SELECTION GUIDE

Data Center
Memory Solutions

Affordably Accommodate the Next Wave of Data Demands

How to reach beyond DRAM and NAND limitations with the ROI-boosting benefits of Intel® Optane™ persistent memory and SSDs

Executive Summary

As the volume of global data continues to rise exponentially, the world's cloud infrastructures, where much of that data will reside, must adapt before they become overwhelmed. Traditional memory and storage architectures often prove too costly or limited in performance for conventional real-time data demands.

Intel® Optane™ persistent memory and Intel® Optane™ SSDs help address the performance and capacity gap that exists between DRAM and NAND storage. By combining groundbreaking Intel® Optane™ technology media with new software, users can employ various configuration modes to achieve a range of benefits, including:

- Massively expanded total system memory capacities
- High virtual machine (VM) density, enabling more VMs per server and/or total server count consolidation
- High application performance partially due to less need for disk swapping
- New tiering possibilities from responsive storage
- Persistent memory for remarkably fast system recovery

Intel Optane persistent memory and SSDs provide an affordable solution for enabling data center and cloud applications with the capacity and performance needed to accommodate the coming generation of data demands.

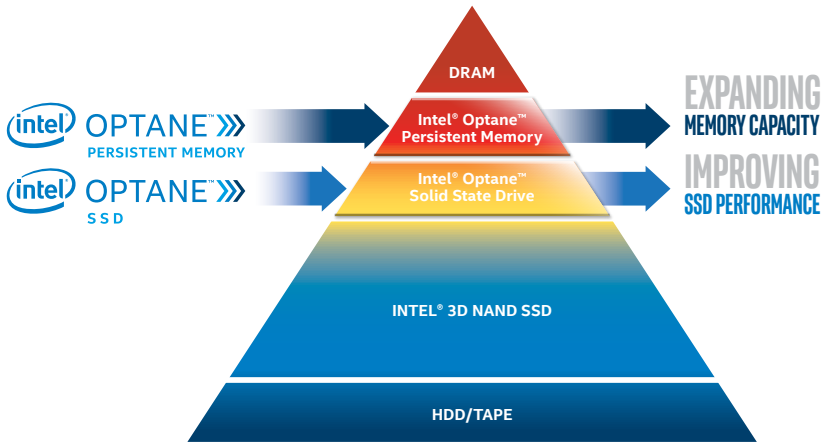

Figure 1. Between fast memory and high-capacity storage, a broad gap has held back storage capabilities in core data applications. Intel® Optane™ technology solutions now address those needs.

Business Challenge

To paraphrase Einstein, the technologies that brought us to today will not solve the problems of tomorrow. Organizations need a fundamental improvement in how they process data to keep pace with the rising volume, speed, and complexity of workloads. This need stems from a range of macro-level, industry-defining trends, starting with the explosive growth of global data. According to IDC, the sum amount of all digital data created will explode from 33 zettabytes (ZB) in 2018 to 175 ZB in 2025, and nearly half of this data will be stored in the cloud.[1] In the midst of this exponential increase and migration to network core, IDC expects that by 2025, 30 percent of data will be available in real-time to those who need it as it is created.[1] Thus, the global need focuses on both the capacity to store all this data as well as ensuring that storage has the speed to provide real-time responsiveness.

According to Cisco, many factors fuel this data explosion:[2]

- 5G rollouts beginning in 2020
- IoT connections will double (to 28.5 billion) by 2022
- Internet gaming traffic will expand by 9x from 2017 to 2022
- Video surveillance will grow by 7x from 2017 to 2022
- Virtual and augmented reality traffic will experience 65 percent CAGR over the same five years

From now until 2025 is startlingly little time to prepare the world to accommodate more than a quintupling of created and captured data. The strain on computing infrastructure will be significant. Imagine if the highways around your home needed to accommodate five times more traffic and, simultaneously, everyone wanted to drive much faster than before. We now face this situation in data centers.

All that data and much more will fuel rampant growth in data workloads, much of which will need real-time processing. Traditionally, addressing such speed demands meant costly measures, such as moving to in-memory databases. However, in-memory systems tend to be prohibitively expensive, in part because per-gigabyte DRAM pricing tends to remain high while module capacities are limited.

This "fast but expensive and capacity-limited" nature of DRAM has led to a fragmented memory/storage market represented by Figure 2.

"Hot" or "active" data needs to reside in memory to provide performance, but organizations struggle to get all the data they want into DRAM. "Warm," or nearline, data has more relaxed responsiveness requirements, so it can reside in storage. However, the data explosion affects all levels. Low-latency demands on nearline storage keep increasing, sometimes beyond the limits of what NAND SSDs can fulfill. It's time for a new technology to fill the performance and capacity gap in today's market so that we can meet the data needs of tomorrow.

The Intel® Optane™ Technology Solution

As Figure 2 implies, performance and capacity have traditionally been viewed as contrasting, almost opposing states, like black and white. In reality, though, they are more like frequency bands on a continuous spectrum. DRAM is not the fastest form of memory; we can peek inside the CPU at L1 through L3 caches, or even registers, to find faster. Similarly, hard drives and tape are not the slowest storage media. Memory and storage exist on a continuum. That continuum is typically viewed through two lenses, cost and latency, which tend to have an inverse relationship. As cost per gigabyte increases, access latency decreases, as shown in Figure 3.

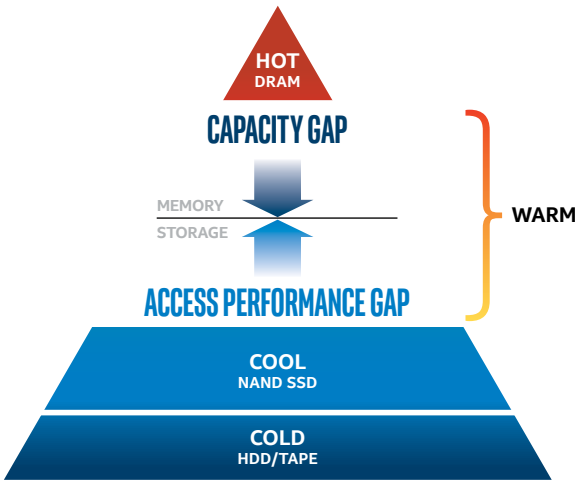

Figure 2. The pressures being placed on data centers have resulted in a need for more capacity in memory resources and a need for higher performance from "warm" storage. At times, pricing may limit the former while legacy technology constraints limit the latter.

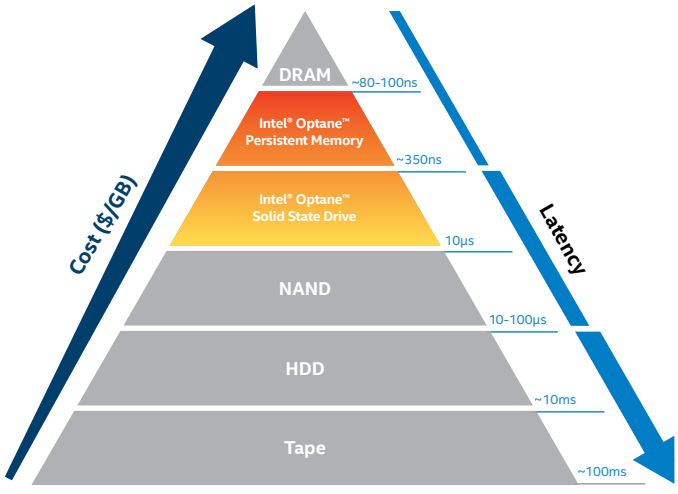

Figure 3. As the cost per gigabyte increases, access latency decreases.

Intel Optane technology media slots into the gap between capacity and performance on this memory/storage continuum. The medium uses a resistive memory technology with inherently fast performance qualities. With NAND, a cell must be erased before it can be written to, and movement of electrons through inhibiting insulation and onto a floating metal gate can make programming slow. Intel Optane technology allows for "write-in-place" programming. There's no need to erase the media before writing. The controller simply sets or resets the material, making it much faster than NAND.

Whether this underlying media is packaged as memory modules or SSDs, Intel Optane technology possesses several characteristics that make it versatile and needed.

Low Latency

DDR4 modules specify latencies in the 10 to 20 nanosecond (ns) range.[3] Using the same physical DDR4 packaging and memory bus, Intel Optane persistent memory offers latency of up to about 350 ns. This is a relatively narrow delta compared to the 10 to 100 microseconds (µs) of NAND SSDs. (Remember that 1,000 nanoseconds make one microsecond.) As shown in Figure 4, even across the chasm of a PCI Express (PCIe) storage bus, Intel Optane SSDs still provide markedly faster access times than nearly all NAND SSDs.[4]

Data Persistence

Like NAND storage, Intel® Optane™ media deployed as storage does not lose data in the event of power failure or cycling. Depending on the mode in which it is configured, this data persistence can also apply to Intel Optane media deployed as system memory.

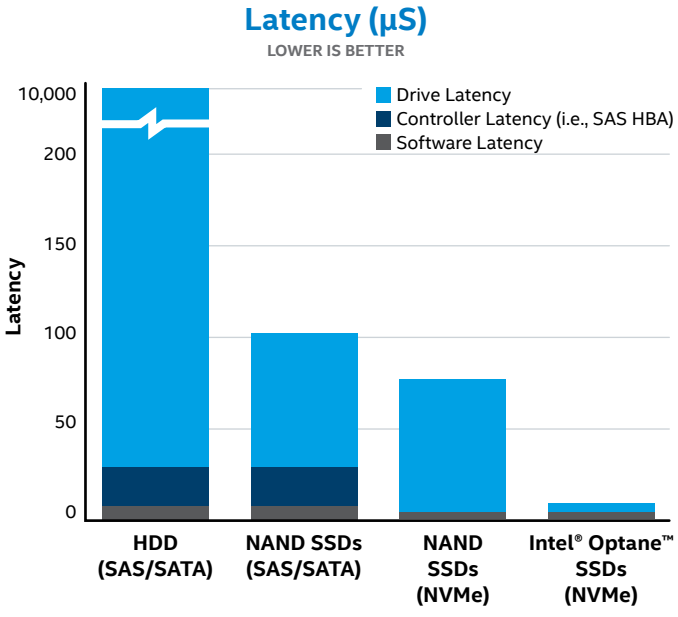

Figure 4. Moving to NVMe helps eliminate latencies caused by controller overhead, but most of the performance benefit seen in Intel® Optane™ SSDs derives from the media itself.[4]

High Endurance

Unlike NAND, Intel Optane media is a write-in-place technology, which eliminates NAND's requirement for erasing a cell before writing to it. This dramatically increases endurance. For example, the 1.5 TB Intel® Optane™ SSD DC D4800X specifies 82.1 petabytes written (PBW), as opposed to the 1.6 TB Intel® SSD P4610, which specifies 12.25 PBW. The ROI benefits, especially for data centers with write-intensive workloads, can be compelling.

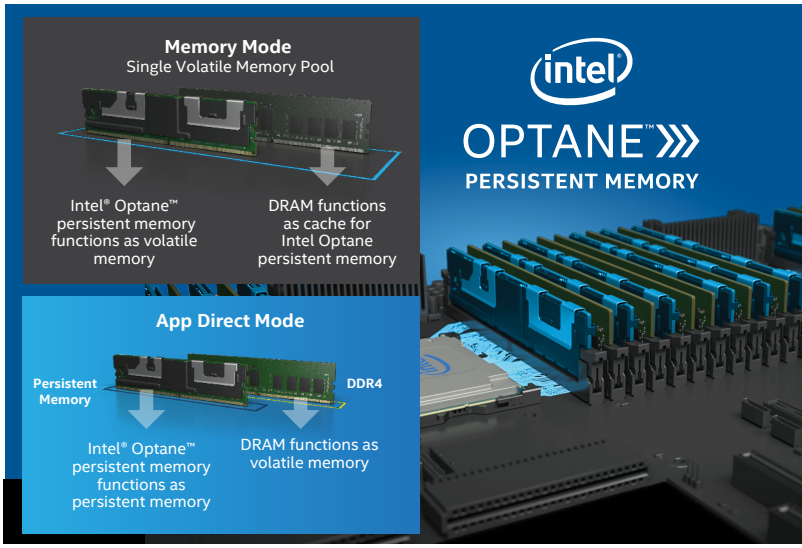

Intel® Optane™ Persistent Memory

Intel Optane persistent memory adheres to DDR4 physical module specifications. Thus, when used on platforms supporting 2nd Gen Intel® Xeon® Scalable processors, Intel Optane persistent memory modules simply drop in alongside conventional DDR4 memory. Intel Optane persistent memory currently ships in 128 GB, 256 GB, and 512 GB capacities. How they operate within the server depends on which mode they are configured in with Intel® software.

Memory Mode enables higher total system memory due to the fact that Intel Optane persistent memory modules feature much larger capacities than their DRAM counterparts. Intel Optane persistent memory modules run alongside DDR4 modules, with the latter functioning as a large L4 cache while the Intel Optane memory modules provide the total memory area accessible by users. Note that Intel® Optane™ memory is volatile in this mode. Memory Mode is ideal for applications and situations that are memory-constrained, such as very large database workloads or virtualized environments that need higher virtual machine (VM) and/or container counts.

App Direct Mode configures Intel Optane persistent memory such that data stays persistent within the modules. DRAM memory, however, remains volatile. Memory capacities are cumulative, enabling even higher system memory counts, although users will need software compatible with App Direct

Mode, and some additional programming may be needed for optimization. Under App Direct Mode, the operating system recognizes Intel Optane persistent memory and RAM as separate memory pools and can treat them differently. Memory can be persistent like storage, byte-addressable like memory, and cache coherent for supporting persistent memory beyond the local node. Because data persists, there is no need to reload data, which can yield benefits in cases such as fast database restarts and accelerated real-time analytics.

App Direct-Dual Mode is a sub-set of App Direct Mode that allows some Intel Optane persistent memory to work in Memory Mode while the rest serves in App Direct Mode. This is effective for letting applications use high-performance storage without I/O bus latency penalties.

To be clear, Intel Optane persistent memory does not replace DRAM; it complements DRAM. Intel Optane persistent memory significantly expands system memory and brings persistence to data that would have otherwise been relegated to slower storage media. Both facets of Intel Optane persistent memory can bring substantial ROI benefits. For example, memory-constrained VMs tend to underutilize CPUs. By expanding the memory available to constrained VMs, CPU utilization can increase for each VM, leading to higher total utilization. Alternatively, higher memory capacity can allow more VMs on each physical system, allowing organizations to run fewer total servers.

Intel® Optane™ SSDs

By placing the same Intel Optane technology media on x4 AIC, U.2, or M.2 form factor devices, a similar range of opportunities open for storage. As noted earlier, the PCI Express storage bus introduces additional latencies and overhead, but read/write performance (typically under 10µs) remains more than sufficient to enable a high storage tier with benefits that are either costly or impossible to reach with NAND SSDs.

We previously noted how Intel Optane media exhibits far greater endurance than NAND. Intel Optane storage also shows significantly faster and more reliable performance than NAND under sustained load. To show this point, Figure 5 illustrates the relationship between latency and IOPS, comparing Intel® 3D NAND SSDs and Intel Optane SSDs.

Interestingly, some of Intel Optane persistent memory's benefits can be realized with Intel Optane SSDs with Intel® Memory Drive Technology (currently available for Linux). This solution pairing uses software to create an extension of memory capacity by bonding Intel Optane SSDs with the DRAM pool, much like in Memory Mode with Intel Optane persistent memory modules. Intel Optane SSDs with Intel Memory Drive Technology provide a cost-effective way to massively expand system memory beyond regular DRAM limits. The technology enables very large workloads while maintaining relatively low DRAM capacities, potentially saving on hardware procurement and operations costs.

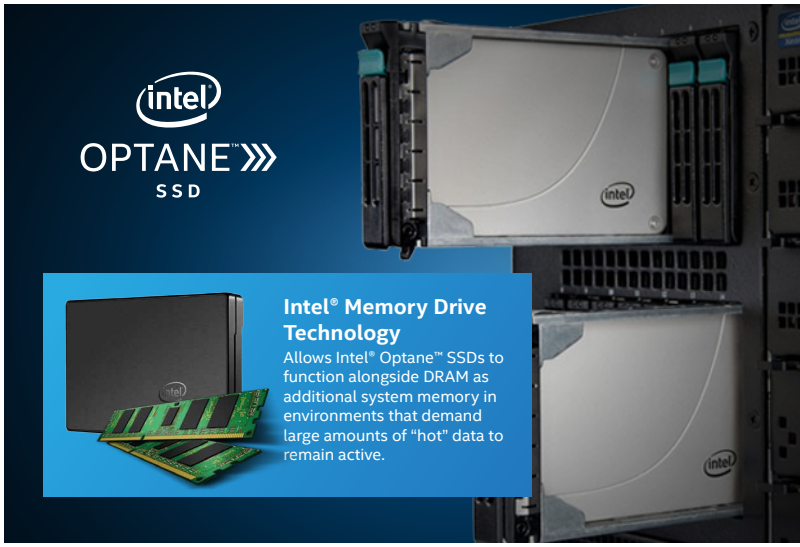

Because the SSD(s) communicate across the PCIe bus, some inevitable I/O latencies remain, which is why Intel Optane persistent memory remains preferable in performance-sensitive environments. However, Intel Optane SSDs with Intel Memory Drive Technology not only provide the outsized system memory sizes needed for many applications and environments, it is also supported by a larger number of today's CPUs.

Just as with all Intel Optane SSDs in general, Intel Memory Drive Technology works with Intel Xeon Scalable processors, Intel® Xeon® processor E5-x6xx v2 or later,[5] and Intel Xeon processor E7-x8xx v2 or later models. Note that data remains volatile under Intel Memory Drive Technology.

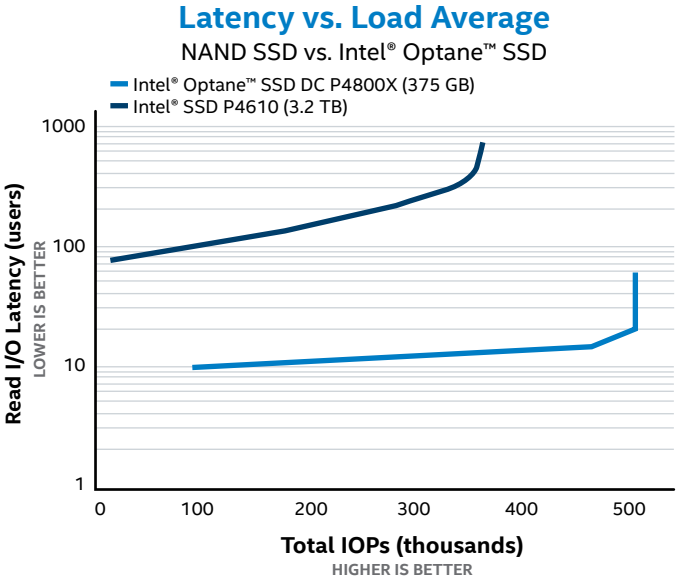

Figure 5. Intel® Optane™ SSDs maintain their low latency performance much more consistently than NAND alternatives as load levels increase.

More System Memory for All Workloads

Many use cases and target applications exist for Intel Optane persistent memory and SSDs, and more will continue to arrive as adoption continues. Both implementations of Intel Optane technology build on the foundation laid by Intel 3D NAND capacity storage in previous- and current-generation Intel Xeon processor-based servers. Intel 3D NAND SSDs establish a baseline for good data performance and reliability across a variety of enterprise workloads.

On platforms running applications that would benefit from lower latency access to storage, adding Intel Optane SSDs can provide better performance compared to 3D NAND SSDs (Table 1). Intel Optane SSDs can provide more than enough capacity and speed to serve in a buffer/caching role. Frequently accessed storage data can reside in Intel Optane SSDs and realize markedly faster response rates while also sparing main SSD capacity storage from unnecessary use and wear. This helps Intel Optane SSDs deliver higher total system performance and lower costs as the storage subsystem reaps the benefits of Intel Optane media's long endurance characteristics.

Additionally, select Intel Xeon processor-based platforms[5] can use Intel Optane SSDs to optimize total system memory capacities. By using Intel Memory Drive Technology, even previous platforms can jump from gigabytes to several terabytes of available memory. This can markedly increase efficiency of memory-constrained applications and workloads, enabling benefits such as larger datasets, higher virtual machine density, and physical server consolidation.

Compared to 3D NAND SSDs, organizations that deploy server platforms based on 2nd Gen Intel Xeon processors can realize the best range of benefits. Rather than expand system memory across the NVMe bus, high-capacity Intel Optane persistent memory modules enable system memory capacities that are several times larger than what is possible with DRAM alone. All this memory coexists on the DDR4 bus and reaps that channel's access time benefits, providing outsized capacity for very large datasets and higher performance (compared to reaching to Intel Optane media across the NVMe bus) for demanding applications that have been enhanced to be Intel Optane persistent memory-aware.

Intel Optane persistent memory modules on 2nd Generation Intel Xeon processors provide additional benefits through the use of App Direct Mode. With App Direct, organizations can optimize configurations for both performance and system memory capacity as workloads dictate. Ultimately, businesses that use 2nd Generation Intel Xeon processor-based servers and that need to expand memory capacities and/or use fast, non-volatile memory should add Intel Optane persistent memory.

Table 1. Any Intel® Xeon® processor-based server can benefit from the advantages of integrating Intel® Optane™ SSDs. However, for servers based on 2nd Generation Intel Xeon Scalable processors, an increasing number of applications can achieve the best range of benefits by using Intel Optane persistent memory.

	Platforms Optimized with Intel® Optane™ SSD	2nd Gen Intel® Xeon® Scalable Processors Enhanced with Intel® Optane™ PMem[5]
STORAGE		
SDS		
- Ceph Block/Object	★★★	
- Hadoop YARN Temp	★★★	
- VMware vSAN	★★★	
Hyper-converged (HCI)		
- MS Azure Stack HCI Storage Spaces Direct	★★★	
- Cisco HyperFlex	★★★	
INFRASTRUCTURE		
Memory		
- VMware ESXi		★★★
- Microsoft Hyper-V		★★★
- VMware VDI		★★★
- Redis Open Source Software KVM		★★★
HPC		
Scratch and I/O Nodes		
- HPC Flex Storage	★★★	

utilizes Intel® Optane™ SSD DC P4800X with Intel® Memory Drive Technology to enable use of SSD as storage or memory.

	Platforms Optimized with Intel® Optane™ SSD	2nd Gen Intel® Xeon® Scalable Processors Enhanced with Intel® Optane™ PMem[5]
DATABASE		
Caching/Persistence		
- SAP HANA		★★★
- Redis Labs		★★★
- Memcached		★★★
- MS-SQL	★★	★★★
- Aerospike		★★★
- KX		★★★
AI/ANALYTICS		
Real-time Analytics		
- SAS		★★★
Machine Learning		
- Apache Spark	★★	★★★
COMMUNICATIONS		
Content Delivery NW		
- Hot Video-On-Demand/ Live Stream		★★★

★★ – Better Solution
★★★ – Best Solution

Intel Optane SSDs target usage models and workloads with widespread presence in the data center marketplace, including hyper-converged and software-defined storage workloads such as vSAN, Ceph, and Microsoft Storage Spaces Direct. In fact, hyper-converged infrastructure is a key focus area and one of the early adopters of Intel Optane SSDs because their low-latency characteristics help to minimize I/O wait times. Reducing I/O wait helps to recover CPU cycles, which in turn can help to increase workload density. Ultimately, increasing density can enable a reduction in node count. Fewer nodes facilitate smaller hardware investment and can lead to reduced software licensing costs.

In the infrastructure space, database workloads such as memcached and Redis are used as a database. Extending memory with a cost-effective solution like Intel Optane persistent memory can do well since networking, rather than the memory subsystem, is the typical bottleneck. With more memory, users can achieve better cache hit rates and thus improve the overall SLA in an affordable manner.

For database uses, SAP/HANA is the current market leader. Intel Optane persistent memory not only delivers high performance at a comparatively low cost, but it also allows for larger memory capacities than possible with DRAM alone. The other unique advantage for this market is persistence in database memory.

Finally, in high-performance computing, Intel Optane SSDs with Intel Memory Drive Technology can be flexibly configured on the fly as either memory or storage. This can bring tremendous versatility to HPC cluster budgeting and design, reducing reliance on high-cost DRAM in large-capacity memory nodes.

Conclusion

A 2019 Flexera survey revealed that 64 percent of organizations will focus on optimizing their existing cloud use for cost savings in 2019.[6] (Note that the #2 cloud priority for organizations, at 58 percent, was moving more workloads to the cloud.) Many of these organizations running cloud applications could benefit from having the memory capacity to host larger data sets, but the high cost of DRAM and limited DRAM capacities impose barriers. Concurrently, organizations could offer more value and service variety to customers if their applications' storage was more responsive, but NAND SSD latency has been largely capped by inherent media limitations that could only be—in some cases—surmounted by pushing data into DRAM. Of course, that approach entails high cost per gigabyte.

Intel Optane persistent memory and Intel Optane SSD technologies give businesses a way to cost-effectively break beyond these barriers using product classes that never existed before. By enabling previously unattainable system memory capacities and introducing new levels of responsiveness and endurance to high-performance storage, Intel Optane technology delivers the affordable solutions data centers and cloud applications need to accommodate the coming generation of data demands.

For more information on Intel Optane technology, visit **intel.com/optane.**

Learn More

Use Cases and Related Applications

For a deeper dive into some of these use cases and related applications, investigate these Intel solution briefs:

- Ceph: Accelerate Ceph Clusters with Intel Optane SSDs.
- SAP HANA: SAP HANA 2.0 SPS 03 contains numerous innovations to the SAP HANA platform—and it is the first major database solution to support Intel Optane persistent memory.
- VMware vSAN: Agile storage systems rapidly enhance analysis and insight from stored data.
- Vexata Systems and Intel Optane SSDs: Vexata delivers enterprise storage arrays based on Intel Optane SSDs and NVMe/PCIe Intel® SSDs to enable transformative application performance for database and analytics workloads.
- Exploiting Intel Optane SSD for Microsoft SQL Server: New NVM-based devices provide unparalleled performance than Flash-based SSDs.

Related Content

If you liked this paper, you may also be interested in these related resources:

- Intel Optane Technology for Data Centers
- Intel Optane Technology FAQ
- Memory and Storage Technical Series: These resources can help system architects, engineers, and IT administrators better understand the limitations of traditional memory and storage options; how those limitations lead to performance and capacity gaps in your data center; and finally how Intel Optane technology is helping to fill those gaps with a new industry-disrupting architecture.
- Overview of the Intel Optane Persistent Memory Video
- Overview of the Intel Optane Persistent Memory Performance Video
- Configuring Intel Optane Persistent Memory for Best Performance Video

[1] IDC/Seagate "Data Age 2025" white paper. November 2018. seagate.com/www-content/our-story/trends/files/idc-seagate-dataage-whitepaper.pdf

[2] Cisco Visual Networking Index: Forecast and Trends, 2017–2022. cisco.com/c/en/us/solutions/collateral/service-provider/visual-networking-index-vni/white-paper-c11-741490.html

[3] wikipedia.org/wiki/DDR4_SDRAM#JEDEC_standard_DDR4_module

[4] 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 processor @ 3.0 GHz (18 cores), RAM 256 GB DDR4 @ 2666 MHz. Configuration – Intel® Optane™ SSD DC P4800X 375 GB and Intel® SSD P4600 1.6 TB. 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. Intel performance test results as of July 24, 2018 may not reflect all publicly available security updates. See configuration disclosure for details.

No product can be absolutely secure.

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 intel.com/benchmarks.

[5] See the Features At-a-Glance chart at intel.com/content/www/us/en/products/docs/memory-storage/solid-state-drives/optane-ssd-dc-p4800x-mdt-brief.html

[6] Flexera blog (formerly RightScale), February 27, 2019, "2019 State of the Cloud Survey." flexera.com/blog/cloud/2019/02/cloud-computing-trends-2019-state-of-the-cloud-survey/#The%20#1%20Priority%20in%202019%20Is%20Cloud%20Cost%20Optimization

Intel technologies may require enabled hardware, software or service activation. No product or component can be absolutely secure. Your costs and results may vary.

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 intel.com/benchmarks.

Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See backup for configuration details. No product or component can be absolutely secure.

Intel technologies may require enabled hardware, software or service activation.

The products described in this document may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Intel does not control or audit third-party data. You should review this content, consult other sources, and confirm whether referenced data are accurate.

Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation in the U.S. and/or other countries.

Other names and brands may be claimed as the property of others.

© Intel Corporation 0420/JGAL/KC/PDF 341076-002US