



Flash Memory Summit

NVMe™ SSD with Persistent Memory Region

Chander Chadha

Sr Manager Product Marketing,

Toshiba America Electronic Components, Inc.

Santa Clara, CA
August 2017



Flash Memory Summit

Agenda

- What is Persistent Memory
- Concept of NVM Express® (NVMe™) SSD with Persistent Memory
- PMR SSD mode of operation
- Key Benefits with PMR SSD
- Use Cases
- Next Steps



Flash Memory Summit

What is Persistent Memory

Key Attributes

- Low Latency
- Data Power Loss Protected
- Utilized for In Memory Applications acceleration
 - Cassandra, MongoDB®, STORM, KAFKA, SPARK ...
- Byte Addressable through CPU Load/Store Memory Instructions
- Block Addressable through software changes
- Today Served by
 - DIMM's Battery backed
 - NVDIMM's with Flash Storage

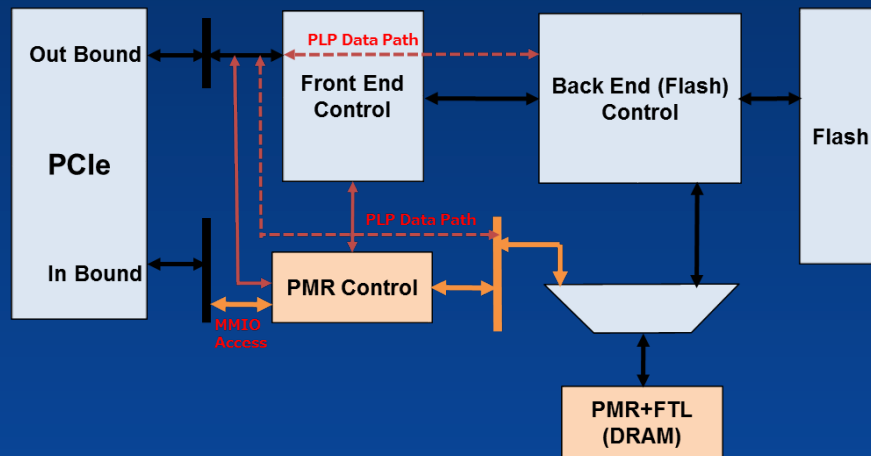
Wikipedia Definition...

In computer science persistent memory is any method or apparatus for efficiently storing data structures such that they can continue to be accessed using memory instructions or memory APIs even after the end of the process that created or last modified them.¹



NVMe™ SSD with PMR : Concept

Functional Blocks for NVMe PMR SSD



Key Blocks for PMR

- ❑ NVMe™ Enterprise SSD
- ❑ Extra DRAM to support PMR
- ❑ Data loss protection
- ❑ Configurable PMR size

Single Device offering for both block storage and PMR needs



Flash Memory Summit

PMR Mode of Operation

- Memory Mapped PMR after enumeration
 - Driver reads capability register and allocates Persistent Memory to Host (application)
- Accessibility through PCIe® bus
 - MMIO Mode for Byte Access
- Writes and Reads Transactions:
 - Writes are “posted writes” based on PCIe “no ACK”
 - Reads are end to end from PMR to Host CPU
- In case of power loss, PMR Data gets saved to Flash
- PMR Data gets restored from Flash on next power up



Flash Memory Summit

Key benefits of SSD-based PMR

- ❑ Single Device with Persistent Byte Memory and Block storage
- ❑ Saves DIMM slots
- ❑ Dual port accessibility for higher reliability
- ❑ Aggregation of PMR's from multiple drives
- ❑ Robust and mature PCIe interface
 - Standard platform
 - Solid debug platform
 - Tools, analyzer fully available



Flash Memory Summit

PMR SSD Use Cases

- Log for software RAID & erasure coding systems
- Commit log device for NOSQL databases as well as Relational (MySQL, etc.) databases
- Journal for file systems
- Buffer for write-coalescing in caching systems
- Metadata
- Staging for de-dupe, compression, etc.
- RDMA transactions



Flash Memory Summit

Thoughts on Next Steps....

Next steps ...

- Effort to standardize PMR
 - Registers definitions for PMR settings
 - Get/Set Features for PMR configuration
 - PMR as Namespace unit for security (Lock/Unlock)
 - Data units boundaries for moving data between PMR and Flash
- Programming Model API for accessing PMR



Flash Memory Summit

Backup

Santa Clara, CA
August 2017



Evolution of Persistent Memory


Emergence of Persistent Memory Options (source SNIA)

NVDIMMS - JEDEC TAXONOMY

NVDIMM-N

Standardized


- Memory mapped DRAM. Flash is not system mapped
- Access Methods -> byte- or block-oriented access to DRAM
- Capacity = DRAM DIMM (1's -10's GB)
- Latency = DRAM (10's of nanoseconds)
- Energy source for backup
- DIMM interface (HW & SW) defined by JEDEC



NVDIMM-F

Vendor Specific

- Memory mapped Flash. DRAM is not system mapped.
- Access Method -> block-oriented access to NAND through a shared command buffer (i.e. a mounted drive)
- Capacity = NAND (100's GB-1's TB)
- Latency = NAND (10's of microseconds)



NVDIMM-P

Proposals in progress

- Memory-mapped Flash and memory-mapped DRAM
- Two access mechanisms: persistent DRAM (-N) and block-oriented drive access (-F)
- Capacity = NVM (100's GB-1's TB)
- Latency = NVM (100's of nanoseconds)

