NVDIMM/NVRAM
Next Generation

Mario Martinez
Netlist

NVM is a Necessity but Transforming

- BBU (Battery-Backed Unit) RAID, SLC SSD, & NVDIMM cache on storage array controllers
- NVRAM (Hybrid)
  - PCIe NVRAM Card: Non-volatile access and stores at DRAM speeds, backs up into NAND only on a power loss, Supercap based, eliminates battery. DMA with block and character devices.
  - PCIe NVRAM SSD: NVRAM + Storage : Predictable and Sustainable Low Latency SSD
- NVDIMM (Hybrid)
  - NVDIMM-N Stores in DRAM, backs up into NAND only on a power loss, Supercap based, eliminates battery
  - NVDIMM-F – maps NAND into memory address space
  - NVDIMM-P - maps NAND and DRAM into memory address space
  - And Beyond! FPGA on NVDIMM: Accelerate dynamically changing workloads

NVM Fit in the Storage Hierarchy

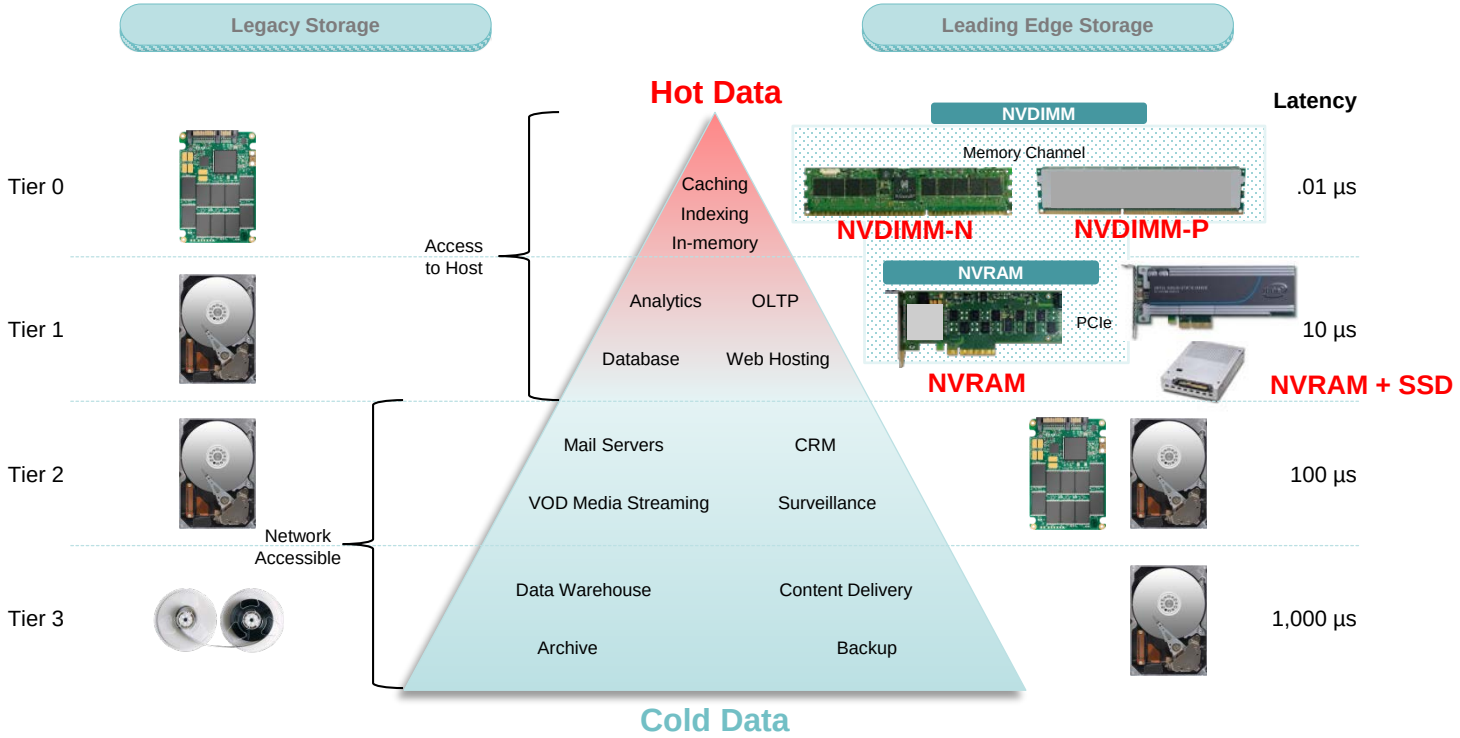

NVM Adoption by System Architecture

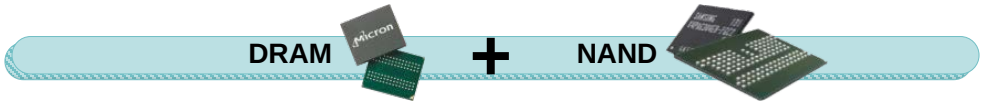

NVM - NVRAM

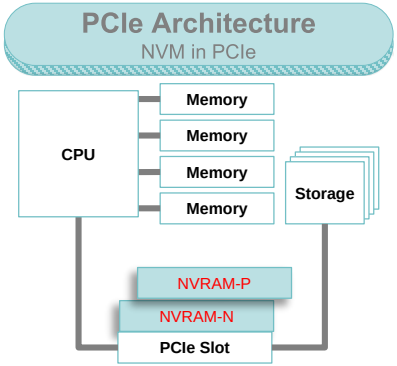

NVM - NVDIMM

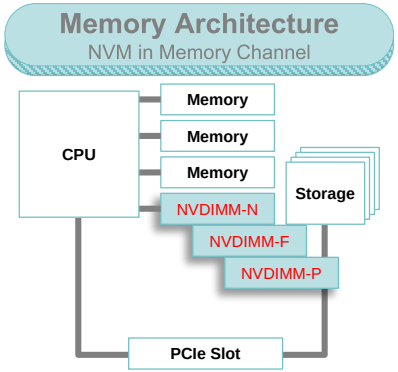

openstack

OpenPOWER

OPEN Compute Project

Microsoft *

*

SUPERMICRO

nimble storage

hp

intel

OPEN Compute Project

SNIA

Rack Scale Architecture

* OCP Specification

Applications that Drive the Need for Faster Storage and/or Disaster Recovery

Write Cache / MetaData	In-Memory Compute	Low Latency Storage	Data Center / Cloud
NVRAM	NVDIMM-P NVDIMM-N	NVDIMM-N	NVDIMM-N NVRAM

NVDIMM-N Next Gen

Type	Features	1st Gen Proprietary	2nd Gen/3rd Gen JEDEC	Change
NVDIMM/ Firmware Hardware	NV controller registers controlled by Host via i2c	Yes	Yes	Same
	DDR4 12V Power Pins (1,145)	Yes	Yes	Same
	DDR4 SAVE_n Pin (230)	Yes	Yes	Same
	NV Controller EVENT# Pin (78)	Yes	Yes	Same
	SPD for NVDIMM representation	In Part number	JEDEC SPD	New BIOS
	NV Controller registers	DDR3 compatible	JEDEC Registers	New Driver, New BIOS
	Memory Interface to Host	RDIMM	2nd - RDIMM 3rd - LRDIMM	Add LRD
	JEDEC Raw Cards	None	2nd - RDIMM 3rd - LRDIMM	2nd - None 3rd - New
	Capacity	8/16GB	2nd - 8/16GB 3rd – 32/64GB	Increased

NVDIMM-N Next Gen

Type	Features	1st Gen - Proprietary	2nd Gen/3rd Gen - JEDEC	Change
System/ OS/ BIOS/ MRC	OS Driver (Block&Load/Store) - Block w/b first	• DDR3 compatible	• New ACPI 6.0 and PMEM library compatible • Hardware Agnostic	New Driver
	NVDIMM Aware Kernel (Direct Access support)	• Intel patch for 3.14 • No support for JEDEC	• 3.20 or higher – • Hardware Agnostic	New Kernel
	Intel MRC Changes to support NV Vendor	• Yes - uses DDR3 MRC on Haswell	• New MRC is required • Hardware Agnostic	New MRC
	BIOS to support NV Vendor	• Yes - Insyde/AMI support Intel MRC	• New BIOS is required • Hardware Agnostic	New BIOS
	Direct Access (DAX) support for NVDIMM-N modules in Ext4	• No	• Yes - eliminates the page cache layer completely. • Hardware Agnostic	New Driver
	NVDIMM aware ACPI	• <v6.0 (No support)	• 6.0 or higher	New Driver
	12V support to connector - Input	• Via Auxiliary	• Yes	Modify HW
	12V support Type	• Source Supercap	• Source Supercap • Backup operation	Compatible
	ADR support	• Yes	• Yes	Compatible
	EVENT support – Output	• Yes	• Yes	Compatible
	SAVE_n support - Input	• Yes	• Yes	Compatible

NVRAM & NVDIMM Next Gen

	NVRAM	NVDIMM-N	NVDIMM-F	NVDIMM-P
Power Consumption	• Active Read: >10W • Idle: <4W	• Active Read: <6W • Idle: <1W	• Active Read: >10W • Idle: <4W	• Active Read: >10W • Idle: <4W
Form Factor /Interface	• PCIe HH/HL Card • PCIe 2.5in Drive • NVMe	• DDR4 RDIMM & LRDIMM JEDEC	• DDR4 RDIMM JEDEC	• DDR4 RDIMM/LRDIMM JEDEC • Custom • PCIe
Performance	• >500K IOPs RandRead • >500K IOPs RandWrite	• >3M IOPs RandRead • >3M IOPs RandWrite	• >140K IOPs RandRead • >100K IOPs RandWrite	• > 3M IOPs – Mode N • > 140K IOPs – Model F
Latency	• <15us	• <1us	• <100us	• <10us – Mode N • <100us – Mode F
Capacity **dWPD** **System Fit** **Configuration**	• +32GB • Unlimited Writes • 1U Server Support • PCIe Scale Out	• +32GB • Unlimited Writes • 1U Server Support • Follows RDIMM Population Rules	• +1TB • <10 WPD for 5yr • 1U Server Support • Restricted Population Rules	• +1TB • >100 WPD for 5yr • 1U Server Support • Restricted Population • Follow Intel Population Rules
MB Support	• Agnostic	• Limited	• Limited	• Limited & Agnostic
CPU Usage	• <5%	• >50%	• >50%	• >50%
Addressable	• DRAM	• DRAM	• Flash	• DRAM & Flash

Next Steps

- Continue to Enable the Industry to support NVRAM and NVDIMM by working through standards bodies:
  - NVM express
  - SNIA NVM Programming TWG
  - UEFI
  - JEDEC
  - OpenPower
- Enable Applications and Operating Systems ISV
  - OpenStack
  - VMWare
  - Linux

Thank You!