NVDIMM Technical Brief

Solving Data Volatility

Contributor:

Smart Modular Technologies

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, visitwww.snia.org/forums/sssi and get involved in the conversation at http://twitter.com/SNIASolidState.

About Smart Modular Technologies

SMART Modular is a technology leader in the design, development, and deployment of current and next-generation memory and storage products. Combining leading-edge design and manufacturing with proven world-class support, SMART Modular leverages its deep expertise in DRAM, SRAM and Flash architectures with world-class board-layout design. SMART Modular delivers high-quality, high-reliability solutions to a broad customer base, including tier one computing and telecom OEMs as well as industrial, medical and automotive customers. In support of these designs, SMART Modular provides award-winning innovations in memory and embedded storage technologies. SMART Modular is part of the SMART family of global companies. For more information, visit http://www.smartm.com/index.asp

Table of Contents

Executive Summary

Non-Volatile Dual In-line Memory Modules (NVDIMMs) merge two leading technologies: DDR3 DRAM[1] and NAND Flash[2], to solve the problem of data volatility in mission critical servers. Let's take a closer look, and explain the strengths and weaknesses of each.

DDR3 DRAM vs. NAND Flash

The prevailing memory technology, DDR3 DRAM, and the leading storage technology, NAND Flash, each has its own set of attributes that makes it perfectly suited for its job.

Speed: DDR3 Wins by 20X

NAND is relatively fast compared to its storage rival the hard disk drive, yet DDR3 is still about twenty times faster than NAND. A single DDR3-1600 DIMM has 12.8GB/sec bandwidth, seventeen times faster than the SATA III maximum of 750MB/sec. Random read and write speeds for a single DDR3 DIMM are approximately two million IOPS (Input/output operations per second), while SATA SSDs are in the neighborhood of 50,000 4KB IOPS. At the system level, the numbers compare about the same.

Typical servers have two CPU sockets, each with three or four DDR3 channels, delivering a total memory bandwidth in a single server of up to 102GB/sec. Eight SATA III SSDs in a RAID 0 configuration reach a theoretical data transfer limit of 6GB/sec, a mere 1/19th that of the DDR3 system. DDR3 is so much faster than NAND, that many SSDs use DDR memory as a cache to accelerate write performance.

Endurance: DDR3 Wins by 1000X

DRAM has practically unlimited endurance in terms of write cycles because the process of writing data to DRAM does not degrade the memory. Program/Erase (PE) cycles on NAND flash, however, slowly degrade components, which is a major constraint of the NAND technology. MLC NAND supports only about 5000 PE cycles per block. Of course, there are other NAND alternatives available such as eMLC or SLC that support significantly more PE cycles, but those come at a much higher price, and still cannot compete with the endurance of DRAM.

Capacity: NAND Wins by 16X

In this age of virtualized servers, no amount of memory or storage capacity seems to be enough. The highest capacity DDR3 DIMM in production today is only 32GB, but 64GB DIMMs are not far away. Enterprise class 2.5" SATA SSDs are readily available in 500GB, and 800GB and 2TB models can be found.

---

[1] an abbreviation for double data rate type three synchronous dynamic random access memory
[2] a type of Flash, a non-volatile read/write semiconductor memory which is used in solid state storage devices

2014 STORAGE NETWORKING INDUSTRY ASSOCIATION

Volatility: NAND Wins by Default

DRAM loses all of its data instantly when power is removed, so DRAM is referred to as volatile memory. Even a brief unexpected power outage will wipe out all data stored in memory. In contrast, NAND flash is non-volatile, meaning that it retains its data when power is removed. In the event of an unexpected power outage, the only data at risk in an SSD is data that is still in the DRAM cache and has not yet been written to flash.

Cost: NAND Wins by 10X

MLC NAND costs about $1 per gigabyte, while DDR3 DRAM costs roughly $10 per gigabyte. These prices fluctuate heavily, but today MLC NAND is about ten times cheaper per gigabyte than DDR3, and that gap will likely increase with time. In summary, DDR3 memory is far superior in speed and durability, while NAND flash wins in capacity, non-volatility and price

Table 1, DDR2 and NAND Flash Comparison

Speed	DDR3	NAND Flash
Sequential (MB/sec)	**12,800**	750
Random (4K IOPS)	**2,000,000**	50,000
Endurance (P/E cycles)	**5,000,000**	5,000
Max Capacity (GB)	32	**512**
Volatile/Non-Volatile	Volatile	**Non–Volatile**
Cost/GB ($)	10	**1**

The Volatility Problem

If memory is volatile and flash is non-volatile, it would seem that servers and applications would be designed to avoid placing critical data in memory. Application developers attempt to do just that, developing robust applications that are not susceptible to inadvertent data loss. Moreover, many modern applications are written to take full advantage of the attributes of both DDR3 memory and NAND based SSDs to deliver the highest possible performance and reliability.

However, it is inevitable that some critical data will be in memory for some applications, and will be exposed to the risk of sudden power failure. For example, in applications such as high speed online transaction processing (OLTP), unexpected data loss could result in the failure of financial transactions, leaving the data unrecoverable or corrupted.

How NVDIMMs Work

During normal operation, an NVDIMM functions as a regular DDR3 RDIMM. It plugs into a standard DDR3 DIMM socket and has JEDEC® standard SPD data. Every NVDIMM has its own on-DIMM DDR3 controller that, when the NVDIMM senses a power failure, immediately copies all data from DRAM into flash. When power is restored, the NVDIMM copies the data from flash back into DRAM and normal operation resumes with no data loss.

The NVDIMM uses the HW SAVE# signal to trigger a backup, at which point the NVDIMM's own DDR3 controller takes over the DDR3 interface from the host and dumps all DRAM data to flash through either side-band signals or the I2C interface. At a backup speed of 128MB/sec, it takes about 10 seconds per 1GB to either backup or restore the data.

NVDIMMs require support from the system motherboard. When plugged in, the BIOS must recognize the NVDIMMs. OEM[3]s who control the BIOS and MRC (memory reference code) can make the necessary code changes to implement NVDIMMs into their servers. The DRAMs must be in self-refresh before a power failure to insure the cache data is backed up properly. NVDIMMs from vendors such as Smart Modular Technologies also feature a USB interface that allows OEMs direct access to an NVDIMM backup image; this can bypass the need for BIOS changes.

Supercaps or Batteries?

NVDIMMs require a stored charge so they can back up their data to flash when power fails. Either super capacitors (supercaps) or battery packs can be used to support these NVDIMMs. Supercap modules provide the backup power for the DRAM on NVDIMMs. Attached to the NVDIMM via an included tether cable, they are designed to store enough charge to keep the DRAM powered up long enough to copy all of its data to flash. As an example, supercap modules from Smart Modular Technologies are available in three different form factors to accommodate different system requirements. These modules include health monitoring features for checking voltage, current and temperature values, and perform ultra capacitor cell voltage balancing and automatic periodic capacitance and ESR self-test.

How It Works

If there is a power failure, the supercap module powers NVDIMM while it copies all data from the DDR-3 to on-module flash

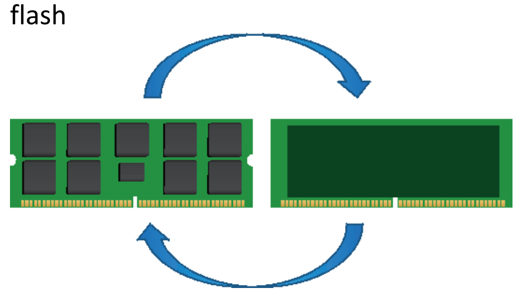

When power is restored NVDIMM copies all data from flash to DDR-3 and normal operation resumes

DDR4 NVDIMMs

Standards for DDR4 NVDIMMs are being developed now, and are gaining industry-wide support. DDR4 NVDIMMs are expected to have some feature enhancements over DDR3, as well as faster performance and lower power consumption. Since NVDIMM adoption is happening simultaneously with the DDR4 production ramp, most of the NVDIMM volume is expected to be DDR4 based.

[3] Original Equipment Manufacturer

2014 STORAGE NETWORKING INDUSTRY ASSOCIATION