WHITE PAPER

Data Center
Storage

Data, Data Everywhere – Storage on the Brink?

Why Intel® Optane™ technology is a key part of enterprise data strategy

Table of Contents

Businesses are deluged by data. And by vendors and commentators reminding them of the fact. It has become commonplace to talk about the exploding volume, variety and velocity of information – each time with iteratively more terrifying projections of the sheer quantity produced in the enterprise. The suggestion is that the systems required to store and secure all this data are on the brink of being overwhelmed if radical action is not taken.

This paper is a bit different. Here, we are going to look at this data deluge problem not only as a *storage problem*, but also as a data optimization opportunity. We are then going to zero in on one specific technology – Intel® Optane™ technology – which can play a significant role in making the most of this opportunity. Intel® Optane™ technology is the first really new memory technology since the 1970s (DRAM originating in 1966 and the precursors of solid-state drives in the early 1970s). It has the capability to change enterprise data performance and the economics of that performance, in fundamental ways as part of a modern enterprise data strategy.

The Need for an Enterprise Data Strategy

It is certainly true that enterprise data storage has challenges. Volume, variety and velocity are among them. But there are other challenges too, which are equally important, and taken together build the argument for a comprehensive enterprise data strategy.

Alongside the "three V's", IT organizations are challenged with making data useful in ways which have never previously been required. The increasingly business-critical nature of analytics and artificial intelligence (AI) applications means that data has to be accessible, interrogable and often near-real-time in ways which were not the case when today's storage architectures were designed. High Performance Computing (HPC) applications need to hold huge data sets in-memory, or in storage with performance as close to memory as possible. Real-time and streaming applications (for example in financial services) need data to be captured and held in ultra-fast storage, very close to the compute node. All of these needs affect directly how certain tiers of data are identified and stored.

As businesses need to do more and more important things, with more data, its location is also starting to become distributed. Internet of Things (IoT) applications – whether autonomous cars, smart cities or retail – need data to be stored and processed close to where it is produced. The data center is, to an extent, disaggregating, and data center-like technologies are appearing at the edge of the network, all requiring data capture, storage, and analysis.

Finally, we have cost. Storage total cost of ownership (TCO) is as important as ever, but it is increasingly difficult to calculate. It is not just about the cost of storing and securing data – vault-like – in cold or cool storage. It is also about the cost per operation of accessing and analyzing that data when it is needed. Benchmarks need to reflect changing real-world workloads, not simplified ideal states.

So, taken in combination, these factors argue for a comprehensive approach to a modern enterprise data strategy.

Building an Enterprise Data Strategy

IT must create an intelligent data strategy that not only efficiently and securely stores data, but also pre-positions the data for real-time or future application use. Understanding application and location requirements – speed and continuity of access, proximity to the compute node etc. – really means understanding the business processes and applications the data strategy must support. Alongside the data deluge, the other great commonplace of IT commentary is that 'IT must get closer to the business'. Nowhere is this truer than in data strategy. By understanding the applications – often built around analytics or AI – which drive competitive advantage, and the ways in which those applications may move towards the edge, IT can build a data strategy which directly impacts business success.

Intel believes that a successful modern data strategy is software-defined, based on application requirements, and intelligently tiered. This paper is going to focus on data tiering, and Intel Optane technology's place in that process.

The strategic approach to data tiering is as follows:

- *Understand current application data requirements:* What is the relationship between data and the applications that use it. How is the data consumed? What are the latency and availability requirements? Who needs access and how? Which applications are business-critical?
- *Project future data requirements:* Which future applications will make challenging data demands – either performance, location, or cost? What will be the effect of the long-term business plan on data requirements?
- *Create a data hierarchy:* Data lives everywhere, both inside and outside an organization. Most application strategies will require varied types of data. It is important to understand the relationships between the various databases, both structured and unstructured. Modeling the data provides a hierarchy across an organization and is often a success factor in the execution of an analytics strategy.
- *Consider enterprise-level impactors:* A data strategy driven by business objectives will need to address factors impacting both the storage environment and analytics strategies. These include such things as data availability, data privacy, data retention, and compliance with data-centric regulations that the enterprise may need to conform to, data sourcing, and data cleanliness.

At the end of this process enterprises should have a 3D picture of their current and future data hierarchies, and the ability to start tiering them. Most IT departments think of data in terms of temperature. 'Cold' data is almost never accessed, but needs to be securely stored. At the other end of the scale, 'hot' data needs to be frequently, rapidly, and securely accessed. And, of course, data needs to be able to move through the tiers according to application and security needs.

Many businesses store 'cold' or 'cool' data in tape storage or in the cloud. Warm(ish) data can be happily stored on disk. The issue is with very warm and 'hot' data. Modern applications need more and more data, and they need it faster. So, to return to the beginning of this paper, the data deluge should have a fourth vector – volume, velocity, variety, and heat. More and more data is – or in an ideal world would be – hot. This causes real issues – both cost and technical – for IT, and this is where Intel Optane technology comes in.

Intel® Optane™ Technology: Disrupting Traditional Data Tiering Models

Traditionally, data has been tiered and stored on media of increasing cost and performance according to the access needs:

- Tape is cheap and slow
- Hard-drive disks are moderately priced and moderately slow
- Solid-state drives (SSDs) are more expensive and faster
- DRAM is extremely expensive and fast
- Processor cache is the most expensive and fastest (but severely limited in size).

This model has been somewhat disrupted at the low end by the use of cloud storage for cold and cool data by some enterprises, but at the warm/hot end, the trade-off between cost and performance remains acute.

Intel Optane technology is an entirely new technology which, while not as fast as DRAM, is fast enough to act as an extended memory pool and more responsive than even the highest performing NAND SSDs. This means that, for many applications, Intel Optane technology can replicate in-memory performance and therefore provide a much better high-performance solution for warm or hot data than disk, and a much cheaper solution than DRAM.

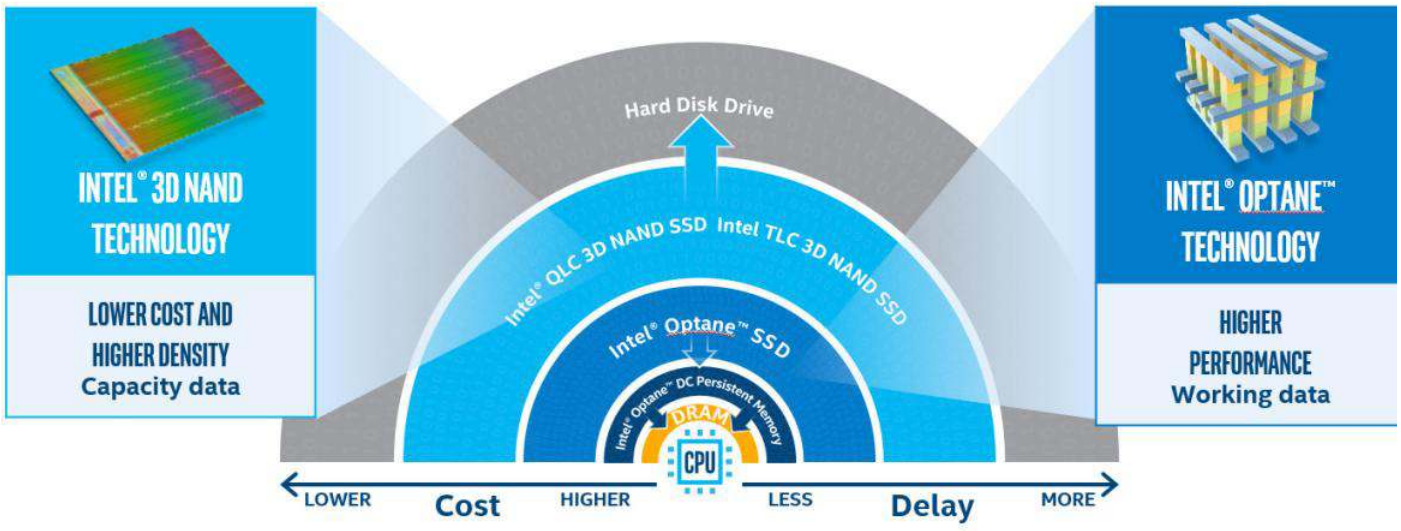

Figure 1. Intel® Optane™ technology introduces a new layer to the storage infrastructure

Effectively, Intel Optane technology combines the attributes of memory and storage with low latency, high endurance, outstanding quality of service and high throughput. This creates a new data tier, and one which is especially useful for the applications which are driving competitive advantage in today's enterprise environment: advanced analytics, AI, and HPC.

In today's architecture, blurring the boundaries between storage and memory means accelerating fast storage and extending cache. Intel Optane technology increases scale per server and reduces transaction costs. Combined with the latest Intel® Xeon® Scalable processors, Intel Optane SSDs enable bigger and more affordable data sets to gain new insights from larger memory pools. Many applications will benefit from using Intel Optane SSDs as a fast storage or cache device. They deliver the performance necessary to enable a new, higher performance cache or tier, as well as presenting the ultimate low latency for use as direct attached storage with the most demanding applications or services. In these applications, Intel Optane SSDs will vastly accelerate performance, break storage bottlenecks, improve workload scaling, and reduce the total cost for deployments. The ability to use a larger-capacity and lower-cost device, compared to DRAM, will create opportunities to save money by replacing some of the DRAM set, or gain new insights by growing data set sizes and complexity by augmenting DRAM to grow into significantly larger memory pools.

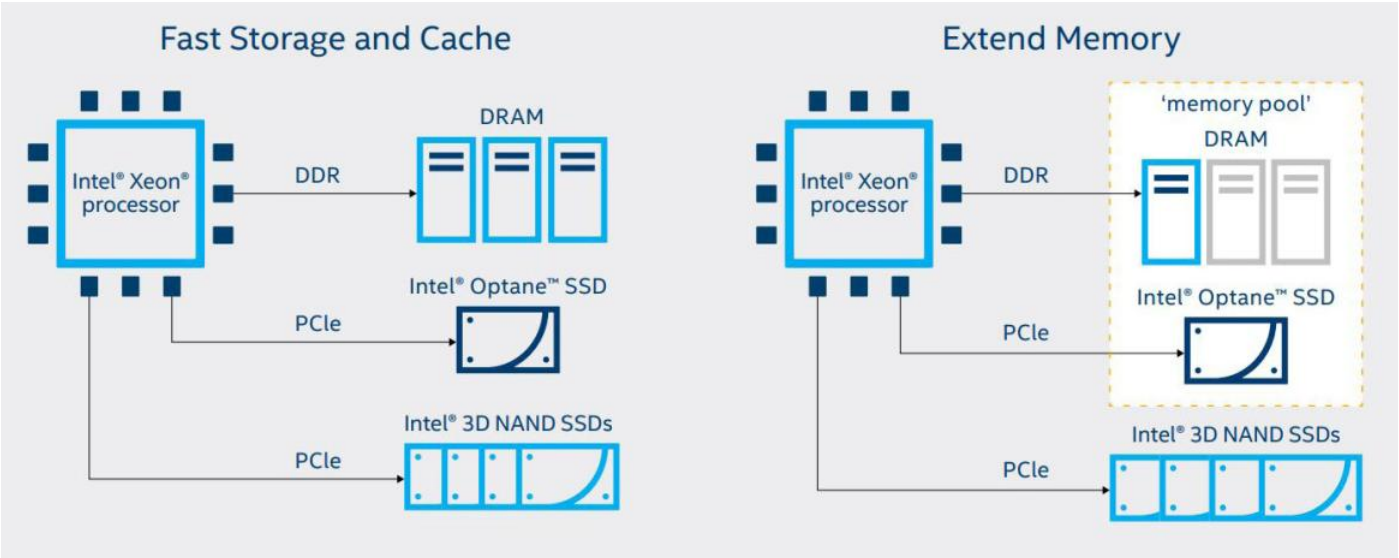

Figure 2. Intel® Optane™ technology enables faster storage and cache, and helps extend memory

Intel® 3D XPoint™ Technology – The Heart of Intel® Optane™ Technology

Intel Optane technology has at its heart a completely new approach to memory, based on massive advances in materials science, called Intel® 3D XPoint™ technology (pronounced cross point).

Intel 3D XPoint technology slices sub-microscopic layers of material into columns, each containing a memory cell and a selector. It connects the columns using perpendicular wires with an innovative cross-point structure that enables memory cells to be individually addressed by selecting one wire on top and one underneath. These memory grids can be stacked three dimensionally to increase density. Unlike DRAM, which requires a transistor for each memory cell – negatively impacting size and cost – each Intel 3D XPoint memory cell can be written to or read simply by varying the voltage in the wires, completely removing the need for transistors.

The state of the memory cells is persistent, meaning they keep their values even if they are powered down. Together, these capabilities provide the fast, dense, non-volatile characteristics which make Intel Optane technology perfect for blurring the lines between storage and memory.

For some workloads, Intel Optane SSDs increase the memory pool by up to 8x[1] and displace DRAM up to 10:1[2]. Over time, Intel Optane technology will be available in other form factors, which will unleash the creativity of the ecosystem to reshape further performance computing as application and location demands evolve.

Intel's Storage Portfolio

Intel® storage solutions are built on standards-based technologies, which reduce risk and provide flexibility to implement the storage architectures of your choice across multiple vendors or open-source solutions.

In combination with Intel Optane technology, Intel® 3D NAND SSDs transform the economics of storage infrastructures with unprecedented balance of scalability, performance and flexibility of choice, accelerating your data-intensive applications, reducing transaction costs for latency-sensitive workloads, and improving overall data center efficiency.

Benchmarking Intel® Optane™ Technology: Real-World Benchmarks Reveal Extraordinary Enterprise Data Center Performance

Storage performance is often evaluated at a high queue depth that is essentially the number of concurrent requests for data from storage at any one time. Most NAND SSDs are measured at their performance peak, which is often at a queue depth of 128 or even higher. Additionally, the performance demonstrated on NAND SSD is typically either 100 percent read or 100 percent write workloads. These are not common workloads. The most typical workloads from enterprise applications are both mixed reads and writes, with queue depths between 8 and 12.

Many enterprise applications - particularly operational databases, storage caches, log files and similar performance- or mission-critical applications - require responsiveness. This means it is common to limit the scale of work presented to any one storage device. In fact, it is not common for these types of workloads to reach queue depths higher than between 4 and 12.

IIntel has tested Intel Optane technology using a very common mixed workload of 70 percent reads and 30 percent writes, which is often seen in databases. We are also showing queue depth ranging from 1 to 16 – outstanding I/Os at the storage device. In this case, the more performance the NAND SSD can deliver, the better. Intel Optane technology delivers better performance where it matters.

- Intel Optane technology performance at QD1 is 6x that of a high endurance NAND SSD (P4600) for 70/30 mixed workloads.[3]
- Intel Optane technology achieves maximum throughput at QD16, whereas most NAND-basede SSDs require queue depths of 100 or more to saturate performance.[4]
- Intel Optane technology maintains consistent read response times regardless of the write throughput applied to the drive. Average read response times remain below <30us while maintaining a 70/30 mixed read/write bandwidth of 23GB/s.[4]
- Intel Optane technology offers 99% read response time is up to 63x better than that of a high endurance NAND SSD under random write workload.[5]

Intel® Select Solutions offer workload-optimized configurations, including modern Intel® SSDs that are verified for Intel® Xeon® Scalable processor–based platforms and are designed to help accelerate storage modernization initiatives. When organizations choose Intel Select Solutions, they get the optimized performance that hyper-converged infrastructures need, without the time and hassle required to tune the stack.

Intel is fostering industry-wide collaboration and innovation to provide customers with solutions that deliver the performance and efficiencies today's businesses demand with storage-optimized platforms, ingredients, and software.

Intel® Storage Builders provides an open, collaborative environment for true storage innovation and advancement. The program brings together forward-looking companies, a broad Intel ecosystem, key foundational technologies, and strategic resources to empower our members to move storage forward.

Take the Next Step to Modernize Your Data Strategy Today

Enterprises need to address the data challenges of the future with a clear vision of their application and architectural needs, and how data fits in. Then, develop a clear roadmap to deliver the storage modernization which will support these enterprises' data strategy. The alternative is choking on data growth, escalating costs and missed business opportunities. A crucial part of the modernization strategy is data tiering, and Intel Optane technology – by blurring the lines between memory and storage – provides completely new tiering opportunities to support real-time and data-intensive applications.

Gain further insights to help address your own data tiering and storage needs:

- Product Brief: Intel® Optane™ Memory – Breakthrough Responsiveness
- Solution Brief: DataON and Intel Select Hyper-Converged Infrastructure (HCI) Maximizes IOPS Performance for Windows Server Software-Defined Storage
- Solution Brief: Intel® Select Solutions for VMware vSAN*
- Webpage: Intel® Optane™ Technology

Solution Provided By:

[1] 2-socket platform All-DRAM memory configuration hardware limited up to 1.5TB (assuming 24 DIMM sticks by 64GB). Intel® Memory Drive Technology software supports up to 12TB addressable space, while DRAM as a cache is only resulting in up to 8x capacity increase.

[2] Intel® Optane™ SSD + Intel® Memory Drive Technology (IMDT) configuration – 2 x Intel® Xeon® CPU E5-2699 v4 @ 2.20Ghz, Intel® Server Board S2600WT, 128GB DDR4 + 4* Intel® Optane® SSD (SSDPED1K375GA), CentOS 7.3.1611. All DRAM configuration – 2 x Intel® Xeon® CPU E5-2699 v4 @ 2.20Ghz, Intel® Server Board S2600WT, 768GB DDR4 CentOS 7.3.1611. Test - GEMM, segment size 18689, factor 22, threads 42.

[3] Intel-tested: 4K 70/30 RW Performance at Low Queue Depth. Test and System Confguration: CPU: Intel® Xeon® Gold 6140 FC-LGA14B 2.3GHz 24.75MB 140W 18 cores CD8067303405200, CPU Sockets: 2, RAM Capacity: 32G, RAM Model: DDR4, RAM Stufng: NA, DIMM Slots Populated: 2 slots, PCIe* Attach: CPU (not PCH lane attach), Chipset: Intel C620 chipset BIOS: SE5C620.8 6B.00.01.0013.030920180427, Switch/ReTimer Model/Vendor: Cable - Oculink 800mm straight SFF-8611 to right angle SFF-8611 Intel AXXCBL800CVCR, OS: CentOS 7.5, Kernel: 4.14.50(LTS),FIO version: 3.5; NVMe* Driver: Inbox, C-states: Disabled, Hyper Threading: Disabled, CPU Governor (through OS): Performance Mode. EIST (Speed Step), Intel Turbo Mode=Disabled, and P-states = Enabled. The benchmark results may need to be revised as additional testing is conducted. Performance results are based on testing as of August 7, 2018, and may not reflect all publicly available security updates. See configuration disclosure for details. No product can be absolutely secure. Tests document performance of components on a particular test, in specific systems. Differences in hardware, software, or confguration will affect actual performance.

[4] Intel drive evaluated - Intel® Optane™ SSD DC P4800X 375GB. Test and System Configuration: CPU: Intel® Xeon® E5-2687W v4 3.0GHz 30MB 160W 12 cores, CPU Sockets: 2, RAM Capacity: 32GB, RAM Model: DDR4 2133MHz, PCIe Attach: CPU (not PCH lane attach), Chipset: Intel C610 chipset, BIOS: SE5C610.86B.01.01.0024.021320181901, Switch/ReTimer Model/Vendor: Intel A2U44X25NVMEDK, OS: CentOS 7.3.1611, Kernel: 4.14.50, FIO version: 3.5; NVMe Driver: Inbox, C-states: Disabled, Hyper Threading: Disabled, CPU Governor (through OS): Performance Mode; EIST (Speed Step): Disabled, Intel Turbo Mode: Disabled, P-states = Disabled; IRQ Balancing Services (OS) = Off; SMP Affinity, set in the OS; QD1 utilizes I/O Polling Mode. Performance results are based on testing as of August 31, 2018 and may not reflect the publicly available security updates. See configuration disclosure for details.

[5] Source – Intel-tested: Average read latency measured at queue depth 1 during 4k random write workload. Measured using FIO 3.1. Common Configuration - Intel 2U Server System, OS CentOS 7.5, kernel 4.17.6-1.el7.x86_64, CPU 2 x Intel® Xeon® 6154 Gold @ 3.0GHz (18 cores), RAM 256GB DDR4 @ 2666MHz. Configuration – Intel® Optane™ SSD DC P4800X 375GB and Intel® SSD DC P4600 1.6TB. Latency – Average read latency measured at QD1 during 4K Random Write operations using FIO 3.1. Intel Microcode: 0x2000043; System BIOS: 00.01.0013; ME Firmware: 04.00.04.294; BMC Firmware: 1.43.91f76955; FRUSDR: 1.43. SSDs tested were commercially available at time of test. The benchmark results may need to be revised as additional testing is conducted. Performance results are based on testing as of July 24, 2018 and may not reflect the publicly available security updates. See configuration disclosure for details.

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more complete information visit www.intel.com/benchmarks.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. No computer system can be absolutely secure. Check with your system manufacturer or retailer or learn more at www.intel.com/content/www/us/en/architecture-and-technology/intel-optane-technology.html

Cost reduction scenarios described are intended as examples of how a given Intel- based product, in the specified circumstances and configurations, may affect future costs and provide cost savings. Circumstances will vary. Intel does not guarantee any costs or cost reduction.

All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest Intel product specifications and roadmaps.

1018/CAT/LD/PD 337640-002EN