

SMART Technical Brief: NVDIMM

With SMART's SafeStor™ Technology

NVDIMMs merge two leading technologies, DDR3 DRAM and NAND Flash, 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 attains approximately two million IOPS (Input/output operations per second), while SATA SSD are in the neighborhood of 50,000 4KB IOPS. At the system level, the numbers are no different. 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 does not compare to 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.

DDR3 versus NAND Flash

	DDR3	NAND Flash
Speed		
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
Vol/Non-Vol	Volatile	Non-Volatile
Cost/GB (\$)	10	1



Applications

- OLTP (online transaction processing)
- Nonvolatile cache backup for RAID storage systems
- Data logging and analytics
- De-duplication

Benefits

- Eliminates need for batteries
- Eliminates need to maintain a duplicate cache memory subsystem separate from main memory
- Simplifies memory system design
- Allows for increased cache backup size
- Saves costs
- Saves space

Features and Specs

- Fits standard 240-pin RDIMM form factor
- Functions as a standard RDIMM module (supports JEDEC® standard SPD data) during normal operation
- System up on power-loss, must place the DRAM in self-refresh operation (in current design offering)
- Backup and Restore initiated using either side-band signals or I2C interface
- DDR3 mode register values are restored during hand-off
- Brown-out or power-glitch (return of power during backup) is supported
- Able to survive multiple power glitches without losing protected data
- Once restore is complete, on-board FPGA will set the flag and turn on the bus switches for system to take over normal operation
- 4GB backup in 34 seconds
- Highly reliable SLC NAND flash
- 10+ year life span operating at 70°C ambient temperature
- 10+ year unpowered data retention
- Able to monitor external power supply controllers for state-of-health monitoring

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.

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.

SMART NVDIMMs require support from the system motherboard. When plugged in, the BIOS must recognize the NVDIMMs. OEMs 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. SMART's NVDIMM also features a USB interface that allows OEMs direct access to an NVDIMM backup image; this can bypass the need for BIOS changes.

How It Works

If there's a power failure, supercap module powers NVDIMM while it copies all data from DDR3 to on-module flash.



When power is restored, NVDIMM copies all data from flash to DDR3 and normal operation resumes.

Supercaps or Batteries?

NVDIMMs require a stored charge so they can back up their data to flash when power fails. For a multitude of reasons outlined below, SMART engineers chose supercapacitors (supercaps) rather than battery packs to support these NVDIMMs, and, have developed an assortment of supercap modules to accommodate various customer needs.

Better Power Efficiency

Supercaps are extremely efficient at charging, holding a charge, and discharging, making them a far more power efficient choice over batteries.

Very Fast Charge Rate

While batteries recharge very slowly, supercaps recharge almost instantly, so they are ready for the next power interruption event within milliseconds.

Broad Op Temp Range

Supercaps are reliable, and can store a full charge over a wide operating temperature range, while batteries tend to suffer a substantial loss of charge at extreme temperatures.

Long Lifespan with Zero Maintenance

Supercaps perform at spec with minimal degradation over time, and will function reliably when needed. Typically, they will also outlast the service life of most servers and storage appliances. To ensure reliable functionality, SMART supercap modules feature health monitoring that checks voltage, current and temperature values.

SMART NVDIMM Features and Specs

Supercaps perform at spec with minimal degradation over time, and will function reliably when needed. Typically, they will also outlast the service life of most servers and storage appliances. To ensure reliable functionality, SMART supercap modules feature health monitoring that checks voltage, current and temperature values.

Supercap Modules

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. SMART's supercap modules 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.



Supercap Module Features

- BBU, PCIe and 2.5" SSD form factors
- MCU controlled isolated SEPIC buck-boost charger
- Programmable input power monitoring
- Charge status indicator
- Max. continuous charge/discharge current: 2.5/7.1 A
- Programmable working mode
- Health monitoring: checks voltage, charge/discharge current, temperature; auto self-test; cell voltage balancing
- Variable # of cells of 50F/2.7V ultra capacitor
- 10 years lifetime @ 25°C
- Operating temperature range: -40°C to +65°C

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.

DDR4 Memory Modules - Ordering Information

288 PIN - DDR4 NVDIMM								
SMART Part Number	Density	Height (mm)	Module Config	Device Config	Speed	Voltage	Command Set	Temp
ST4097NJ420472SBL2	32GB	31.25	4Gb x72	2Gx4	2400MT/s	1.2V	JEDEC	0°C to +70°C
ST2047NJ420493SCL2	16GB	31.25	2Gb x72	2Gx4	2666MT/s	1.2V	JEDEC	0°C to +70°C
ST2047NR420493SCL2	16GB	31.25	2Gb x72	2Gx4	2666MT/s	1.2V	SMART	0°C to +70°C
ST1027NJ410493SE	8GB	31.25	1Gb x72	1Gx4	2666MT/s	1.2V	JEDEC	0°C to +70°C
ST1027NJ410472SE	8GB	31.25	1Gb x72	1Gx4	2400MT/s	1.2V	JEDEC	0°C to +70°C

Corporate Headquarters

39870 Eureka Dr.
Newark CA 94560
Tel: (+1) 800-956-7627
Tel: (+1) 510-623-1231
Fax: (+1) 510-623-1434
Email: info@smartm.com

Customer Service

Tel: (+1) 978-303-8500
Email: customers@smartm.com

Latin America

Tel: (+55) 11 4417-7200
Email: sales.br@smartm.com

EMEA

Tel: (+44) 7825-084427
Email: sales.euro@smartm.com

Asia/Pacific

Tel: (+65) 6678-7670
Email: sales.asia@smartm.com



www.smartm.com