Capabilities and System Benefits Enabled by NVDIMM-N

Bob Frey

Arthur Sainio
SMART Modular Technologies

August 7, 2018

NVDIMM-N Maturity and Evolution

- If there's one takeaway you should have from this session it is that the NVDIMM-N is a mature persistent memory solution. You could call it the grandfather of “modern” persistent memory solutions, but it continues to evolve.

NVDIMM-N Basic Components

- ***Plugs into JEDEC Standard DDR4 DIMM Socket***
- ***Appears as standard DDR4 RDIMM to host during normal operation***
- ***Energy Source may be Device Managed (tethered) or Host Managed (supplied through 12V DDR4 Power Rail)***

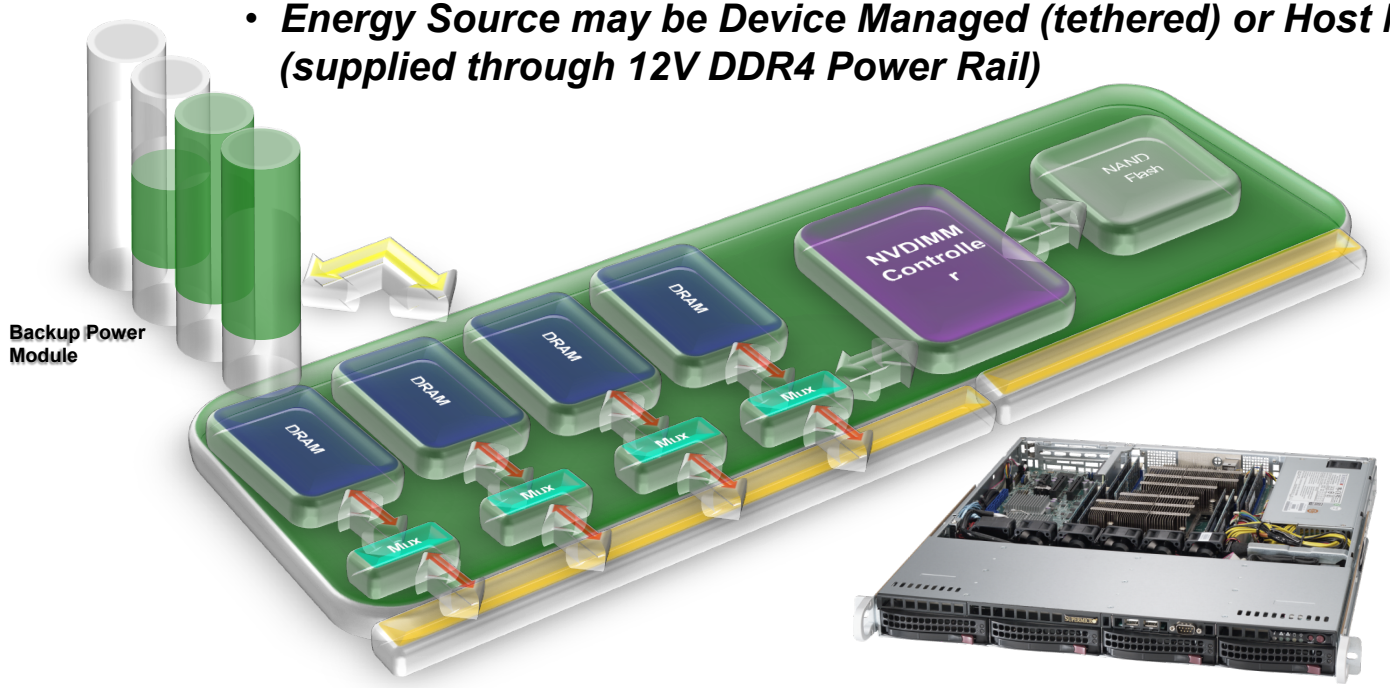

Adapted from SNIA presentations by AgigA Tech

NVDIMM-N Ecosystem

- **Standards: JEDEC, SNIA, and UEFI**
- **JEDEC Naming: DRAM backed by -N (NAND).**
- **JESD245 I2C Command & Status interface.**
- **Intel DDR4 Skylake MRC JEDEC JESD245 Support**
- **Windows Server 2016 and Linux 4.4+ kernels (/dev/pmem) include DAX support.**
- **Open source library for applications Load/Store Access (PMEM.IO)**

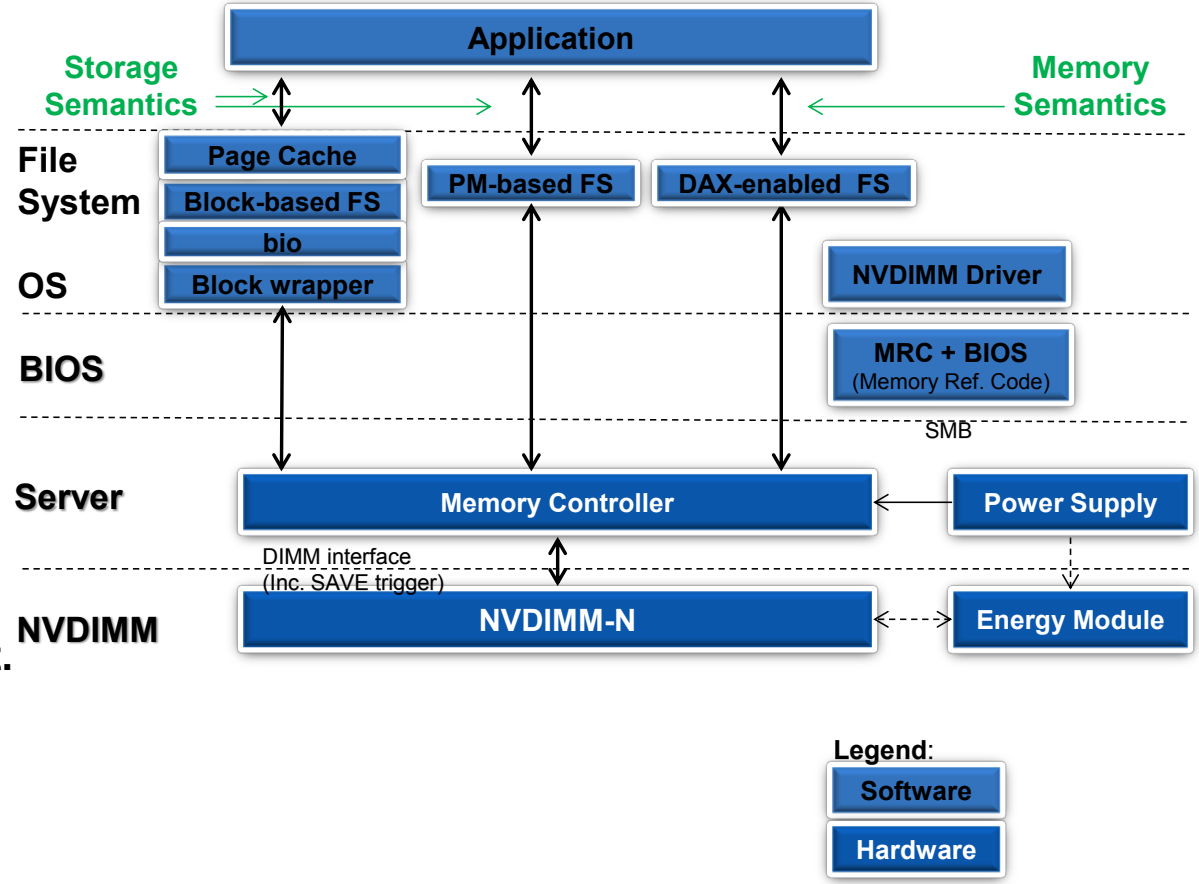

Platform Hardware Requirement

- DDR4 12V Power Pins (1, 145) standardized
- DDR4 SAVE_n Pin (230) standardized
  - Bi-directional SAVE_n to indicate SAVE completion
- EVENT_n asynchronous event notification
- I2C Device Addressing
- 12V in DDR4 simplifies NVDIMM power circuitry and cable routing
  - One cable needed between NVDIMM and BPM (Backup Power Module)
  - No cable needed if Host provides 12V backup power via DDR4 12V (Host Managed)

Persistent Memory Standards

- **JEDEC JESD 245A, 245B: Byte Addressable Energy Backed Interface**
  - Defines the host to device interface and features supported for a NVDIMM-N
- **ACPI 6.2**
  - NVDIMM Firmware Interface Table (NFIT)
  - NVM Root and NVDIMM objects in ACPI namespace
  - Uncorrectable memory error handling
  - Notification mechanism for NVDIMM health events and runtime detected uncorrectable memory error

NVDIMM-N Basic DRAM/NAND Backup/Restore Functionality

- Host Catastrophic failure or Reboot occurs, places DRAM's in Self Refresh, and asserts SAVE_n signal or NVDIMM controller is triggered and autonomously puts DRAMs into Self-Refresh.
- CPU L1/L2/L3 caches typically not flushed! Posted writes, Read-back, Serialization.
- NVDIMM-N switches power rails
- NVDIMM-N switches all ADD/CMD/CTRL & DATA mux's
- NVDIMM-N reads DRAM and writes NAND flash.
- Next power cycle checks for backup image and initiates data restore
- NVDIMM-N switches control back after restore and NVDIMM is armed.

NVDIMM-N Back-up Scenarios and System Implementation

- Back-up on SAVE_n, RESET_n, CKE
- Back-up on Reset/Reboot/IPMI power command
- Back-up on Power Failure
- Back-up on HW Fault (THERMTRIP, CATERR)
- Back-up on OS Hang/Watchdog Timeout
- DRAMs in self-refresh? Data integrity maintained?

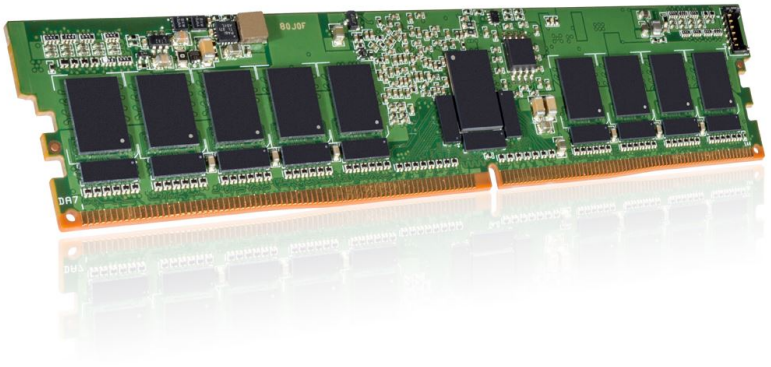

NVDIMM-N /dev/pmem vs. SCM SSD Random and Sequential Writes

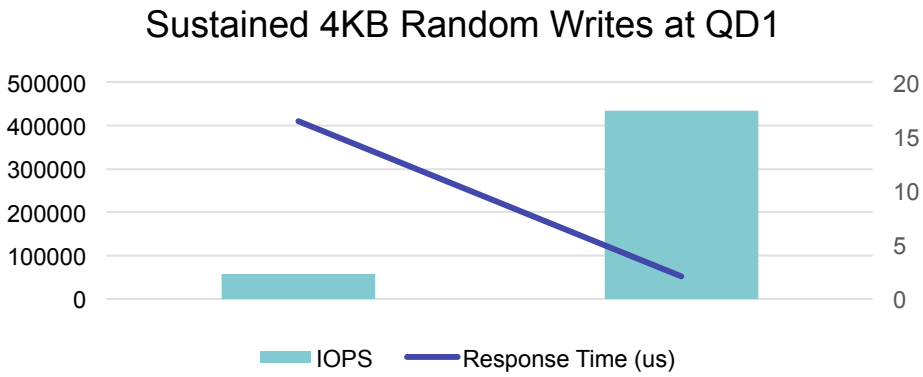

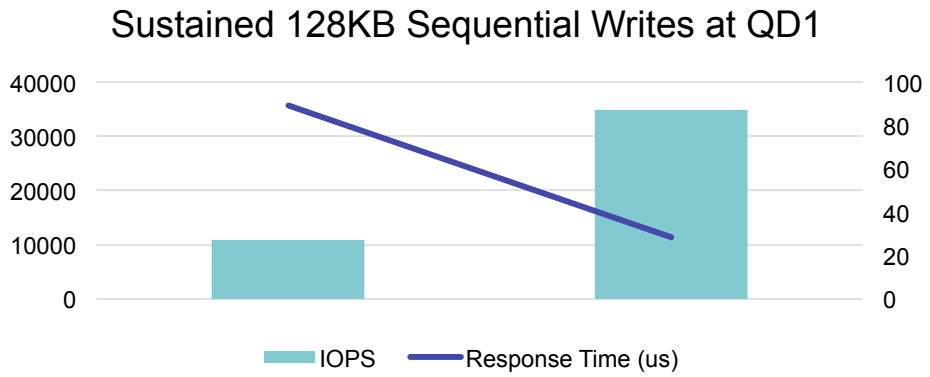

Sustained Random Writes	SCM SSD	NVDIMM
IOPS	57,526	434,290
Response Time/ Latency (µs)	16.43	3.09

Sustained Sequential Writes	SCM SSD	NVDIMM
IOPS	10.981	34,926
Response Time/ Latency (µs)	89.37	28.45

- Test Setup: Supermicro X11DRi, ezfio benchmark tool, 480GB SCM SSD, 16GB NVDIMM

NVDIMM Use Cases

Flash Memory Summit

- **Enterprise Storage**
  - Tiering, caching, write buffering and meta data storage
- **Traditional Database**
  - Log acceleration by write combining and caching

- **In Memory Database**
  - Journaling, reduced recovery time, tables
- **Server Data Continuity**
  - Retain working set after OS hang or crash

NVDIMM-N Future Outlook

- DDR Technology Speed Increases
  - Intel: Sky Lake DDR4-2666, Cascade Lake: DDR4-2933, Ice Lake: DDR4-3200
  - AMD: Epyc DDR4-2666, DDR4-3200
  - Power9: DDR4-2666
- ASR (Autonomous Self-Refresh)
- Encryption: SNIA, JEDEC
- Replacement of (-N) NAND Flash with alternative PM.
- Gen-Z, OpenCAPI
  - NVDIMM-N mapping

NVDIMM-N Summary

- Today there are multiple NVDIMM-N vendors with qualified customers using NVDIMM-N in shipping products. If you visit all of the storage appliance vendors exhibiting here at FMS you will find that nearly all of them have adopted NVDIMM-N modules to improve the performance of their solutions.
- New use cases are being defined and supported, replacements for PM being evaluated, and mapping to new server architectures underway.
- As long as DRAM devices maintain a performance advantage over other solid-state storage devices, then there will be a demand for NVDIMM-N type solutions which provide non-volatile backing storage for the volatile DRAM.

Thank You!