


THREE USE CASES FOR STORAGE CLASS MEMORY (SCM)



Hussein Baligh

Senior Advisor, Sales Engineer Analyst
Dell Technologies
Hussein.baligh@dell.com

Lina Eissa

Associate Sales Engineer Analyst
Dell Technologies
Lina.eissa@dell.com

Ebrahim Serag

Associate Sales Engineer Analyst
Dell Technologies
Ebrahim.serag@dell.com

Abdelrahman El Nakeeb

Data Center Account Manager
Dell Technologies
Abdeirahman.elnakeeb@dell.com



The Dell Technologies Proven Professional Certification program validates a wide range of skills and competencies across multiple technologies and products.

From Associate, entry-level courses to Expert-level, experience-based exams, all professionals in or looking to begin a career in IT benefit from industry-leading training and certification paths from one of the world's most trusted technology partners.

Proven Professional certifications include:

- Cloud
- Converged/Hyperconverged Infrastructure
- Data Protection
- Data Science
- Networking
- Security
- Servers
- Storage
- Enterprise Architect

Courses are offered to meet different learning styles and schedules, including self-paced On Demand, remote-based Virtual Instructor-Led and in-person Classrooms.

Whether you are an experienced IT professional or just getting started, Dell Technologies Proven Professional certifications are designed to clearly signal proficiency to colleagues and employers.

[Learn more at www.dell.com/certification](http://www.dell.com/certification)

Table of Contents

1. Introduction	4
1.1 Evolution of Storage.....	4
1.2 Evolution of Memory	5
1.3 The need for a new technology	7
2. A new tier – What is Storage Class Memory?	8
3. Candidate Technologies for Storage Class Memory Devices	10
3.1 Phase Change Random Access Memory (PCRAM).....	10
3.1.1 Review on PCRAM.....	11
3.1.2 Case Study of PCRAM.....	11
3.2 Resistive Random-Access Memory (ReRAM or RRAM)	13
3.2.1 Review on RRAM.....	14
3.2.2 Case Study of RRAM.....	14
3.3 Magnetic Random-Access Memory (MRAM)	15
3.3.1 Review on MRAM.....	16
3.4 Spin-Transfer Torque Magnetic Random-Access Memory (STT-MRAM or STT-RAM)	16
3.4.1 Review on STT-MRAM.....	17
3.4.2 Case Study on STT-MRAM.....	17
3.5 Brief Technology Comparison	18
4. Three Use Cases for Storage Class Memory (SCM).....	19
4.1 Use Case 1: Tiering at the Storage-Side:	20
4.2 Use Case 2: Caching at the Storage Side	21
4.3 Use Case 3: Persistent Memory at the Server-Side	22
4.3.1 PMEM in Memory Mode.....	24
4.3.2 PMEM in Application Direct Mode	26
5. Conclusion.....	30
Appendix: Terminologies	31
References	32

Disclaimer: The views, processes or methodologies published in this article are those of the author. They do not necessarily reflect Dell Technologies' views, processes or methodologies.

1. Introduction

A data center is only as fast as its slowest link. For the past 50 years, Moore’s law has continued to deliver faster and faster processors; however, the storage side has struggled to keep up. Historically, innovation in storage came gradually: first there were magnetic tapes, and then came spinning hard disk drives (HDDs) with their mechanically limited speeds. The last big leap was the adoption of NAND Flash SSDs, which cut the latency from milliseconds to microseconds. Another development was the adoption of NVMe as the default SSD interface instead of SAS. However, the fastest component after the processor continued to be the volatile memory on the CPU side, the DRAM. An exciting new technology arose in the last few years to bridge this gap between DRAM and SSD: Storage Class Memory, also known as SCM. This paper will give you a brief on the history of memory and storage, the concept of SCM, SCM technologies, and SCM use cases.

1.1 Evolution of Storage

The need to store information existed a long time before the development of computers. In the 1700s, punch cards were used to give machines a sequence of instructions; the holes in the punch cards act as on and off switches. Back then, it was used in textile looms and player pianos. Years later, Herman Hollerith invented a calculator with movable parts that used punch cards not only to give instructions but also to store data; it was called the Analytical Engine, and it was used in the U.S. Census to speed up the process. It took them eight years to complete the 1880 census but thanks to Herman’s invention, it only took them one year for the 1890 US. Census. By 1950, punch cards had become an essential part of American industry and government. The warning, “Do not fold, spindle, or mutilate,” originated from punch cards.¹




Figure 1: Punch card from the mid-20th century [Ref. 2]



Figure 2: SLR Tape (1986 – 2015) [Ref. 3]

Then, another invention, the Magnetic Tape, was patented in 1928 by Fritz Pfleumer; Magnetic tapes are still used today for archival storage but it’s very slow in data retrieval. Then came Magnetic Disk Storage; its first form was the magnetic drum invented by Gustav Taushek. It used read/write heads to access data that is stored in drum tracks. IBM was the primary driver in the development of magnetic disk storage. Both the floppy disk drive and the hard disk drive (HDDs) were invented and improved by IBM engineers. These technologies were much faster than tape, but a new invention, Optical disc, was even faster. Optical discs were invented by James T. Russel, who had the idea of using light to record data. He was not taken seriously until Sony paid Russel millions of dollars to finish off his project and it finally came to life in 1980. From there, we started seeing CD, DVD and Blu-Ray in the market.¹



Figure 3: Hard Disk Drive (HDD) [Ref. 4]



Figure 4: A stack of CDs [Ref. 4]

Today, the main storage media used in enterprise environments are HDDs and tapes, with the addition of SSDs which will be described in the next section. HDD Magnetic disks is sometimes called the backbone of the data center, as it remains the dominant media for non-critical production, backup, and active archives. Despite its long lineage, Tape also remains a popular, cheap media for long term retention, while advancements in tape technology continues to improve its capacity and endurance.⁵

1.2 Evolution of Memory

Shift Registers

The earliest form of memory. Using J-K flip-flops and 8-bit shift registers, Shift Registers were built to hold one byte of information, but data was accessed by shifting around data in a circle. The latter saved connections, but it made it very challenging and slow to read the data. This was the earliest form of RAM.⁶

SRAM

The shift from Shift Registers to Static RAM (SRAM) was preceded by a continuous stream of inventions.

Frederick Viehe invented the Magnetic Core memory, which uses very small doughnut-shaped magnets, called Ferrite Cores, and address and sense wires that are connected to the magnets. By changing the magnetic polarity of the small iron oxide loops, 0s or 1s switches are created. One bit of data could be stored per ring and retrieved at any time. In 1953, MIT purchased the Magnetic Core memory patent and invented Whirlwind using its technology. Even though Magnetic Core was fast and efficient, manufacturing them was very time-consuming because they were hand-made.

In 1957, Andrew Bobeck invented the Twistor Magnetic that was very similar to Magnetic Core, but with magnetic wires replacing the circular magnets. It worked by having magnetic wires interwoven with current carrying wires; each intersection represented one bit of data. Twistor Magnetic proved to be more efficient than Magnetic Core memory because it was cheaper, weighed less, and required less current. After the Twistor's success, Bobeck developed another magnetic memory technology in the 1980's, known as Bubble Memory, a thin magnetic film that uses small-magnetized areas.

Finally, the Static RAM (SRAM) as we know it today was invented; each bit was still a J-K flip flop (like the Shift Registers model) made up of bipolar transistors. However, random-access memories (RAM) were a more appealing substitute to Shift Registers because you could access data in random order instead of shifting data in a circle.⁶

DRAM

In 1948, Professor Fredrick Williams, and colleagues, used vacuum tubes to develop the first Random Access Memory (RAM) for storing frequently used programming instructions to increase overall speed of the computer. Williams and co used an array of cathode-ray tubes, a form of vacuum tube, where the data was stored on the face as electrically charged electrons fired and represented as 1s and 0s to act as on/off switches.

Finally in 1968, Robert Dennard invented RAM as we know it today. It is a solid-state memory and known as dynamic random-access memory (DRAM), where transistors and capacitors are used to store bits of data. The concept of using capacitors (like a solid-state version of the Williams tube) to store data was a huge milestone in the development of RAM. It is much better than J-K flip flops because it is much cheaper; it reduced the number of transistors from a dozen to one transistor.⁶ It got even cheaper when MOS started being used instead of bipolar transistors. However, the problem was that it leaked off charge. The entire device had to be read and re-written to “refresh” the bits back to their original states. Although it leaked data, it was still very appealing to users because of its reduced cost. From there, the DRAM was born and the older version that was discussed earlier was renamed as “SRAM”⁷ Table 1 shows a comparison between the two.

BASIS FOR COMPARISON	SRAM	DRAM
Speed	Faster	Slower
Size	Small	Large
Cost	Expensive	Cheap
Used in	Cache memory	Main memory
Density	Less dense	Highly dense
Construction	Complex and uses transistors and latches.	Simple and uses capacitors and very few transistors.
Single block of memory requires	6 transistors	Only one transistor.
Charge leakage property	Not present	Present hence require power refresh circuitry
Power consumption	Low	High

Flash Drives and SSDs

Flash technology evolved from the concept of the RAM. Below are the different technologies that followed the SRAM and DRAM technologies; they were stepping stones to the development of Flash drives and Solid State Drives (SSDs).⁷

- **ROM** used one transistor per bit with a metal link to set the bit during the manufacturing process. The memory could be read from, but not written to.
- **PROM** emerged from the idea that programmers didn't want data to be permanently stored during manufacturing. A PROM's data is programmed into them after manufacture and it is used to test on a set of testing devices before writing data into all the devices manufactured.
- **Erasable PROM (EPROM)** was invented where data could be erased using UV light. However, the users didn't like that they had to remove the chip and apply UV light to erase the data. Therefore, **Erasable EPROM (EEPROM)** was developed to have two transistors instead of one where the data could be erased by one transistor and programmed by the other, without needing to pull out the chip. The price doubled because the transistors doubled but users were satisfied.

Finally, flash technology was born by trying to get the features of EEPROM with the low cost of EPROM. Hence, they created a new technology that only had one transistor responsible for programming and erasing. Moreover, to shrink the chip further, they removed half the signal lines and made all the bits communicate through their neighbors. From there **NAND flash technology** was born.⁷

As Flash drives and SSDs appeared in the market, these technologies have become essential data storage media for storing information. These technologies don't have moving parts; they use chips and transistors to store data. Over time, these chips and transistors are being enhanced to suit our increasing need for faster, denser storage systems.

1.3 The need for a new technology




Figure 5: Timeline of Memory Class Introductions [Ref. 20]

DRAM remains the gold standard in the handling of data in terms of access speed where it is as a memory by CPUs and caches. Storage media – NAND Flash, disk and tape – are slower but much cheaper and can scale to larger capacities than DRAM.

Unfortunately, this model is no longer suitable for everything. Emerging technologies such as Artificial Intelligence (AI), bio-informatics and in-memory databases all suffer from not having enough memory available from DRAM alone. These applications have very large working data sets that necessitate placement on the storage, but even NAND Flash, the fastest existing storage media, is still orders of magnitude slower than DRAM. Thus, the applications do not get the desired run time that could theoretically be reached if its data sets were all on DRAM.

To bridge this gap between DRAM and SSD, an exciting new family of technologies arose in the last few years, one that features memory-like performance at storage-like cost. These technologies are all under the name Storage Class Memory (SCM).

2. A new tier – What is Storage Class Memory?

“Ideally one would desire an indefinitely large memory capacity such that any particular ... word would be immediately available. ... It does not seem possible physically to achieve such a capacity. We are therefore forced to recognize the possibility of constructing a hierarchy of memories, each of which has greater capacity than the preceding but which is less quickly accessible.”



Figure 6: Herman Goldstine and John von Neumann at The Princeton Institute for Advanced Study [Ref. 9]

Prediction of the necessity of SCM in 1946. Taken from the Preliminary Discussion of the Logical Design of an Electronic Computing Instrument by Arthur Burks, Herman Goldstine and John von Neumann.⁹

Storage Class Memory (SCM) is a new tier of memory/storage that is at the top of DRAM and at the bottom of NAND flash. Figure 7 illustrates the position of SCM in the storage hierarchy. SCM, or as it is sometimes known, Persistent Memory (PMEM), offers fast non-volatile memory to the processor, with speeds slightly below DRAM but still vastly above those of even the fastest NAND Flash storage, while at the same time having capacities at the scale of these NAND Flash drives.¹⁰

IT system architects have always wanted to get primary data access latencies as low as possible while knowing the DRAM capacity and cost constraints and limitations. NAND SSDs started to shift their interface from SATA and SAS to NVMe (Non-Volatile Memory express) which uses the PCIe bus in servers as its transport. SSDs using NVMe have access latencies as low as 100 microseconds. However, that is still slower than DRAM. The need for a new technology to cover the market demands was becoming essential. SCM is taking place in the market to fill that gap.

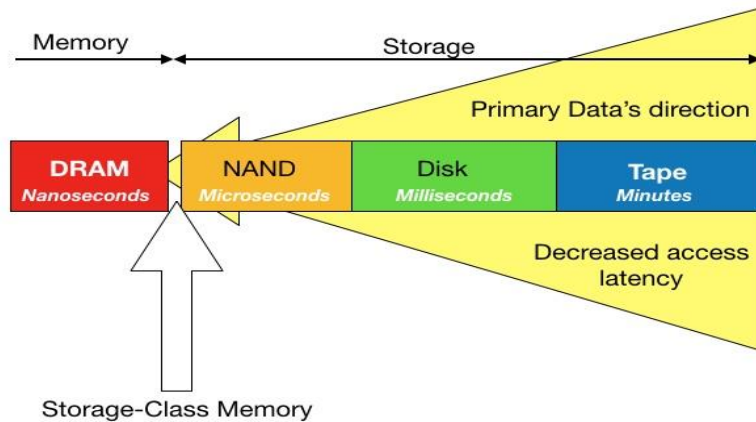


Figure 7: Diagram 2: The position of SCM in the storage hierarchy [Ref. 10]

IT system architects have always wanted to get primary data access latencies as low as possible while knowing the DRAM capacity and cost constraints and limitations. NAND SSDs started to shift their interface from SATA and SAS to NVMe (Non-Volatile Memory express) which uses the PCIe bus in servers as its transport. SSDs using NVMe have access latencies as low as 100 microseconds. However, that is still slower than DRAM. The need for a new technology to cover the market demands was becoming essential. SCM is taking place in the market to fill that gap.

SCM is a persistent memory that acts as a compromise between SSD and DRAM features. Even though both DRAM and NAND SSDs are made of solid-state chips and both are under the umbrella of solid-state storage, they have completely different roles. In the past, data storage and memory had the same definition; now, data storage is defined as the long term and memory (RAM) is defined as short term. SSDs are mainly used for storage and RAM is used to perform calculations and operations from the storage retrieved from the primary storage.

By default, RAM and SSDs are faster than other storage alternatives such as hard disk and tape. But if we compare RAM with SSDs, RAM is much faster. NAND Flash SSDs are persistent and relatively cheap, but slow, while DRAM is a volatile memory that is a thousand times faster and benefits from its proximity to the CPU. The reason they have different speeds is due to their different topologies. NAND are divided into blocks of cells that are at least 128KB in size. In order to write in NAND, the data in the blocks have to be erased, then the new data is added and the whole block is written back again. For instance, if you want to write a 4KB in SSD and there is already data in the SSD, the data in one of the blocks has to be completely erased, then the 4KB is added to the already existing data, and the existing data and the new 4 KB are written back again in the empty block. On the contrary, the DRAM doesn't have this limitation where you

must erase before you write. The DRAM consists of capacitor and transistors where you can charge or discharge it. Therefore, DRAM's topology is much faster and also allows random accesses, enabling users to read from or write in whatever order. DRAM storage is dynamic so it needs a new electronic charge every few milliseconds to recompense for charge leaks from the capacitor. It always needs power to keep going and that is why it is very expensive compared to SSDs.

Moreover, another feature used to compare SSDs with DRAM is their volatility. DRAM is volatile because it is temporary and when a computer loses power, the data is lost and is irretrievable. SSDs are non-volatile storage; it doesn't need power to store data. Other non-volatile media includes flash drives, NVRAM, ROM, floppy disks, magnetic tape, optical discs, and punch cards but SSDs remain the fastest of the latter.

SCM technology seeks to take the best from both worlds by being a non-volatile persistent storage that is faster than SSD but slower than DRAM, and costs more than SSD but less than DRAM. Tech companies and customers can use SCM as a much faster SSD or a bigger DRAM. SCM can scale higher than DRAM, reaching up to 15 or 20 times the capacity per die at a lower cost, without the need to add more CPUs and more memory slots, and without incurring the huge power costs of DRAM that follow. However, the most important advantage that SCM has is granularity: SCM can be both byte-addressable (like DRAM) and block-addressable (like NAND SSD). Therefore, the processing time to convert from block to bytes will not be a burden as before. SCM has huge repercussions for the place and use of SCM in the data center. It can truly unlock their remarkable potential to handle different workloads and demands of the market.

3. Candidate Technologies for Storage Class Memory Devices

Technologies like Phase Change Random Access Memory (**PCRAM**), Resistive Random-Access Memory (**ReRAM**), Magnetic Random-Access Memory (**MRAM**), spin-transfer torque magnetic random-access memory (**STT-MRAM**) and others are showing potential candidacy to be adopted as the technologies that power SCM.¹¹

To achieve the best solid-state memory and traditional hard disk capabilities, low-cost Non-volatile memory (NVM) technology must be researched and developed. NVM is not new to researchers, having been considered as a future substitute for NAND Flash technology. With the recent enhancements to NAND, the urgency to replace it has subsided, opening the door to other applications for NVM technology – namely SCM – that will likely capitalize of its higher performance, advanced addressing, superior endurance and greater reliability compared to NAND.¹²

3.1 Phase Change Random Access Memory

Phase Change Random Access Memory (PCRAM) capitalizes on the properties of phase change materials (PCMs). Such materials exist in either a formless phase scientifically referred to as an “amorphous phase”, or a phase of form scientifically referred to as “crystalline phase(s)”. The phase switching of these materials can be triggered via optical or electrical heating. In PCRAM, the crystalline phase is achieved from the amorphous phase by electrically heating the material for a long enough period at the appropriate temperature, a process viewed as a “set” operation. The crystalline phase has relatively lower electrical resistance compared to the amorphous phase which can be “reset” by applying a short impulse of electrical current to the material, heating it beyond its melting point.




Figure 8: Phase change materials and phase change memory [Ref. 13]

3.1.1 Review on PCRAM

There are many PCMs being researched to perfect the operation of PCRAM. Since PCM is the main component of this memory technology, any adjustments in it would greatly impact reliability and performance of the PCRAM.

One common PCM – $\text{Ge}_2\text{Sb}_2\text{Te}_5$ (GST) – retains data at a state of no-power for up to ten years. Performance is relatively faster than NAND SSDs, but greatly impacted by the significantly longer duration of time a set operation takes.

Other PCMs have shown promise in reducing the penalty of the set operation without affecting the reliability and the nonvolatility of the PCRAM.¹⁴

3.1.2 Case Study of PCRAM

3D X-Point was jointly developed by Intel and Micron, a good example of Memory Technology based on PCRAM. This technology has been shipped with Intel’s Optane Drives which have been outperforming SSDs ever since their release. Although Micron have only unveiled an SSD product – the QuantX family – Intel’s Optane family is already shipping both flavors of their attempt at “Universal storage”: NVMe-PCIe SSD as well as NVDIMM.

3D X-Point SSDs are compatible with current and future hardware or operating systems that support NVMe. This enabled faster adoption of these drives in attempt to close the memory/storage gap from the storage side.

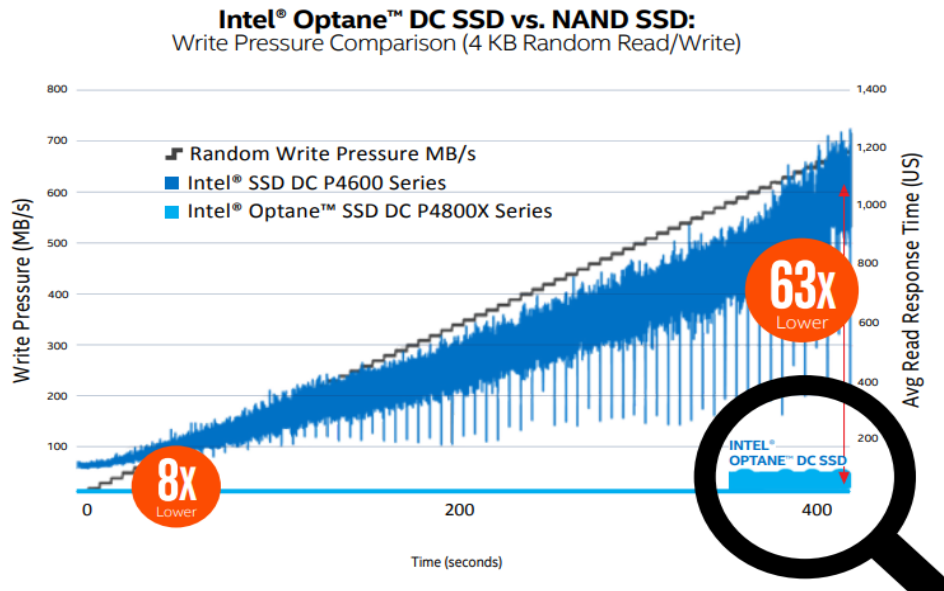


Figure 9: Intel Optane DCP4800X provides consistently low latency compared to an intel 3D NAND SSD [Ref. 15]

NVDIMM or Persistent Memory Use of Optane is another story, as it requires newer hardware (chipset and motherboard) specially designed to support the operation of Optane in a Persistent Memory role. This role allows for much more economical growth of the DRAM, supporting more data to be processed in-memory. According to Intel, Optane’s flavors can enhance data tiering from Memory to cold storage.

Memory and Storage Data Tiers in the Data Center




Figure 10: Memory and Storage Tiers (red = "hot" data, yellow = "warm" data, and blue = "cold" data) [Ref. 15]

Micron QuantX X100 NVMe SSD promises high-performance local storage with 2.5M IOPs support. It claims to be 3 times faster than SSDs in the market with latency under 8 μ s and 9GB or reads and writes per second.

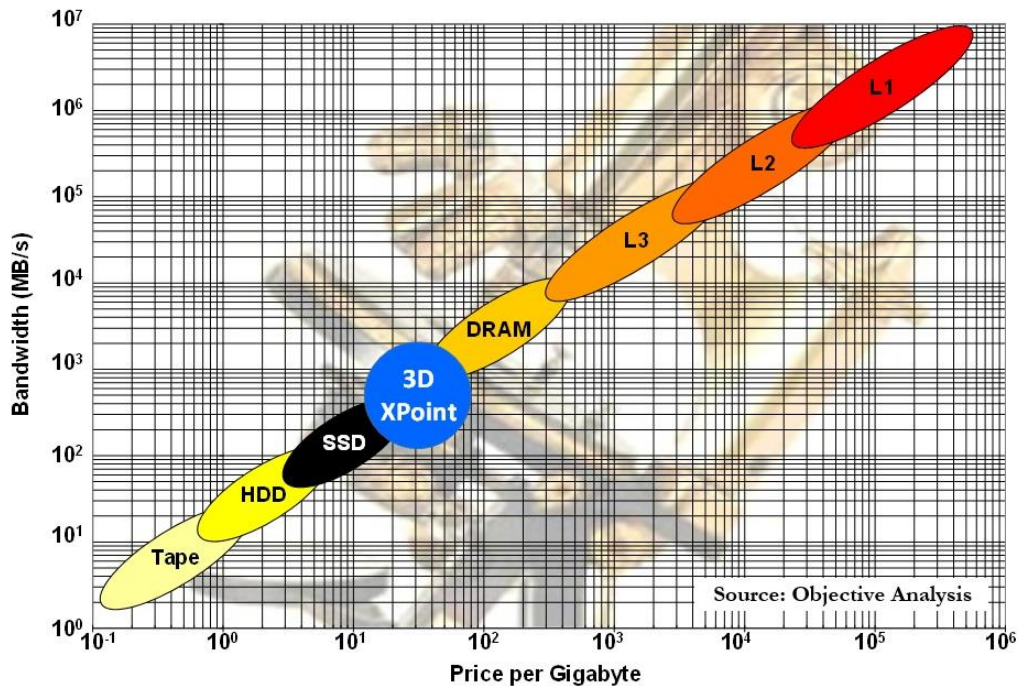


Figure 11: Where 3D X-Point Technology should fall in the bandwidth vs Price per Gigabyte graph to be competitive and adopted [Ref. 16]

A main adopter of Intel’s Optane is Dell EMC’s PowerMax. The Intel Optane SSD DC D4800X in combination with PowerMax yields 50% faster response time and lower latency and can scale up to 4 PB with 256 front-end ports and reach 15M IOPS.²¹

3.2 Resistive Random-Access Memory (ReRAM or RRAM)

Resistive Random-Access Memory (ReRAM or RAM) exploits the physical phenomena of Resistive Switching (RS). In this phenomena, a dielectric object experiences change of its resistance across its terminals due to being subject to electric current or field. Switching from a high resistance state to a low resistance state is the SET operation, while switching from low resistance state to high resistance state is the Reset operation. The Memristors or the cell of the RRAM has a simple structure metal-insulator-metal with a conductive filament.

Some may conclude that Phase change memories are a subset of the RRAM, as they share some of the characteristics of the RRAM. This makes the RRAM family a worthy of consideration for SCM.




Figure 12 - Basic RRAM cell structure [Ref. 17]

3.2.1 Review on RRAM

RRAM has attracted a lot of attention over the years, as it offers a smaller and simpler cell structure than most SCM candidates which is important for scalability. It is potentially inexpensive and operates at very high speed. If research on improving endurance of the cell and the data retention of the memory yield success, RRAM can be generalized across all tiers of digital memory.¹¹

3.2.2 Case Study of RRAM

Crossbar Inc. is a California based company that specializes in developing ReRAM products, with applications that range from Data Centers to Consumer Electronics.

Crossbar lists the ReRAM advantages in three categories; Energy, Performance and Density.

- Energy: being non-volatile and not requiring power to retain data, ReRAM consumes energy 20 times less than NAND.
- Performance: Crossbar claims ReRAM has 100 times lower read latency, 1000 times faster writes, 10 years of retention and 1000 times greater endurance of its cells compared to NAND.
- Density: Crossbar ReRAM is twice as dense as 3D NAND, stackable, scales up to Terabytes of capacity while its cell scales down to 10 nm in size.

Crossbar’s T series NVDIMM is available in densities up to 1TB. The NVDIMM shows greater performance at a considerably low energy consumption.

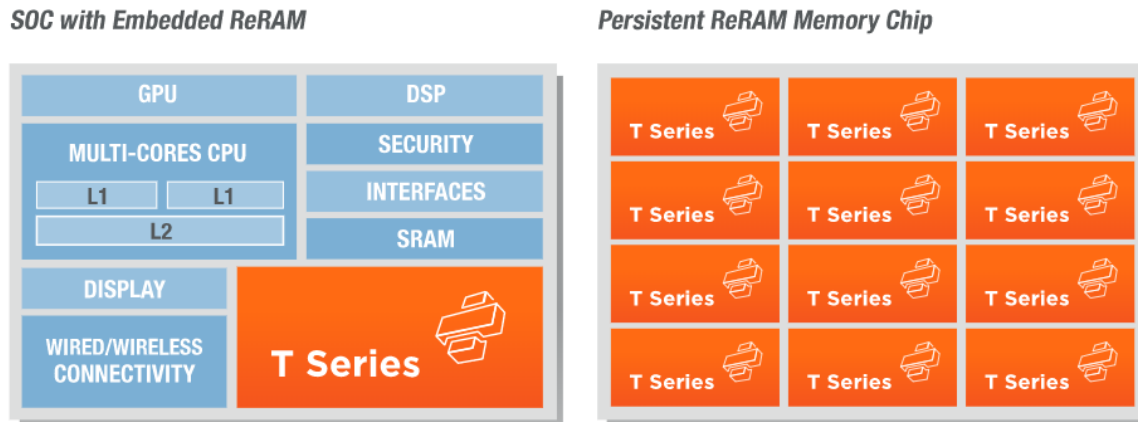


Figure 13 - Embedded Persistent Memory Array & Stand-Alone Chip [Ref. 18]

At a lower cost than DRAM per Gigabyte, and at read equivalent performance, Crossbar believes the T Series to be a cost-effective solution for customers that require high density, high performance, low power, non-volatile data storage. Crossbar is also targeting data centers growing market trend of Hyper-converged servers where computing, storage and networking become designed in a server form factor.¹⁸

ReRAM technologies do not require erase operation and garbage collection. They support bit/byte overwrite capability as well as being less complex for the storage controller, enabling better performance compared to Flash-based Solid-State solutions.

The ReRAM-based NVMe SSD storage solution read/write latency is 10x lower and the IOPS @ 512B IO is 8x faster compared to Flash NVMe SSD.¹⁸

Others like Fujitsu, Panasonic and Adesto have developed ReRAM memory products as well.¹⁹

3.3 Magnetic Random-Access Memory (MRAM)

Magnetic Random-Access Memory (MRAM) stores data using magnetic charges. MRAM cells consist of two magnetic elements, one of which has a fixed magnetic polarity while the other has a switchable (i.e. programable) polarity. The magnetic elements are placed on top of each other with an insulating tunnel separating them.

Writing and Erasing of the cell is achieved by passing an electrical current through the write line to induce a magnetic field across the cell. Reading examines the electrical resistance of the cell as a current passed through it. An “On” occurs when low resistance is established, this is due to the parallel magnetic moments of the magnetic elements; magnetic tunnel effect. The opposite occurs when the magnetic moments are antiparallel, leading to a state of High resistance between magnetic elements.




Figure 14 - Basic MRAM Cell structure [Ref. 11]

3.3.1 Review on MRAM

MRAM shows great potential in terms of speed and power consumption; it has access time in nanoseconds. However, some obstacles have limited its adoption. Scalability was obstructed by interference of magnetic field between neighboring cells. Another obstacle was its high cost. This called for efforts from researchers and developers to construct variations of MRAM. Such variants would eventually become strong candidates for SCM.¹¹

3.4 Spin-Transfer Torque Magnetic Random-Access Memory

Spin-Transfer Torque Magnetic Random-Access Memory (STT-MRAM or STT-RAM) is a variant of MRAM that substitutes using magnetic fields with electricity to change magnetization. Like MRAM the parallelism of the magnetic moment between the magnetic elements results in a low resistance state and allows for current to flow. High resistance state is achieved by antiparallel magnetic moments. Switching of the magnetization is induced by a write electric current that is connected via a transistor.

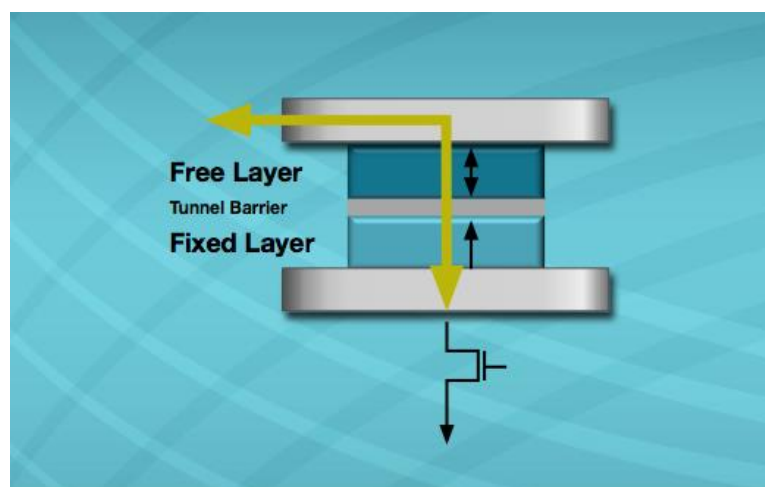


Figure 15 - Single STT-MRAM cell

3.4.1 Review on STT-MRAM

STT-MRAM is such a powerful Non-volatile memory that it is being purposed for cache. Recent developments that have ensured high endurance, years of retention as well as ability to withstand harsh temperatures have indicated the candidacy of STT-MRAM to replace static random-access memory (SRAM) as cache. Nevertheless, there are still many obstacles to be mitigated by developers, namely the write speed of STT-MRAM that is still inferior to SRAM.¹⁴

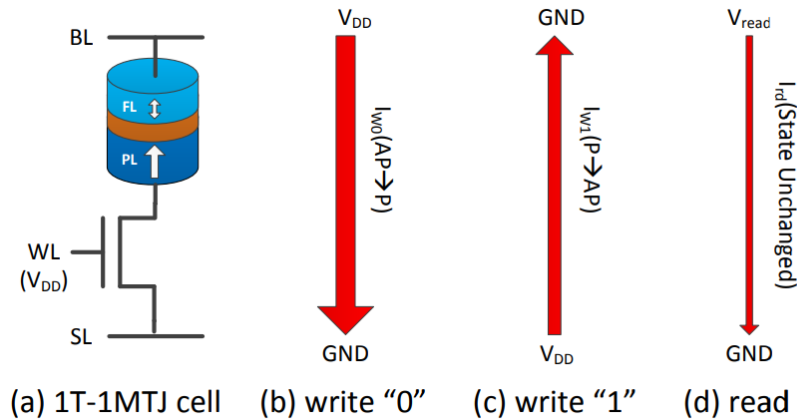


Figure 16 - The write and read operations on STT-MRAM cell "1T-1MTJ" [Ref. 22]

3.4.2 Case Study on STT-MRAM

Everspin developed nvNITRO Storage Accelerator based on STT-MRAM Technology, utilizing its high speed performance to accelerate enterprise compute and storage systems that address demanding applications. The nvNITRO storage accelerator applications include database and application acceleration, all flash storage acceleration, file system acceleration, Shared Remote Persistent Memory, and more. This is thanks to its ultra-low access latency that goes as low as $2\mu\text{s}$, and 1.5M+ random 4KB IOPS.²³

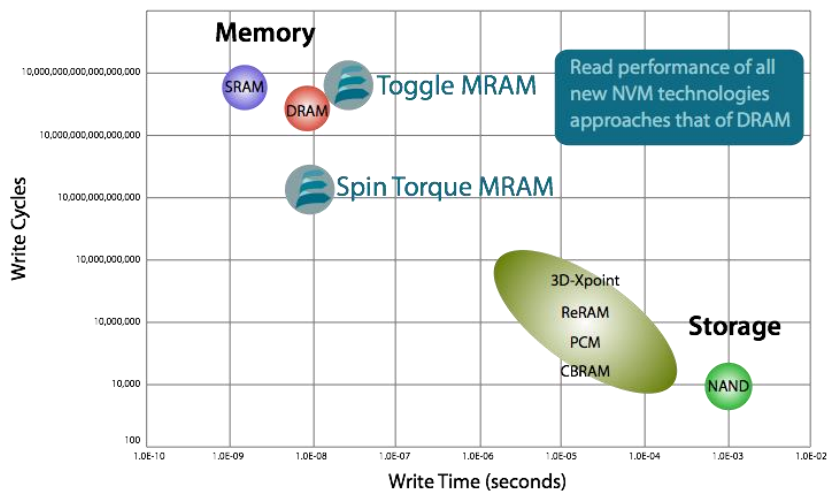


Figure 17: Read performance of all new NVM technologies [Ref. 24]

3.5 Brief Technology Comparison

	SCM Technologies					Flash (NAND)
	SRAM	DRAM	PCRAM	ReRAM	STT-MRAM	
Non-volatile	No	No	Yes	Yes	Yes	Yes
Cell size [F^2]	50-120	6-10	4-19	6-10	6-20	5
Read time [ns]	≤ 2	15	≈ 2	1-20	2-20	10^3
Write/time [ns]	≤ 2	15	10^2	50	≈ 10	10^6
Write Power	Low	Low	Low	Medium	Low	High
Endurance (cycles)	10^{16}	10^{16}	10^{10}	10^6	10^{15}	10^5
Future Scalability	Good	Limited	Promising	Promising	V.Promising	Limited
Relative Cost per GB	\$\$\$\$	\$\$\$\$	\$\$	\$\$\$	\$\$\$	\$

Table 1: Compiled Comparison of some of the Technologies

Commentary on the comparison

Although SRAM, DRAM and Flash are not emerging technologies, they have been included in the comparison to highlight where the SCM candidate technologies fall in the Memory-Storage Spectrum.

The Desirable values for a perfect candidate would be small cell size with likely scalability, low power consumption, high endurance, as well as fast reads and writes. Being cost effective is also desirable and could greatly affect desirability.

PCRAM is overall good, but it is thrown off by its very slow write operation compared to other candidates. It makes up for this by being cost effective. ReRAM is also acceptable except for endurance which is relatively low.

STT-MRAM seems to be the most balanced of all three, but would face adoption issues and competition if cost is not mitigated.

4. Three Use Cases for Storage Class Memory

Intro: How and where to use SCM

Tech vendors and customers have great flexibility in the method they can choose to integrate SCM in their existing products. Since SCM can act either as a (much faster) SSD, or as a bigger DIMM, it follows that companies would capitalize on this backwards compatibility and start using SCM for its performance advantages with minimal disruption or modifications. On the other hand, to gain more advantages from their investment, companies have started developing and updating their infrastructure to use SCM in ways that can truly unlock their impressive potential to handle the difficult workloads and demands of the latest applications on the cutting edge of technology.

One of the biggest points of differentiation between Memory and Storage is how the data is accessed. Storage media is block-addressable and is accessed using read/write commands from the application that must pass through a driver and software/hardware stack necessary to convert from blocks to bytes before reaching the storage. This processing takes time and makes it much slower in comparison to Memory access. Memory is accessed using load and store operations, which are commands to move data between CPU registers and memory, eliminating the overhead and thus giving memory a much faster path to the application.²⁵

SCM offers a radical new capability in that it can be addressed at either the byte OR the block level. Initially this positioned SCM as a potential future replacement to NAND flash as a main tier, or as a persistent replacement to DRAM as a memory. However, NAND flash has maintained its popularity, being inexpensive, scalable, and easy to manufacture. On the memory side, DRAM continues to be the gold standard in speed, and the persistence and capacity gains that could be achieved with SCM are offset by its slow speed, so DRAM is in no danger of being replaced either.²⁶

However, this means that there is a strong opportunity for SCM to find success in implementations that are not simply “replacing” NAND Flash or DRAM, but by instead as acting as a completely new layer whose purpose is to complement and improve existing memory and storage configurations.²⁷

These unique properties have enabled the server and storage vendors to use SCM in a variety of ways that might come to dominate the enterprise market in the coming period. In this article we will focus on three specific use cases that we believe show the most promise.

- **Tiering at the Storage-Side:** SCM can be integrated with the storage array alongside NAND Flash drives by standardizing NVMe bus technology for both. In this case, the SCM acts as an even faster “Tier 0” inside the array to handle the hottest data and relegate the rest to SSD. While there will still be network delays, NVMe-over-Fiber technology exists to alleviate this issue.
- **Caching at the Storage-Side:** SCM can be used as a basic read cache for the external storage array, taking advantage of its incredibly strong random read and write IOPS capabilities to accelerate the performance of the array. However, the network delay between the compute side and storage side will mean that we can’t take full advantage of SCM using this method.

- **Persistent Memory at the Server-Side:** Also known as PMEM, this is the use case with the least development but biggest potential. It comes in two modes:
 - SCM can plug directly as an NV-DIMM, extending the capacity of the existing DRAM in a configuration known as **Memory Access Mode**.
 - Another approach uses SCM to create a second tier of memory next to DRAM, similar to the tiered storage concept. In this case, called **App Direct Mode, Direct Access Mode (or DAX)**, the application can directly access the non-volatile SCM, thereby having extremely low latencies.

The following sections will deal with each of them in turn.

4.1 Use Case 1: Tiering at the Storage-Side

Description

Storage Class Memory can be used in Storage Arrays as a Tier or as a Cache. These two use cases are vastly different despite their seeming similarity.

Tiering means that SCM drives constitute a persistent storage tier in a storage array next to standard NAND flash drives, and directly supported by the array's OS. Inside the array, SCM SSDs can be mounted in standard disk enclosures using NVMe interface. Data placement can be handled by automated storage tiering that periodically moves data between the storage tiers based on data access statistics collected over that period. The algorithm then places the most active data on the fastest tier and vice versa. SCM drives will act as "Tier 0" which houses the most active data on the system. Due to complexity of identifying what data would most benefit from being put into Tier 0, and which data can be relegated down to the slower NAND flash, machine learning (ML) algorithms using predictive analytics and pattern recognition can be used for optimal data placement and SCM utilization.

Advantages of Tiering: Tiering improves performance and total cost of ownership (TCO) by targeting only the hot data for the SCM tier. This is in contrast to Caching which only improves performance. In Caching, SCM operates as a high capacity secondary cache that sits between system cache (or DRAM) and the NAND flash drives in the system, in order to house the frequently accessed data instead of the NAND flash storage. Thus, caching is great when the whole data set is to be uploaded into the SCM for read-heavy activities. However, active data sets larger than the cache can handle can result in thrashing where the cache has to constantly swap in/out cache objects, where they expire and are forced to be dropped to make space for the next piece of active data, leading to a constant state of paging and page faults. This results in a huge, sudden drop in I/O performance.

With Tiering, read and write I/O are performed at the speed of the SCM as long as sufficient capacity exists. Hot, active data is placed on the SCM resulting in the most IOPS being serviced at faster speed, while the cooler, less active data is moved down to the flash, thereby improving performance and TCO.²⁸

Tiering in Action: Dell EMC PowerMax

Dell Technologies introduced the PowerMax array in 2018 with an all NVMe architecture supporting both NAND Flash and SCM drives in the same disk enclosures. With the addition of Intel Optane drives as a SCM storage tier in PowerMax, along with AI-ML-powered data placement techniques, PowerMax can reach up to 350GB/s and up to 15 million IOPS, all with 50% lower latency.²⁹

The PowerMax array is the first Dell EMC hardware platform to offer end-to-end NVMe from Servers (HBAs) to PowerMax drives (SCM/ Flash) using low latency NVMe-oF. NVMe flash drives constitute the majority of a PowerMax array's capacity, with SCM left to handle the most demanding apps requiring low latency.

PowerMax OS provides data-at-rest encryption and deduplication to the Optane drives, however compression is not supported on the SCM tier. Unusually for Intel drives, these Optane SSDs will be specifically dual-ported to ensure resiliency against single port failures. SRDF (Symmetrix Remote Data Facility) replication offers another layer of protection.³⁰

PowerMax can also be configured completely with SCM drives only, making it the first 100% SCM system. In these systems, both compression and deduplication are supported on the Optane drives.

PowerMax systems support two Intel Optane drive models; the 750GB and 1.5TB dual-port Optane DC 4800X SSDs with NVMe interfaces.

As a result of the new SCM drives and support for NVMe-oF with 32Gb Fibre Channel at the front-end connectivity, Dell says customers will see extreme performance improvements and lower latency:³¹

- Up to 350GB/sec bandwidth – more than 2X increase
- Up to 15 million IOPS – 50% improvement
- Up to 50% lower response times vs. older VMAX flash drives
- Under 100 microseconds read latency

4.2 Use Case 2: Caching at the Storage Side

Description

In this case, SCM is used in the storage system back end where it acts as the caching tier for hot data with ultra-low latencies. SCM operates as a high capacity secondary cache that sits between system cache (or DRAM) and the NAND flash drives in the system, in order to house the frequently accessed data instead of the NAND flash storage.³²

Historically, whenever a new, faster storage technology was introduced, it was usually first utilized as a storage accelerator in a cached design. This was the case when SSD's were first added to storage arrays: due to their relatively high cost at the time, a small number were added to storage arrays to be used as a Cache in front of the existing HDD's to capitalize on their blazingly fast performance, plus to make the most out of the limited number of SSD's. Years later, we now see the same scenario repeating itself, but with SSD as the dominant storage type in the array, and SCM in the role of the expensive but fast Cache layer that can transform the array's performance.

Advantages of Caching

The main advantage of using caching method is that a relatively small number of SCM drives can enable a storage array to present applications with SCM speed and performance by dynamically caching only the hot active data, while storing the bulk of the data on much slower and less expensive media.

Another point to consider is that Caching is more dynamic than tiering, since it responds to workloads immediately instead of waiting for a tiering algorithm to move active data to the high speed SCM tier on a scheduled basis. Due to its periodic nature, a system using tiering may miss out on accelerating spikes that are not a workload with a maintained high, so caching, with its quick response time is better for applications with dynamic, unpredictable workload changes.

In both Tiering and Caching, the intelligent algorithms that automate data placement are critical. Most storage vendors have introduced artificial intelligence and machine learning (AI/ML) in their storage arrays to dynamically optimize the process. A new idea being floated is to add NVMe and SCM into enterprise storage as both tier AND cache, and then use AI/ML powered algorithms to make informed decisions on the most available SCM capacity, thereby delivering the best performance in the most cost-effective way.³³

Caching in Action: HPE Memory-Driven Flash

HPE is one of the vendors that chose the caching approach to integrating SCM in their arrays. The HPE 3PAR 20000 storage system supports read caching on Intel Optane 750 GB NVMe SCM Module add-in cards, calling it “HPE Memory-Driven Flash”. The main storage of the array is NVMe NAND Flash drives.

HPE Memory-Driven Flash integrates SCM and NVMe with software intelligence to deliver unprecedented application performance, up to 50% faster in comparison to all-flash arrays with standard NVMe SSDs. This ensures unmatched predictability, enabling all critical workloads to be serviced with the best performance, with an average latency under-200us and near 100% IO within 300us.³⁴

HPE 3PAR 20850 arrays equipped with this technology delivered:³³

- Less than 200 microseconds of latency for most IO
- Nearly 100% of IO in under 300 microseconds
- 75 GB per second of throughput
- 4 million IOPS

4.3 Use Case 3: Persistent Memory at the Server-Side

Description

When Storage Class Memory is deployed at the server, it is more commonly called **Persistent Memory** – PMEM for short. PMEM has two modes: Memory Mode and Application Direct Mode (App Direct mode).

As mentioned in a previous section, SCM is not intended to replace DRAM, which is why both modes involve using PMEMs and DRAM together, albeit with different configurations.

For now, **Memory Mode** is the one more commonly used, which makes the PMEM appear as if it's a larger than usual DRAM. In this mode, the PMEMs are the main bulk of the memory stores with the DRAM DIMMs as a faster buffer to handle the needed data as rapidly as possible. This offers better performance than the prevailing configuration that now uses DRAM to buffer NVMe NAND flash SSDs. Memory mode is backwards compatible, and can work with any application or OS. However, it is not actually persistent, as will be explained in the next section.

To have PMEM live up to its name and be used as a Persistent Memory requires **App Direct Mode (aka Direct Access Mode, or DAX)**. App Direct Mode makes a PMEM act like a slower tier of DRAM disk, but at the same time data is never erased when the power is shut off. This enables fast restarts by minimizing downtime and data loss. The fact that SCM offers much bigger memory sizes make it easier to run any in-memory database easier. But unlike Memory Mode, App Direct Mode isn't backwards compatible since it requires potentially huge modifications to the application and possibly the file system.

A mixed mode configuration is also possible. The PMEM's capacity can be divided to be used as both modes at the same time; part of it is used in Memory Mode and part in App Direct Mode.³⁵

The next two sections discuss Memory Mode and App Direct mode in more detail.

4.3.1 PMEM in Memory Mode

Convincing customers to adopt new technologies is never easy, however one thing that helps is to guarantee backwards compatibility with their existing architecture. For that reason, PMEM in Memory Mode is the go-to mode for customers hesitant to adopt PMEM. In this mode, SCM is plug-and-play and can work with the existing software and applications with no modification.

PMEM DIMMS cannot be used as the sole memory in a system, they must accompany standard DRAM. In practice, installing a small DRAM DIMM accompanied by a (by definition) large capacity PMEM DIMM in your server makes the OS and applications believe the whole capacity is simple DRAM. Each PMEM is paired with a DRAM where the DRAM acts as a cache – a small but fast device in front of a large but relatively slower PMEM and works to accelerate its operations by holding the most frequently accessed data. This whole process is invisible to the user, which doesn't even see the DRAM capacity; it just sees the PMEM capacity alone but assumes that it is DRAM. For example, if you have a system with 64GB DRAM DIMM plus a 512GB PMEM DIMM, the total memory size that appears to the OS is only 512GB.

The role of the faster DRAM in this configuration is to temporarily store the most active data that usually resides on the slower PMEM, and the cache controller intelligently places the data such that the effect is that you have a large PMEM DIMM (512GB in our example) performing at the much higher speed of the 64GB DRAM. Of course, eventually there is a read or write miss, where the data must be retrieved from the PMEM, but this happens so rarely that on average the whole system is running at 95% or 99% of the speed of DRAM; almost no difference.³⁶

Memory Mode is backwards compatible, but volatile; the data is not persistent and is erased with any loss of power, even though by its very name PMEM is persistent. Memory Mode creates a single hybrid pool of memory composed of DRAM plus PMEM, which gives it its backwards compatibility, however, this also means that the OS and applications sees the SCM as regular system memory, and as such are designed to disregard any data in the memory after a power loss since there is no way to recover it. The other big advantage of Memory Mode is its the large capacity. The biggest DRAM DIMM currently sold is only 128GB, while the PMEM is available in sizes from 512GB and up, at lower prices. That, combined with the backwards compatibility, means that PMEM in Memory Mode can work seamlessly with legacy applications like Big Data analytics.

Memory Mode is best for customers with legacy applications that require large amounts of memory such as virtualized databases. It is also more suitable with applications that have consistent or predictable data access patterns that can be analyzed by the controller to proactively bring the most frequently used data to DRAM. Naturally, the opposite case, applications with random workloads will not be a suitable fit for PMEM Memory Mode and would benefit more being used in a pure DRAM system.³⁷

Memory Mode in Action: VMware vSphere using memory mode

VMware is one of the first companies to embrace Persistent Memory and utilize it best. Using Intel Optane drives as PMEM in Memory Mode, the new technology greatly increased the capacity of the memory while requiring no changes to the existing software, guest OS, or VMs.

The way that Memory Mode benefits servers is simple: since there is a limited number of DIMM slots, if all are carrying DRAM then you have a finite limit to the amount of memory available to the hypervisor. However, if you use a mixture of PMEM and DRAM, then the maximum memory carried by the same number of slots is multiplied. For example, while a 2-socket server with 24 DIMM slots populated with 64GB DIMMs gives a maximum of 1536 GB, if you replace some of the DRAM with 512GB DIMMS you can reach up to 6114 GB of memory without making any adjustments to the VMs.

The expanded memory of the hypervisor can also result in tripling the number of VMs that can run concurrently. VMware has tested the above scenario using a future version of vSphere and found that you can triple the number of virtualized SAS analytics workloads on your 2-socket server, which would usually require using a 4-socket server. The results are shown in Figure18. Again, no changes were needed to either the software or application to get these benefits.³⁸

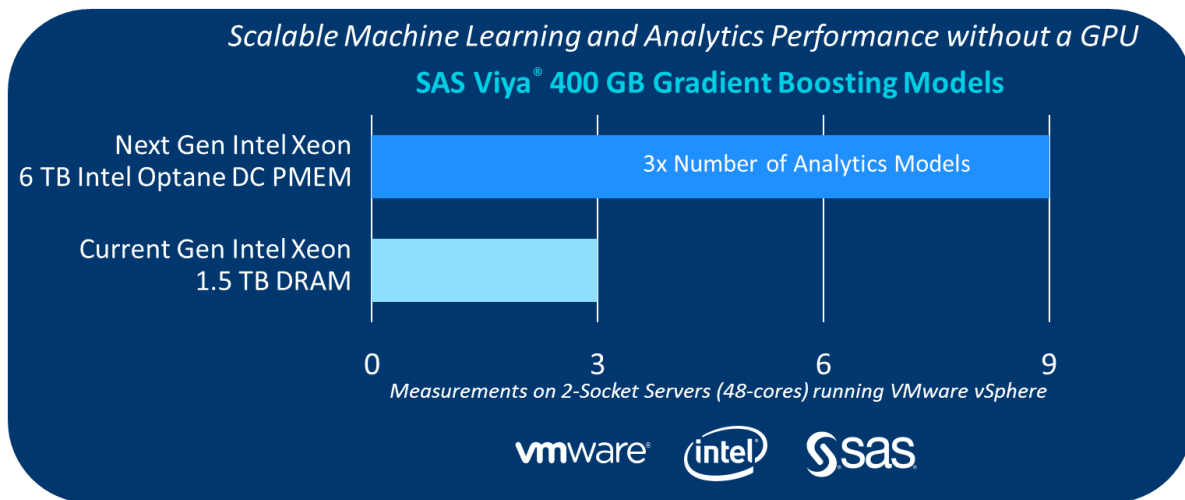


Figure 18: VMware performance improvement using Intel Optane PMEM in Memory Mode [Ref. 38]

4.3.2 PMEM in Application Direct Mode

PMEM in App Direct Mode (aka Direct Access Mode, or DAX), is when the OS and applications are designed to be aware of there being two tiers of system memory – DRAM which is fast yet volatile and PMEM which is slower but persistent. In this case the PM-aware software can intelligently target the best type when it needs to write data. It can use DRAM for applications that require the lowest latency but can tolerate data loss, and use PMEM for workloads that need persistence or data sets that are too large for DRAM. This ability to target certain writes to the persistent storage and others to the DRAM is the crucial difference between Application Direct Mode and Memory Mode.

The first adopters of this technology will likely be customers that have full control of their closed software stack from top to bottom, like some custom applications and closed hyperscale data centers. Having full control of your software makes it easier for the company to modify it to accept Persistent Memory and update it regularly to fix any issues as they arise.

In due time we may see new software being released that can automatically work with PMEM in App Direct Mode without needing costly modifications. When that happens, these software will probably have intelligent methods to detect whether the system has PMEM or regular DRAM, and thus will be compatible to work on both new and old servers. This extra processing may cause the software some miniscule delays, but that is made up for by the massive performance benefits of PMEM in App Direct mode.

In App Direct Mode, the DRAM capacity is not invisible as in Memory Mode, so for example if you have a server with 512GB of PMEM modules and 64GB of DRAM, the OS will register the total of both (576GB) as available, while being aware of there being two types with different properties.³⁷

App Direct Mode finally lets us use PMEM for its intended purpose: as a persistent, non-volatile memory. This opens a wide world of opportunities to improve the performance of standard applications by redesigning them to take advantage of that fact. For most of computing history, the model has been that data can only be secure and persistent in slow, cumbersome storage such as tape, HDD, or even the comparatively good SSD. However, using SCM (or PMEM) now enables persistence of data at memory speeds.

The major advantage of applications finally having a persistent memory is because of the obvious benefits in recovery from power outages. Until now, programmers of critical systems such as in the financial sector have to take a number of steps to check and re-check that a power failure would not lose a single transaction. The rigorous process to do this involves keeping the volatile memory alive using batteries, flushing the data to persistent storage (SSD or HDD), read it to ensure it matches, then writing a flag to indicate that the data is valid and read that back to ensure the valid flag is valid! Finally, after all this is done the transaction can be said to be complete and the data is safe. This process takes 4 IO's to do any single write, but so far was the only way to be sure that no transaction is lost from the memory.

Now we can see how attractive a persistent memory would be to these critical applications. The persistence means that all transactions are safe, with no data loss, ever. The speed boost comes from eliminating the writing and validating steps, plus the inherent speed of the PMEM itself. Even greater improvements are possible if PMEM is accessed as memory instead of as storage which entails passing through the file system.³⁶

The Storage Networking Industry Association (SNIA) first published its Nonvolatile Memory (NVM) Programming Model (NPM) in 2013. The SNIA NPM is a framework that enables applications to take advantage of nonvolatile memory (aka persistent memory) in any system that contains it³⁹. Applications can access the PMEM using standard OS storage calls passing through the file system, in the same method that slower storage devices such as SSDs and HDDs are accessed. This model was developed to address the proliferation of new NVM technologies. A programming model is necessary to coordinate the community of NVM producers and consumers to advance together through the coming architecture paradigm changes.

According to the SNIA model, a PMEM can be accessed in five different ways, starting with the previously mentioned main division of block-addressable Memory Mode or byte-addressable App Direct Mode.

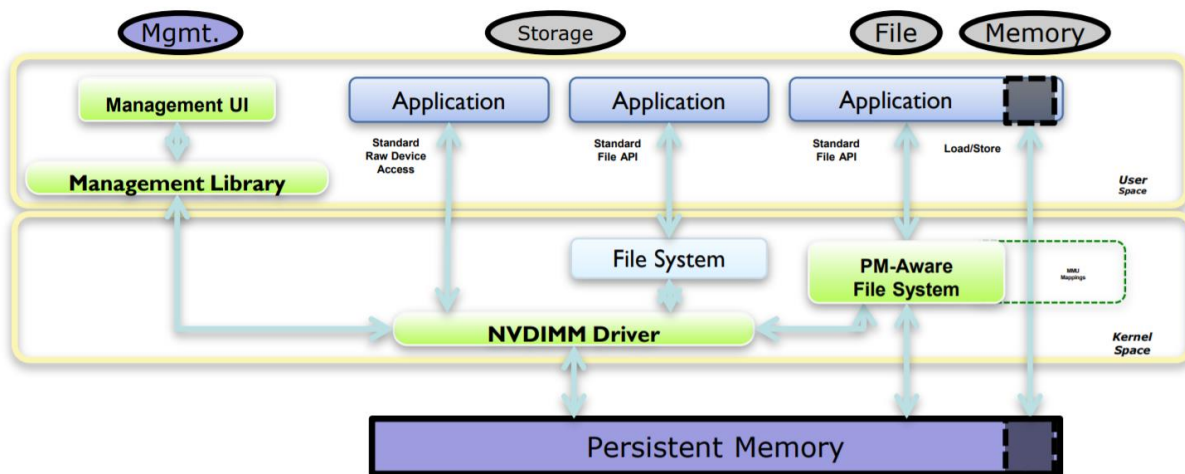


Figure 19: SNIA NVM Programming model: Exposing Persistent Memory to Applications [Ref. 41]

App Direct Mode has three divisions within it; Memory Access, Raw Device Access, or via a File API. The File API access is further divided into two sub-options; via a File System or via an NVM-aware File System.⁴⁰

A quick run-down of these options:

- **Memory Mode:** A non-persistent, block-addressable mode that presents the PMEM as a normal DRAM
- **App Direct Mode – Raw Device Access:** The application reads and writes directly to the PMEM driver in the host OS, which accesses the PMEM. This is faster than going through a file system interface but is not as direct as Memory Access and therefore slower than if the IO's go straight to the DIMM with no interruption. This mode is compatible with current file systems.
- **App Direct Mode – File System:** The PMEM address space is accessed by the application's file IO calls utilizing a file system API. The file system program then passes through the NVDIMM driver to the PMEM below. The file system overhead causes this access type to be slower than both raw device access and memory access modes.
- **App Direct Mode – NVM-aware File System:** This requires a modification to the file system to make it PMEM-aware, which will make it run faster than a traditional file system.

- **App Direct Mode – Memory Access:** In this mode the application uses memory access commands, i.e. load and store instructions, to directly access the PMEM’s address space with no drivers or software getting in the way (Load and store commands are those used to move data between CPU registers and memory). This is the fastest way an application can use PMEM.

Application Direct Mode in Action: SAP HANA using App Direct Mode

SAP HANA is one of the first major database platforms specifically optimized for operation with Intel Optane DC persistent memory, made possible by years of collaboration and cooperation between SAP and Intel. SAP HANA is optimized to leverage the unique characteristics of Intel Persistent Memory.⁴²

Unlike most applications that depend on the OS for memory allocation and management, SAP HANA is an In-Memory Database, and thus is designed to have more control over its memory management, to guarantee a higher level of optimization. This made SAP HANA an excellent candidate for seamless integration with PMEM.

SAP HANA automatically detects the presence of the Intel Optane PMEMs and adjusts itself by automatically knowing which data structures would most benefit from placing on persistent memory, and which should stay in DRAM. This is only made possible by using the **Application Direct mode** of Intel Optane’s persistent memory.

SAP HANA controls what is placed in Persistent Memory and what remains in DRAM.

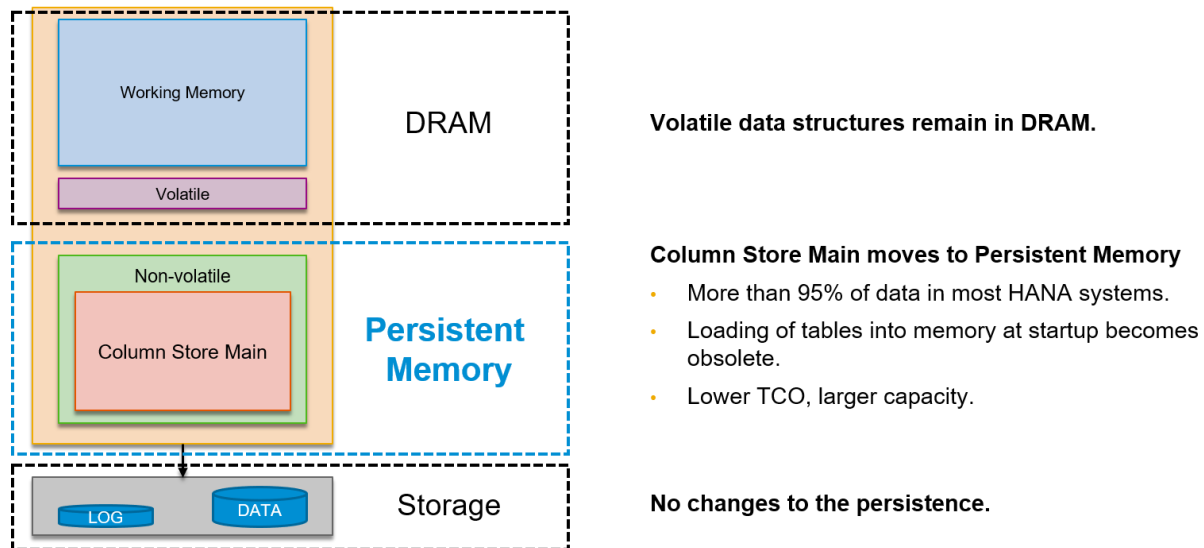


Figure 20: Placement of Persistent Memory in SAP HANA architecture [Ref. 43]

Using Application Direct mode, **persistent memory** is perfectly suited for placement of non-volatile data structures, as is obvious even in the name itself. This makes SAP HANA’s column store main data structure an excellent fit for PMEM. This data structure is optimized for compression, which makes it stable and non-volatile. The main store usually contains the majority (90%) of a typical SAP HANA database footprint, which makes the most of its placement on the large Intel Optane device. In addition, it is rarely reconstructed or modified during the delta merge, which does not happen often in the first place.

This made using PMEM a perfect fit for SAP HANA’s architecture. SAP HANA make a distinction between write-optimized delta and read-optimized main stores, and the unique characteristics of each makes them a perfect fit for placement on DRAM and PMEM, respectively.

SAP tested a configuration based on a 6 TB instance of SAP HANA and discovered that adding Intel Optane Persistent Memory results in a huge improvement in startup times, which improved by a factor of 12.5 – from 50 minutes with regular DRAM to just 4 minutes with persistent memory, even when taking data loading time into consideration. This is excellent news for administrators scheduling planned downtime events (for example due to an upgrade), since now they only need a few minutes instead of almost an hour. Minimizing business downtimes by a magnitude such as this is usually only possible by using expensive features like SAP HANA system replication.

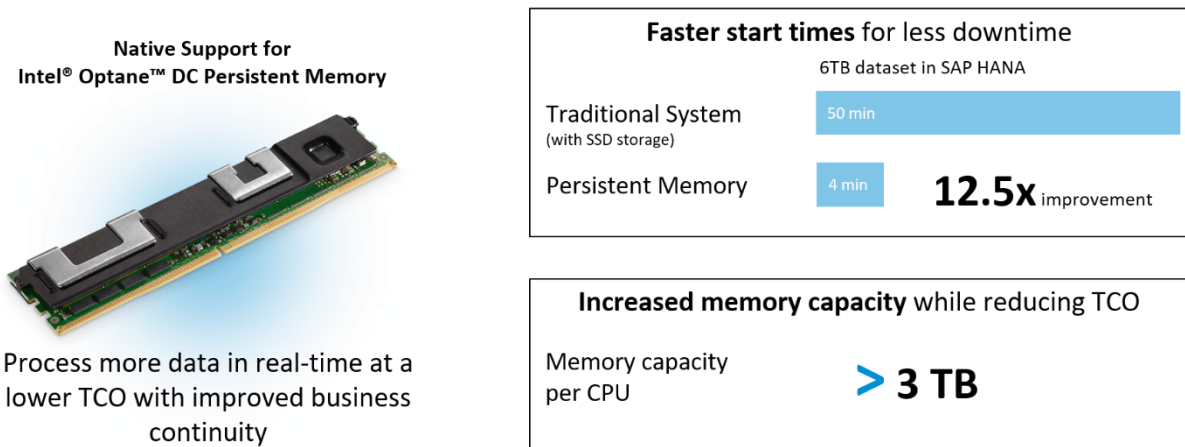


Figure 21: SAP HANA performance improvement by using Intel Optane PMEM in App Direct Mode [Ref. 43]

5. Conclusion

Storage-class memory (SCM) is a new class of storage that is able to retain its contents like NAND flash memory but at the speed of DRAM. It takes advantage of the gap in the market for a non-volatile memory solution with large scale and better cost. Many companies are investing heavily to capture this market potential, with advancements in PCRAM, MRAM and RRAM technologies some of the most promising candidates.

The ability of SCM to work as storage or as memory gives vendors the flexibility to deploy it in a variety of ways on servers and storage arrays. This article highlighted three use cases that we believe have the greatest chance of wide adoption: Caching on Storage Arrays, Tiering on Storage Arrays, and Persistent Memory on Servers.

SCM is projected to be adopted on the back-end in almost all major storage arrays within 2020, with adoption on the server side coming soon after. This shows that SCM is no longer a mere pipe dream or science fiction; it is the next stage of flash evolution, with strong support from major technology players. With its unique combination of speed, low cost, and high performance, SCM promises a transformation in how to handle read-intensive emerging applications such as AI workloads, bio-informatics, Big Data analytics, IOT analytics, in-memory databases, and high-performance computing.

Appendix: Terminologies

Useful definitions to know:

- **NVMe = NVM Express**: I/O interface to PCI Express (PCIe) SSDs
- **NVMe over Fabrics (NVMe oF)**: NVMe for networked fabrics (e.g. Fibre Channel, Ethernet)
- **Non-Volatile Memory (NVM)**: Memory media that persists data without power.
 - In this paper, this is any post-magnetic media such as NAND flash, 3D X-Point, resistive RAM, etc.
 - Media refers to the underlying non-volatile technology that stores data
- **Storage Class Memory (SCM)**: New low latency NVM subset, e.g. Optane, Z NAND flash
 - Block read/write (NVMe) or bytes load/store (DDR) interface
- **Persistent Memory (PMEM)**: Byte (load/store) addressable NVM
 - Usually implies Storage Class Memory (e.g. Optane), but includes flash-based NVDIMMs
- **Packaging/Interface**:
 - Drives (**SSDs = Solid State Drives**): I/O interface (read/write blocks), e.g. NVMe
 - Memory (**NVDIMMs = Non-Volatile DIMMs**): Memory interface (load/store bytes, cache lines)
 - **DIMM** = Dual In line Memory Module: Industry standard memory packaging, e.g. DDR4 interface

References

1. <https://www.dataversity.net/brief-history-data-storage/#>
2. <https://www.computerhope.com/jargon/p/punccard.htm>
3. <https://obsoletedia.org/scalable-linear-recording/>
4. <https://www.unm.edu/~tbeach/terms/storage.html>
5. <https://searchstorage.techtarget.com/definition/storage-medium>
6. <https://www.computerhope.com/jargon/r/ram.htm>
7. <https://themoryguy.com/why-do-memories-have-those-names/>
8. <https://techdifferences.com/difference-between-sram-and-dram.html>
9. <https://www.snia.org/sites/default/files/PM-Summit/2019/presentations/01-PMSummit19-Hady-Keynote.pdf>
10. <https://blocksandfiles.com/2018/11/28/2019-the-year-of-storage-class-memory/>
11. Meena et al., *Overview of emerging nonvolatile memory technologies*, *Nanoscale Research Letters* 2014, 9:526, <http://www.nanoscalereslett.com/content/9/1/526>
12. Chen et al., *Emerging Nanoelectronics Devices, First Edition. 2015 John Wiley & Sons, Ltd. Published 2015 John Wiley & Sons, Ltd., p 498-509*
13. Raoux et al., *Phase change materials and phase change memory*, August 2014, *MRS Bulletin*, 39(08):703-710
14. Zhang et al., *Memory materials and devices: From concept to application*, <https://onlinelibrary.wiley.com/doi/pdf/10.1002/inf2.12077>
15. <https://www.intel.com/content/dam/www/public/us/en/documents/technology-briefs/optane-reimagining-data-center-tech-brief.pdf>
16. <https://themoryguy.com/intels-optane-dimm-price-model/>
17. Hudec et al., *3D resistive RAM cell design for high-density storage class memory - a review*, http://cpanel-199-19.nctu.edu.tw/~thhou/publication/2016_SCIS_59_061403.pdf
18. <https://www.crossbar-inc.com/assets/resources/white-papers/T-Series-Product-Brief.pdf>
19. <https://semiengineering.com/what-happened-to-reram/>
20. <https://www.hardwarezone.com.sg/tech-news-intel-and-micron-announces-3d-xpoint-first-new-class-memory-25-years>
21. <https://itpeernetwork.intel.com/intel-optane-technology-and-dell-emc-powermax-the-revolution-in-enterprise-storage/%23gs.u8pjuh>
22. Wu et al., *A Survey on STT-MRAM Testing: Failure Mechanisms, Fault Models, and Tests*, <https://arxiv.org/pdf/2001.05463.pdf>
23. <https://www.everspin.com/nvme-storage-accelerator-series>
24. <https://www.everspin.com/spin-transfer-torque-mram-technology>
25. <https://www.tomshardware.com/reviews/3d-xpoint-guide,4747-8.html>
26. <https://www.networkworld.com/article/3026720/the-next-generation-of-storage-disruption-storage-class-memory.html>
27. <https://blocksandfiles.com/2019/12/20/western-digital-sivaram-solid-state-storage/>
28. <https://www.architecting.it/blog/max-data-application-performance/>
29. <https://www.storagereview.com/node/7707>
30. https://www.theregister.co.uk/2018/12/21/scm_power_to_the_max/
31. https://www.dellemc.com/ro-ro/collaterals/unauth/white-papers/products/storage/h17118_dell_emc_powermax_family_overview.pdf

32. <https://searchstorage.techtarget.com/tip/2-storage-class-memory-use-cases-to-consider>
33. <https://dcig.com/2019/03/caching-vs-tiering-with-storage-class-memory-and-nvme-a-tale-of-two-systems.html>
34. <https://www.hpe.com/us/en/storage/hpe-memory-driven-flash.html>
35. <https://searchstorage.techtarget.com/opinion/Are-PMEM-and-storage-class-memory-worth-the-price>
36. <https://thesdguy.com/intels-optane-two-confusing-modes-part-1-overview/>
37. <https://searchstorage.techtarget.com/feature/An-in-depth-look-at-Optane-persistent-memory-and-its-operating-modes>
38. <https://octo.vmware.com/vmware-and-intel-optane-dc-pmem/>
39. https://www.snia.org/tech_activities/standards/curr_standards/npm
40. <https://blocksandfiles.com/2019/07/02/optane-dimm-access-modes/>
41. <https://www.snia.org/sites/default/files/PM-Summit/2019/presentations/03-PMSummit19-Rudoff-NVM.pdf>
42. <https://www.sap.com/documents/2019/10/6a467fcb-6d7d-0010-87a3-c30de2ffd8ff.html>
43. <https://blogs.sap.com/2018/12/03/sap-hana-persistent-memory/>

Dell Technologies believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED “AS IS.” DELL TECHNOLOGIES MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying and distribution of any Dell Technologies software described in this publication requires an applicable software license.

Copyright © 2020 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners.