



NVDIMM Messaging and FAQ



Contributor:

AgigA Tech

January 2014

About the SNIA

The Storage Networking Industry Association (SNIA) is a not-for-profit global organization, made up of some 400 member companies spanning virtually the entire storage industry. SNIA's mission is to lead the storage industry worldwide in developing and promoting standards, technologies, and educational services to empower organizations in the management of information. To this end, the SNIA is uniquely committed to delivering standards, education, and services that will propel open storage networking solutions into the broader market. For additional information, visit the SNIA web site at <http://www.snia.org> or <http://www.snia.org/international> for a list of International Affiliates

About the SNIA Solid State Storage Initiative

The SNIA Solid State Storage Initiative (SSSI) fosters the growth and success of the market for solid state storage in both enterprise and client environments. Members of the SSSI work together to promote the development of technical standards and tools, educate the IT communities about solid state storage, perform market outreach that highlights the virtues of solid state storage, and collaborate with other industry associations on solid state storage technical work. SSSI member companies represent a variety of segments in the IT industry (see <http://www.snia.org/forums/sssi/about/members>). For more information on SNIA's Solid State Storage activities, visit www.snia.org/forums/sssi and get involved in the conversation at <http://twitter.com/SNIASolidState>.

About AgigA Tech

Headquartered in San Diego, California, AgigA Tech is a leading developer of high-speed, high-density, battery-free non-volatile memory solutions. The company's flagship AGIGARAM® product family represents a new class of non-volatile memory created to address the fundamental need for higher-density, higher-performance memory in enterprise-class applications. AGIGARAM integrates NAND Flash, DRAM and an ultracapacitor power source into an innovative, highly-reliable non-volatile memory subsystem that can deliver unlimited read/write performance at the fastest DRAM speeds, while also safely backing up all data when power is interrupted. AgigA Tech is a subsidiary of Cypress Semiconductor Corp. The company's products are available worldwide through the Cypress sales force and its distributor partners. More information on the company is available at www.agigatech.com.



Table of Contents

Executive Summary.....	1
Key Messages about NVDIMM Technology	1
Key Messages about the NVDIMM Marketplace	1
About NVDIMM FAQs	2
Ecosystem FAQs	4
System Integration FAQs.....	5
Applications FAQs.....	6

Executive Summary

This document provides a summary of key messages on NVDIMM technology and the marketplace, and provides information on frequently asked questions about NVDIMM.

Key Messages about NVDIMM Technology

- NVDIMMs are a storage class memory module
- NVDIMMs provide the speed/latency/endurance of DRAM with the persistence of flash
- NVDIMMs operate as a standard JEDEC DIMM
- NVDIMMs provide a more reliable alternative to batteries
- NVDIMMs are a complementary memory tier to flash/SSDs
- NVDIMMs provide a separate failure domain for improved system reliability
- NVDIMMs enable a lower total cost of ownership (TCO)
- NVDIMMs provide improved scalability for datacenters

Key Messages about the NVDIMM Marketplace

- NVDIMMs are shipping today
- NVDIMMs are offered by multiple vendors
- NVDIMM standardization is in-progress
- Intel will support NVDIMMs on Grantley (and limited support on Romley)

About NVDIMM FAQs

What is an NVDIMM?

A Non-Volatile DIMM (NVDIMM) is a dual in-line memory module that operates as standard DRAM while also having the persistence of Flash. A typical NVDIMM module incorporates DRAM, Flash, control logic and an independent power source to retain critical in-memory data through unexpected power loss events, system crashes or planned shutdowns.

How does an NVDIMM work?

During normal operation, the NVDIMM appears to the host system as a standard JEDEC¹ DRAM memory module, providing the speed, latency and endurance benefits of DRAM. In the event of an unexpected power loss or system crash, the critical data residing in the DRAM is saved to onboard Flash. When power is returned, the in-memory state of the DRAM is restored from Flash. For most typical NVDIMM solutions, the backup power to transfer the DRAM contents to Flash is provided by supercapacitors.

Does an NVDIMM work like a standard JEDEC DIMM?

Yes, an NVDIMM plugs into a standard JEDEC DIMM socket and operates as a standard DRAM DIMM during normal runtime. The difference of course is that an NVDIMM will retain its data through a power loss event or system crash.

Are NVDIMMs the same mechanical size as a standard JEDEC DIMM?

Yes, most NVDIMM solutions meet the mechanical dimensions defined by JEDEC. Please refer to specific vendor NVDIMM specifications for exact dimensions that are offered.

Are there any technical limitations of an NVDIMM?

Not really. Since an NVDIMM uses standard off-the-shelf DRAM and Flash devices as its underlying technology, the only technical limitations are related to the memory technologies themselves. There may of course be practical limitations, such as the size of the supercapacitors required to backup large density NVDIMM.

¹ JEDEC stands for the JEDEC Solid State Technology Association, formerly known as the Joint Electron Device Engineering Council (JEDEC), an independent semiconductor engineering trade organization and standardization body

What are the benefits of an NVDIMM?

NVDIMMs provide the low latency and nearly infinite endurance of DRAM, along with the non-volatility of Flash. This allows main system memory to be persistent in the event of unexpected power failure or system crash. By having this persistence, applications can run at far higher speeds - I/O performance in a host of applications, including storage and database acceleration are dramatically improved.

Additionally, there is a significant value to applications that are sensitive to down-time. In the event of power failure without NVDIMMs, a typical environment would hold up the entire processor complex using a UPS or generator long enough to transfer the critical data safely to a non-volatile location (e.g. an SSD). This can take many minutes to complete, and during a time that makes it difficult to control the environment when faced with a catastrophic failure. On a system recovery, the subsequent “rebuild time” can take hours. With NVDIMMs the data is saved in seconds, and upon power being restored, as soon as the system reboots the data is immediately available.

What are the applications for an NVDIMM?

Any application that can benefit from the latency and endurance of a persistent DRAM solution would be an ideal target. Some examples include:

- Write Caching
- Metadata Storage
- Tiered Storage
- In-Memory Database
- Big Data Analytics
- Message Queuing
- On-Line Transaction Processing
- Whole System Persistence
- UPS Replacement

What kind of performance can I expect from an NVDIMM?

During normal runtime, an NVDIMM will provide the equivalent performance of standard DRAM.

Are NVDIMMs available today?

Yes, NVDIMMs are currently offered by several vendors

NVDIMM Messaging and FAQ



Is there an NVDIMM standard?

Not yet, but this effort is on-going, with a focus on DDR4. There are several industry efforts to enable NVDIMMs as a standard technology, so that they can simply be plug-and-play in next generation platforms. Below is a summary of the activity to date:

- JEDEC – developing standards around the NVDIMM system interface (e.g. hardware, SPD, common command set)
- Intel – providing platform and BIOS support
- SNIA – defining programming models for persistent memory usage
- NVDIMM Special Interest Group – technology promotion and central clearinghouse for above efforts

Why use supercaps instead of batteries?

Supercaps in general are more reliable and perhaps more importantly wear out more predictably than batteries. Batteries of course are a less expensive solution, but NVDIMMs can save more over time due to the lower maintenance requirements. Most NVDIMM solutions do not preclude the option to use a battery for the backup power source, but the benefits of a supercap solution for applications that require data reliability make supercaps the preferred choice.

Ecosystem FAQs

Are NVDIMMs Plug-and-Play?

Not quite yet. There are a number of standardization efforts (see above) that are on-going that target this as the ultimate goal. The expectation is that Intel Haswell platforms will support NVDIMMs plug-and-play for DDR4.

Does Intel provide support for NVDIMMs?

Yes. For Haswell, NVDIMMs will be supported. There is also limited support for Sandy Bridge and Ivy Bridge platforms.

Who are the vendors that support NVDIMMs?

The vendors who are currently publicly promoting NVDIMM solutions are:

- AgigA Tech
- Micron
- Netlist
- Smart Modular
- Viking Technology
-

It is expected that several more vendors will enter this space in the near future.

System Integration FAQs

What is required to integrate an NVDIMM into my application?

NVDIMMs rely on the host to:

1. Configure NVDIMMs
2. Detect AC Power Loss
3. Flush Write Buffers (or Processor Cache)
4. Place DRAM In Self-Refresh
5. Trigger BACKUP (or SAVE)
6. RESTORE Data On Boot

What do the integration requirements mean from a hardware and software perspective?

Below is a summary of the system modifications that would be required:

- BIOS
 - Detect and configure NVDIMMs (check SPD)
 - NVDIMMs are marked as a specific type (12) in E820 memory table
 - RESTORE data if SAVED image is present
 - (Optional) ARM NVDIMM for next Backup
- Hardware
 - AC power loss detection – signal routed from PSU
 - PSU hold-up requirements
 - <100us for ADR, ~10ms for processor caching
 - ADR connections
 - SAVE# pin routing – pin 167 for DDR3, pin 230 for DDR4
 - Supercap mechanical placement
 - (DDR4 only) 12V power to NVDIMM
- Software
 - SMB driver – for status monitoring
 - System Memory Map
 - Data Access: Byte or Block addressable
 - Application Interface: Health Monitor, ARM/DISARM, Status

What is ADR?

ADR stands for Asynchronous DRAM Refresh. ADR is a feature supported on Intel chipsets that triggers a hardware interrupt to the memory controller which will flush the write-protected data buffers and place the DRAM in self-refresh. This process is critical during a power loss event or system crash to ensure the data is in a “safe” state when the NVDIMM takes control of the DRAM to backup to Flash. Note that ADR does not flush the processor cache. In order to do so, an NMI routine would need to be executed prior to ADR.

NVDIMM Messaging and FAQ



What BIOS support is available?

Intel has modified the MRC (Memory Reference Code) in both Romley and Grantley to support NVDIMMs. Several NVDIMM vendors also provide BIOS integration support as needed.

How does the BIOS configure the NVDIMM(s)?

The MRC will detect an NVDIMM through an SPD identifier and reserve that memory address space as a specific type in the E820 memory table. For DDR3, it is typically reserved as type 12.

How does an NVDIMM save the data?

Upon system power loss, an ADR sequence will be triggered. ADR will flush the write-protected data buffers in the memory controller and place the DRAM in self-refresh. When ADR is complete, the SAVE# pin is triggered (side note: other triggers can also be supported, but SAVE# is the typical implementation). The NVDIMM will then take control of the DRAM and move the contents to the on-board Flash. When the transfer is complete, the NVDIMM goes into a zero power state. The data can be retained for the spec duration of the flash, which is typically many years.

How does an NVDIMM restore the data?

When power is restored, the BIOS will check whether a valid image is present (if a SAVE is in progress, it will also wait for it to complete). If there is a valid image, the BIOS will restore the data from NAND back into DRAM. When the memory is handed back off to the host, the in-memory state prior to the power-loss event will be present. This is similar to a SUSPEND/RESUME sequence of a PC.

Can NVDIMMs and standard DIMMs operate in the same system?

Yes, but not in the same channel. If interleaving is required, the same NVDIMM configuration must be used across the channels that are interleaved. The standard DIMM population rules also apply to NVDIMMs.

Where does the ultracap module go?

Wherever it can fit in the system. Most NVDIMM vendors offer standard form factors intended to fit into common chassis, such as a 2.5" drive bay.

Applications FAQs

How does an application utilize an NVDIMM?

An application can use the NVDIMM address space as byte or block addressable. The NVDIMM itself is agnostic to the use case. Since the address space is allocated as a reserved memory type, an application can target write transactions to the NVDIMMs.

NVDIMM Messaging and FAQ



Can I use an NVDIMM as a RAMdisk?

Yes, an NVDIMM can be set up as a RAMdisk. However, file system overhead will likely have a performance impact.

How do NVDIMMs compare to SSDs?

Comparing NVDIMMs to SSDs is somewhat apples and oranges. In a head-to-head matchup, NVDIMMs will have much lower latency and much higher endurance than an SSD. An SSD on the other hand can support much higher densities. It is more practical to consider NVDIMMs and SSDs as complementary technologies in a tiered storage strategy. NVDIMMs in fact can improve the reliability and endurance of SSDs by lowering the access frequency of the NAND. “Superhot” data can be targeted to the NVDIMMs, while the “hot” data can be held in the SSD.

How does an NVDIMM compare to an HDIMM or UltraDIMM?

Also apples and oranges as explained above. However, one should acknowledge that these solutions do compete for the same real estate. So it really depends on the application requirements as to which one is the better suited solution.

Does an NVDIMM compete with other non-volatile technologies like PCM, MRAM, STT, etc...?

Possibly in the future, but these emerging technologies have not yet offered a viable DRAM replacement. Depending on who you ask, they are several years away, and in fact may never get there as DRAM technology continues to widen the gap in terms of performance.