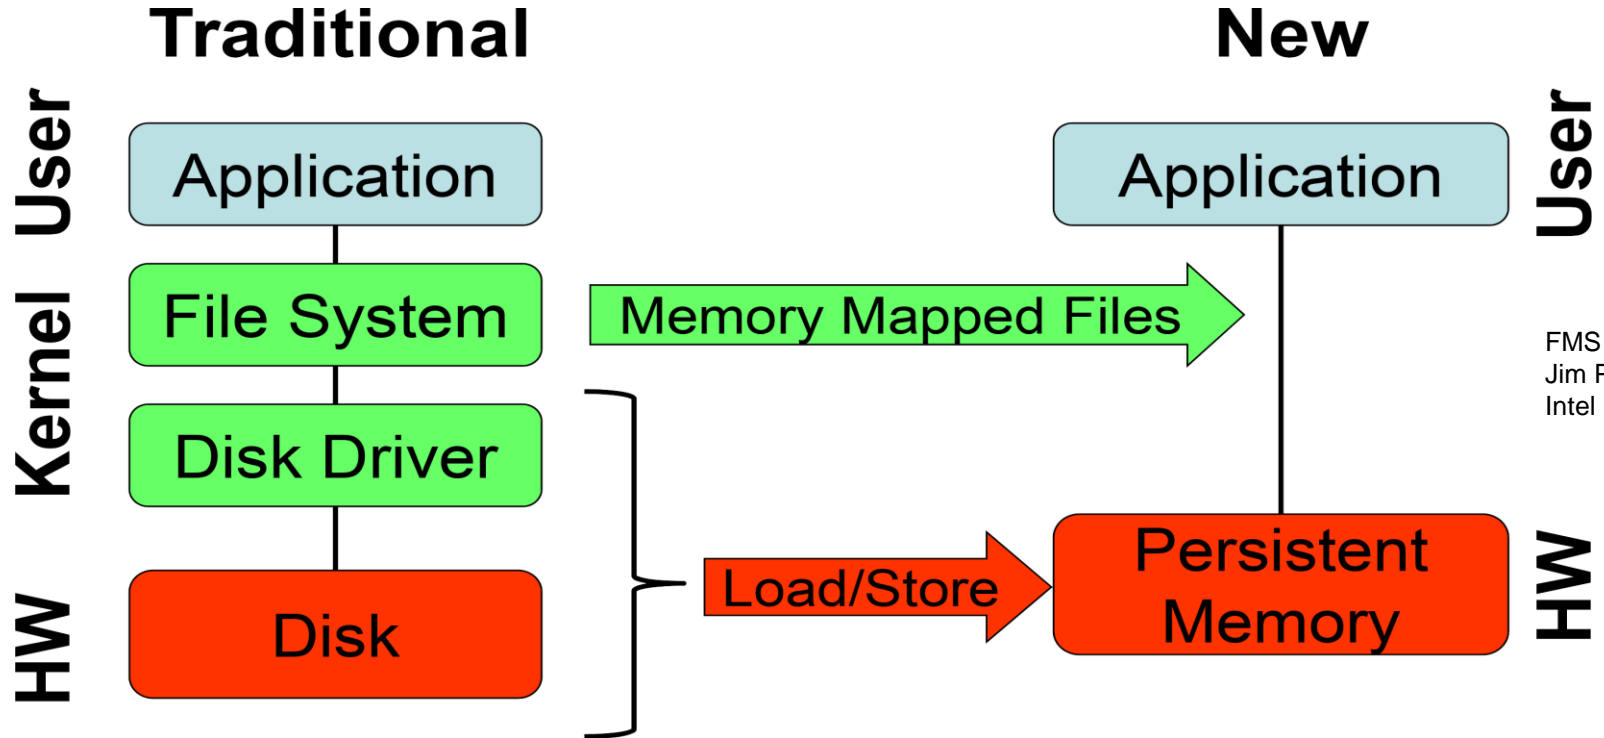
Emerging NVM Technologies

- Performance
- Performance
- Perf... okay, cost

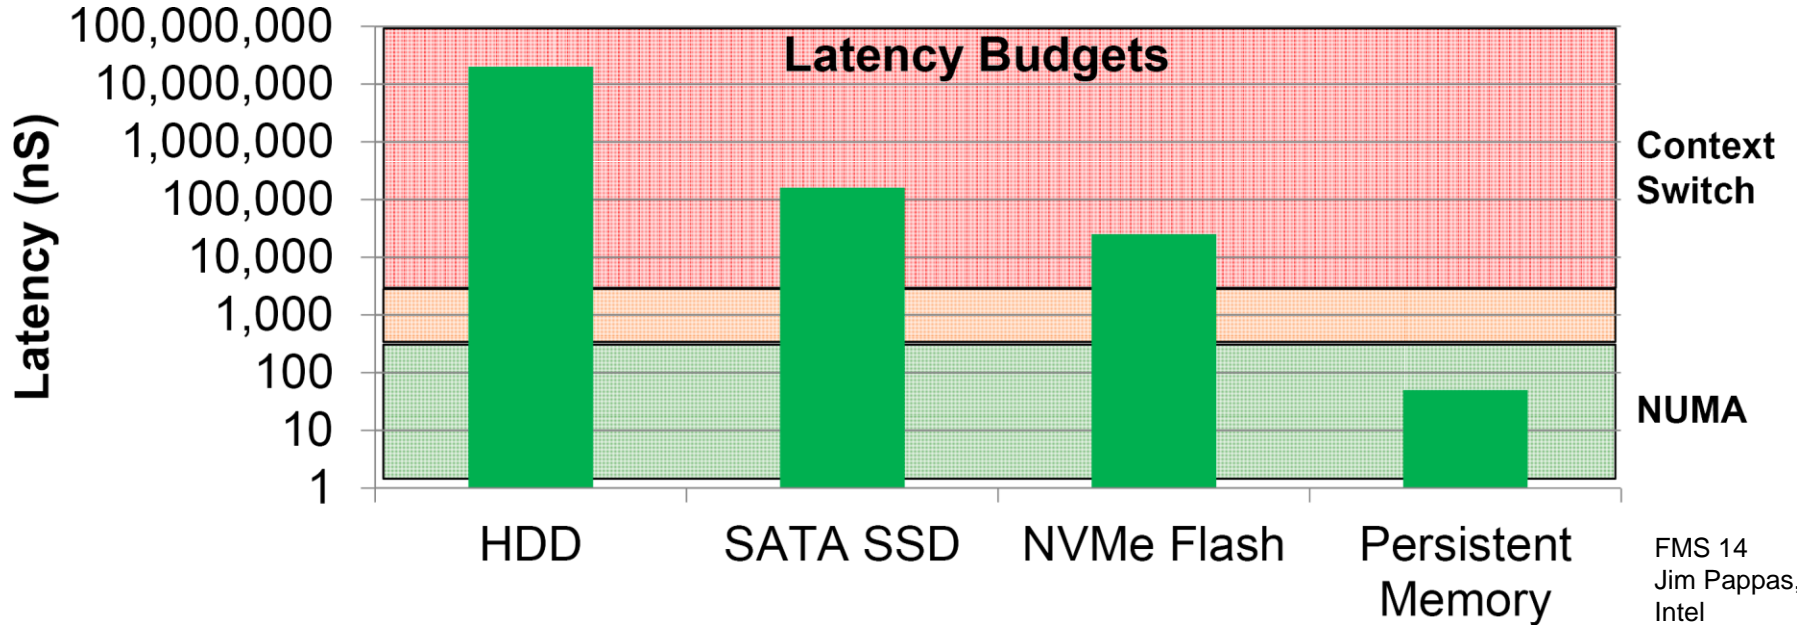
	Traditional	Persistent Memory
User View	NVM.FILE	NVM.PM.FILE
Kernel Protected	NVM.BLOCK	NVM.PM.VOLUME
Media Type	Disk Drive	Persistent Memory
NVDIMM	Disk-Like	Memory-Like

FMS 14
Jim Pappas,
Intel

Eliminate File System Latency with Memory Mapped Files




FMS 14
Jim Pappas,
Intel



Data Access Architecture Changes

- In-Memory Database (IMDB)
 - Places entire Database in DRAM (DIMMs)
 - Minimizes data Latency
 - Maximizes data Bandwidth
 - Lowers system power
 - Data optimization required
 - Currently limited to 6 TB (64GB DIMMs)


Lenovo System x3850 X5



Hypothetical IMDB DIMM Instantiation

- IMDS DRAM main store [256 GB to 96 TB]
 - 32 DIMMs [256 GB] to 96 DIMMs [96TB]
 - ✓ DRAM portion enabled by introduction of 64Gb, x4, 4 high TSV, LRDDR4, 128GB DIMMs
 - ✓ DRAM Legacy support [12 TB]
- Persistent Memory portion enabled by introduction of 3D XPoint
 - ✓ Byte Addressable
 - ✓ 128Gb XPoint, TSV stacked x4, x8, (64GB, 128GB per stack)
 - ✓ Cost decline per Terabyte PM \leq (?)
 - ✓ Persistent Memory mandatory for IMDS (durable state retention)
 - ✓ No longer a need for “flush on fail” or recovery
 - ✓ Refresh reduction lowers Data Center power requirements

Hypothetical IMDB DIMM Instantiation



Near Data Acceleration (NDA) Posit

- Near Data Acceleration (NDA) allows massively parallel search function (binary, ternary)
- Through Silicon Via (TSV) pipes highly parallel DRAM & XPoint data to the logic interposer below stacked devices.
 - Simple memory accelerator on the logic chip
 - FPGA search engine for application specific execution
 - Compatible with commodity DRAM & XPoint
 - System compatible with MapReduce
- Reduces system power by 63% - 96%
- Improves performance by 3.2 to over 60 times
- Scales linearly with memory size

Near Data Acceleration (NDA) Posit



Near Data Processing (NDP)

- Near Data Processing performs best in a linear address space
 - ✓ NDA is limited in this respect
- Hybrid Memory Cube (HMC) is an attractive candidate for NDP
 - ✓ Placing XPoint on an HMC stack with DRAM under consideration
 - ✓ XPoint dramatically increases density - increased latency and reduced bandwidth is the drawback
 - ✓ The idea of a wider matrix of HMC CPU “hives” interconnected with photon pipes is under discussion/development
 - ✓ Persistence is not part of the HMC 2.0 Spec.
 - ✓ Intel is not a declared member of HMC Consortium

Near Data Processing (NDP) Posit

- Legacy compatibility is required for each successive DRAM generation – this may be ending
 - Silicon Photonics will replace copper but is not close to being cost effective
 - HPC budgets may cover high costs but the cost needs to come down an order of magnitude before widespread adoption begins
 - Introduction of XPoint memory opens a new episode if not an entirely new segue in computing
 - DRAM “generational step” performance has been decreasing
 - DDR4 may be the last generation of capacitor based storage devices

Summary

- IMDB & NDP is a practical near term solution for managing Big Data
- Moving from CPU Centric to Data Centric will rearrange the entire computing technology landscape
- Effect of Storage Class Memory on NAND-Flash SSDs:
 - ✓ Little or no affect in the near term
 - ✓ Storage Class Memory 2 will begin to affect the NAND-Flash SSD market starting in 2017

Thank you for coming!