

## Achieve Greater Insight from Your Data

### Intel® Optane™ Persistent Memory 200 Series Enables Fast Tiered Memory, Delivering 32 Percent More Bandwidth on Average<sup>1</sup> with up to 6 TB Total Memory per Socket<sup>2</sup>



Over 90 percent of enterprises are in the midst of their digital transformation journey<sup>3</sup> to become data-centric businesses. Thus, they will need to capture, analyze, and secure increasing amounts of data.

Data growth is accelerating the demand for performant, data-intensive computing. As the demand for compute grows, memory capacity in the system typically needs to scale along with it.

Large pools of DRAM help accelerate computing with low latency, but DRAM is limited in capacity, is volatile, and expensive. Indeed, DRAM is becoming one of the highest priced components in modern server bills of materials. Alternatively, block storage is large, cheap, and persistent, but slow to bring data to the CPU.

Intel® Optane™ persistent memory (PMem) bridges the gap with an innovative memory technology that delivers a unique combination of affordable large capacity and support for data persistence. With 3rd Gen Intel® Xeon® Scalable processors and Intel Optane PMem 200 series, organizations can optimize capacity, performance, and cost of their workloads by enabling tiered memory and improved tiered storage.

PMem is compatible with 2-socket and 4-socket platforms and is helping to turn more data into actionable insights.

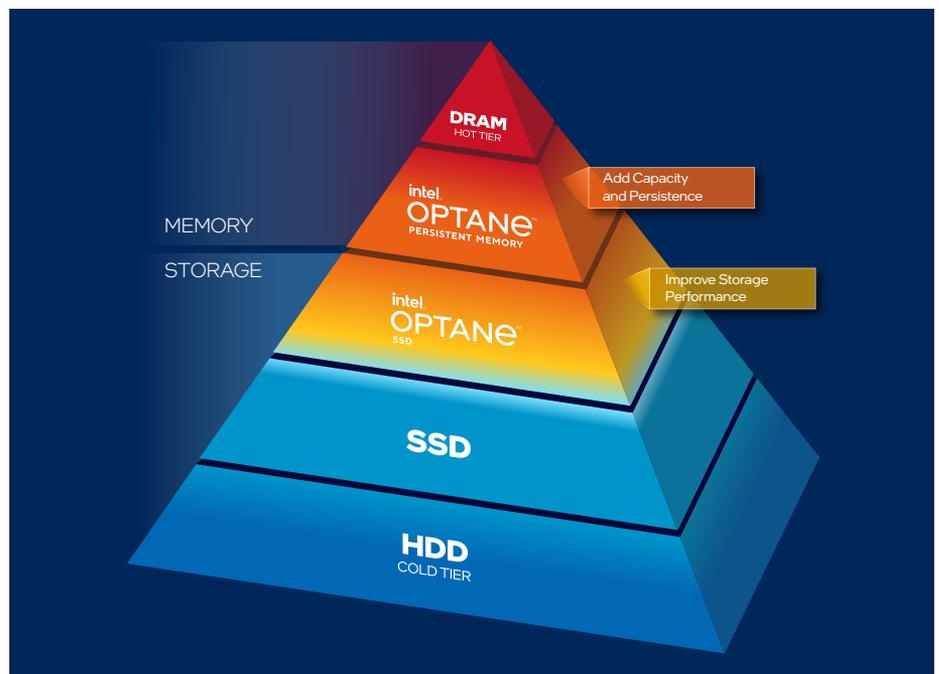


Figure 1. Intel® Optane™ persistent memory enables hierarchical architectures for high-performance, large memory computing.

## Optimize Workload Performance and Reliability with Intel Optane Persistent Memory 200 Series

Both DRAM and Intel Optane PMem 200 series sit on the DDR bus and can be used together, complementing system memory and enabling high-performance, tiered memory architectures. Available in 128 GB, 256 GB, and 512 GB modules, it offers both large capacity and persistence. PMem 200 series delivers an average of up to 32 percent more memory bandwidth than the previous generation.<sup>1</sup> These capabilities:

- Enable tiered memory and storage architectures that help improve overall system performance, cost, and efficiency
- Help accelerate large-memory computing by keeping more data closer to the CPU
- Accelerate restart times with reduced I/O by persisting data in memory and not requiring data reloads from storage
- Reduce power consumption for large-memory nodes

PMem brings important advantages to data-intensive, compute-intensive, and capacity-demanding workloads. From cloud to databases, in-memory analytics, virtualized infrastructure, content delivery networks, and more, these architectures can easily take advantage of large-scale, tiered memory using Intel Optane PMem. Enabling large memory pools with Intel Optane PMem helps accelerate time to insights, cost savings, and new revenue.

### A Powerful Foundation for Optimized Platforms

3rd Gen Intel Xeon Scalable processors and Intel Optane PMem 200 series create a powerful foundation for optimizing platforms for multiple types of workloads. With 3rd Gen Intel Xeon Scalable processors for 2-socket and 4-socket platforms, each memory channel can support one Intel Optane PMem 200 series module (six channels on 4-socket platforms and eight channels on 2-socket platforms). Each module draws only a maximum 15 watts of power, offering high capacity without high power demand and creating large-memory platforms to power data-centric businesses.

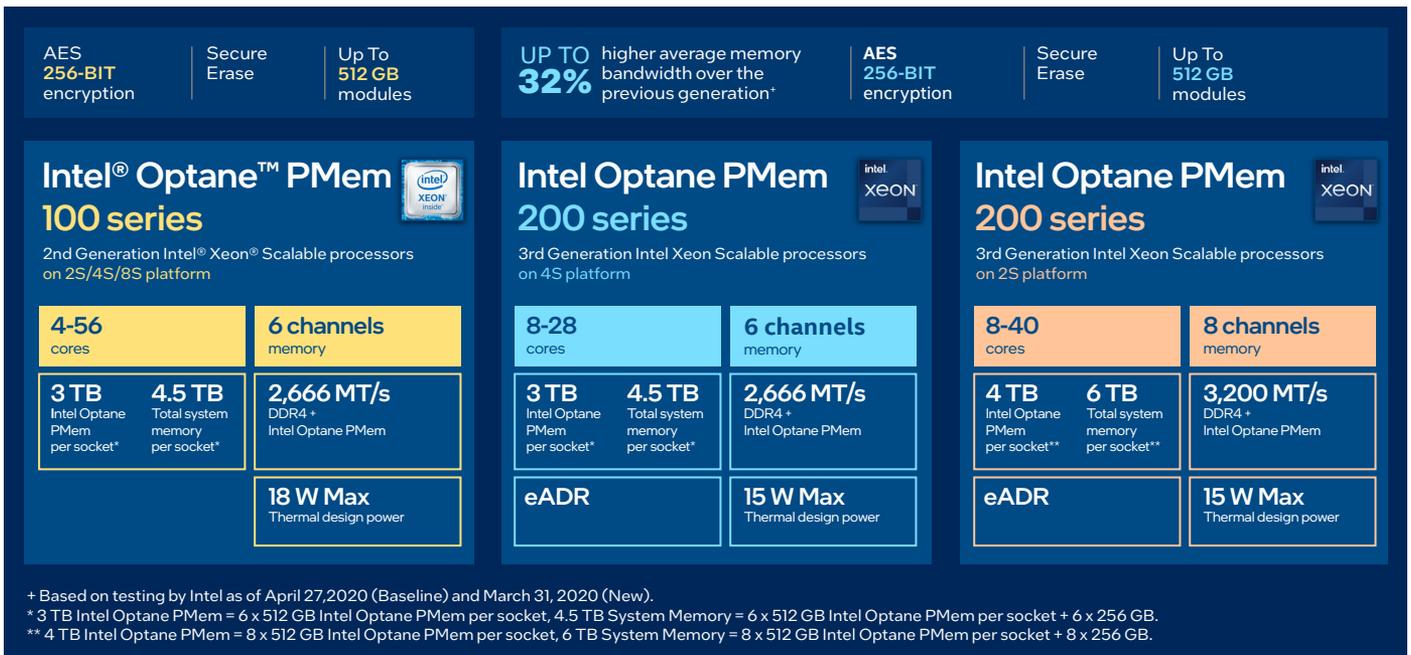


Figure 2. 3rd Gen Intel® Xeon® Scalable processors and Intel® Optane™ PMem 200 series create a powerful foundation for large-memory computing.

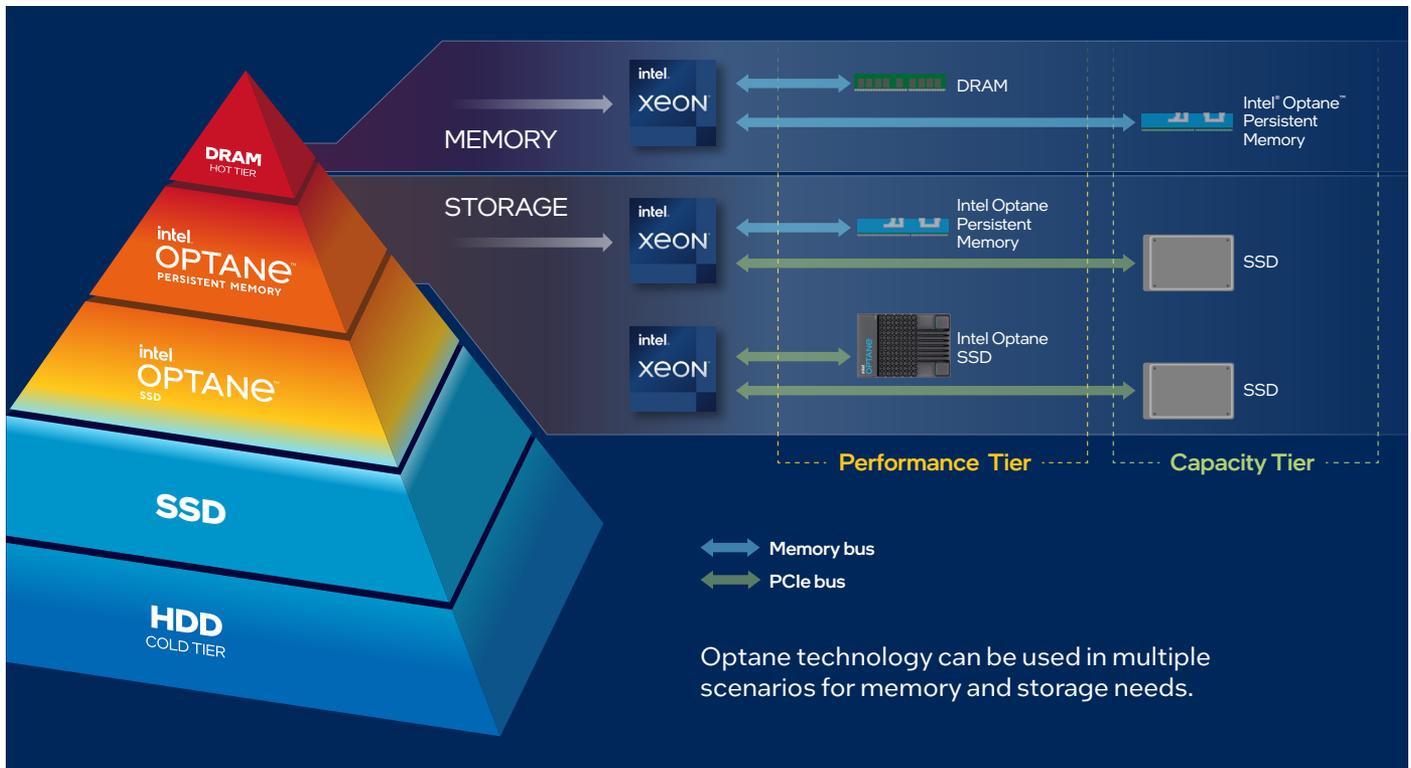


Figure 3. Intel® Optane™ Technology creates multi-tiered memory and storage hierarchies to enable optimized workloads.

### Tiered Memory and Storage

Intel Optane PMem offers system architects and application developers new options to create hierarchical memory and storage tiers to address workload challenges. A tiered approach can better optimize the resources of the platform for data access and transport. Programmers can leverage the speed and proximity of technologies closest to the CPU, while taking advantage of the overall capacity available in the system.

**Tiered memory** – For a tiered memory model, low-latency DRAM offers fast access to local memory, while Intel Optane PMem 200 series creates large capacity to store and protect data with DRAM-like speed. Depending on the application, PMem can be persistent or volatile.

The combination of DRAM with persistent memory:

- Allows fast in-memory computing on massive data sets
- Allows for consolidating more virtual machines on a platform in a virtual environment with large, performant capacity
- Reduces risk of losing critical data when the application is persistent-memory aware
- Speeds time-to-solution from large computations where intermediate results can be persisted and reloaded for final analysis

**Tiered storage** – With a tiered storage model, Intel Optane PMem 200 series creates a performance tier, delivering fast, byte-addressable access to most frequently needed data. Other technologies, such as SSDs, offer slower access to warm data storage in a capacity tier where speed of access is less critical.

Tiered memory and storage architectures help optimize speed, latency, capacity, and cost. Choosing the right combinations for each application can help optimize systems and their workloads.

### Affordable Large Capacity

When deployed as tiered memory to augment DRAM, Intel Optane PMem 200 series enables consolidation and reduced server footprint. These lead to lower software licensing costs, reduced power consumption, and other operational efficiencies.

**In-memory and large data computing** – Intel Optane PMem 200 series supports large in-memory databases with fast access to more data, without loading it from storage. Workloads processing massive data sets, such as scientific or data warehousing and analytics, can work continuously without repeatedly loading and storing data locally.

**Virtualized infrastructure** – Intel Optane PMem can offer greater memory capacity per socket than DRAM for virtualized infrastructures. PMem leaves more headroom for virtualizing future workloads that might require large memory capacity, eliminating the need to run them on bare metal.

### Drive Application Innovation and Explore New Data-Intensive Use Cases

With Intel Optane PMem 200 series, developers have new opportunities to optimize tiered memory and storage for their applications. They can drive innovation and capabilities using the same persistent memory programming model introduced with the first generation of PMem.

Rapid adoption is easy, and customers are able to take full advantage of PMem capabilities with a growing global ecosystem that includes the following:

- ISVs and OSVs
- Virtualization solution providers
- Database and enterprise solution vendors

- Data analytics vendors
- Open source solution providers
- Cloud Service Providers
- Hardware OEMs
- Standards bodies, such as the Storage Network Industry Association (SNIA), ACPI, UEFI, and DMTF

### Seamless Migration from Previous Generation

PMem 200 series is compatible with the software ecosystem already established for previous-generation Intel Optane PMem 100 series. Migrating to or adding new systems built on 3rd Gen Intel Xeon Scalable processors with Intel Optane PMem 200 series is transparent to software designed for the previous generation.

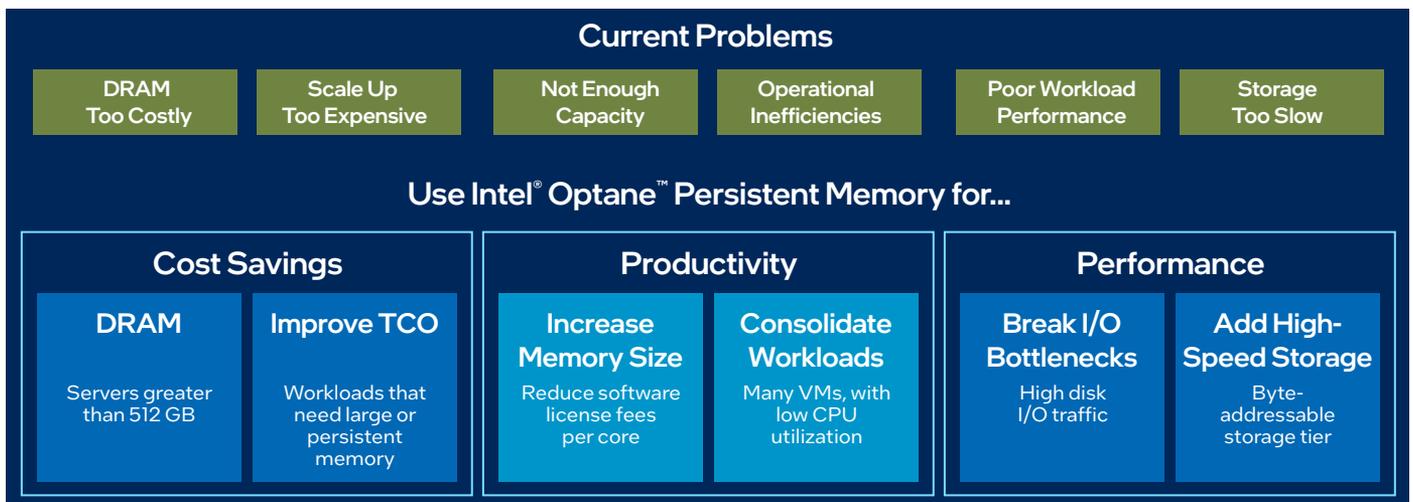


Figure 4. Intel® Optane™ PMem 200 series solves several key challenges in computing today.

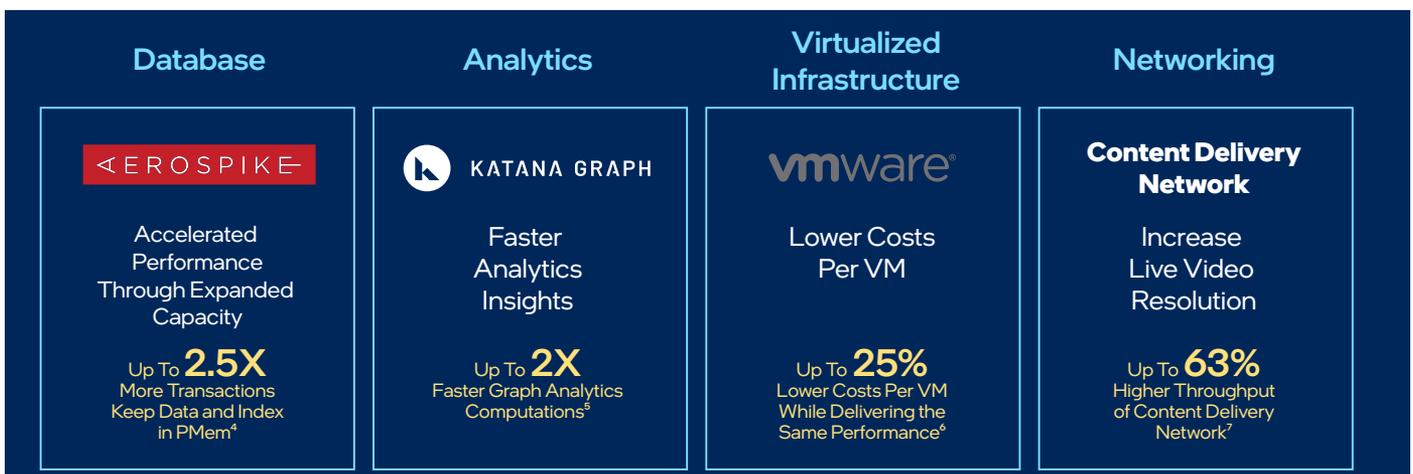


Figure 5. Intel® Optane™ PMem 200 series boosts performance across a wide range of enterprise applications.

## Intel Optane PMem 200 Series Addresses Key Data Center Challenges

Intel Optane PMem 200 series addresses many of the computing challenges that data centers face today. These challenges include high DRAM costs for large-memory nodes, data protection during outages and maintenance, emerging workloads that take advantage of tiered memory and storage architectures, and more.

## Active Workloads with PMem Deliver Real Value Today<sup>1</sup>

Many system solutions already take advantage of PMem, leveraging capacity and persistence to enhance workload performance and reliability across multiple use cases and domains.

## Programming Model

The software interface for Intel Optane persistent memory was designed in collaboration with dozens of companies to create a unified programming model for the technology. This software interface is independent of any specific persistent memory technology and can be used with Intel Optane PMem 200 series or any other persistent memory technology. The Storage Network Industry Association ([SNIA, https://www.snia.org](https://www.snia.org)) published a specification of the model.

The model exposes three main capabilities: management, storage, and memory map paths.

- **Management path** allows system administrators to configure persistent memory products and check their health.
- **Storage path** supports the traditional storage APIs, where existing applications and file systems need no change. They simply see the persistent memory as very fast storage.
- **Memory-mapped path** exposes persistent memory through a persistent memory-aware file system. Applications have direct load/store access to the persistent memory. This direct access does not use the page cache like traditional file systems. It has been named DAX by the operating system vendors.

The Persistent Memory Development Kit (PMDK), available at [pmem.io](https://pmem.io), provides libraries meant to make PMem programming easier. Software developers only pull in the features they need, keeping their programs lean and fast on PMem. These libraries are fully validated and performance-tuned by Intel. They are open source and product-neutral, working well on a variety of PMem products. The PMDK contains a collection of open source libraries, which build on the SNIA programming model. The PMDK is fully documented and includes code samples, tutorials, and blogs. Language support for the libraries exists in C and C++, with support for Java, Python, and other languages in progress.

With the industry-standard persistent memory programming model and downloadable PMDK, developers can build simpler and more powerful applications today to benefit operations now and into the future.

## Operational Modes

Intel Optane PMem 200 series has multiple operating modes:<sup>8</sup>

**Memory Mode** delivers large memory capacity without application changes or system modification and with performance close to that of DRAM, depending on the workload. In Memory Mode, the CPU memory controller sees all of the Intel Optane PMem 200 series as volatile system memory (without persistence). The CPU uses DRAM as a fast cache to Intel Optane PMem. In Memory Mode, data in the modules is protected with a single encryption key that is discarded upon power down, making the data inaccessible.

Memory Mode's large capacity allows workloads that are constrained by memory capacity to operate more efficiently. Virtual solutions can also scale easier on a platform, with more utilization from the same server, helping to reduce TCO.

**App Direct Mode** enables large memory capacity and data persistence. Software can access DRAM and persistent memory as two separate pools of memory. In App Direct Mode, software and applications that are enabled for the industry-standard persistent memory programming model can talk directly to PMem. Direct access reduces the complexity in the stack and takes full advantage of byte-addressable persistent memory with cache coherence, ensuring multiple copies of the same data across the node are in sync. (See <https://www.minitool.com/lib/cache-coherence.html>.) This capability extends the usage of persistent memory outside the local node and provides consistent low latency, supporting larger datasets.

App Direct Mode can also be used with standard file APIs to access persistent memory address space (called Storage over App Direct). Access is done without modifications to the existing applications or the file systems that expect block storage devices. Storage over App Direct presents Intel Optane PMem as high-performance block storage, without the latency of moving data to and from the I/O bus.

In App Direct Mode, data is encrypted using a key stored on the module in a security metadata region, which can only be accessed by the Intel Optane PMem 200 series controller.

The modules are locked at power loss. A passphrase is needed to unlock and access the data. If a module is repurposed or discarded, a secure cryptographic erase and DIMM over-write operation keeps data from being accessed.

## Data at Rest Secured for Enhanced Protection

With Intel Optane PMem 200 series, data at rest is better protected by strong, industry-standard security measures. All data is encrypted by 256-bit Advanced Encryption Standard (AES-256). With hardware-based cipher key processing, encryption is transparent to application software. Thus, no software code changes are needed when adding PMem. Hardware-enabled encryption on 3rd Gen Intel Xeon Scalable processors results in strong, industry-standard data security with low impact on performance.

## Data Persistence and Automatic Cache Flushes

Unlike DRAM, Intel Optane PMem 200 series, when operating in App Direct mode, retains data during a planned or unplanned restart, avoiding time-consuming data reloads. Maintaining data in PMem, means less down time, fewer losses from system outages, and increased operational efficiency.

Applications that manage data structures in persistent memory routinely call cache-flushing commands to move data stored in the CPU cache to persistent memory. Waiting for these cache flushes to complete can reduce performance. Applications can now avoid those waits altogether with extended asynchronous DRAM refresh (eADR). eADR is a new platform feature available with Intel Optane PMem 200 series with 3rd Gen Intel Xeon Scalable processors. The application automatically detects eADR and skips cache flushing, even if it experiences a system crash or power failure. This means application “Stores” are considered persistent as soon as they are visible to the application. This ability brings “lock-free” programming to persistent memory. eADR helps free up resources and improves performance.

Using eADR with Intel Optane PMem requires platform hardware support. The system requires additional stored energy available to allow cache flushes to occur on power failure. Coupled with application checking for the enablement of eADR and following the persistent memory programming model will ensure a solution is enabled fully for eADR support.



**Create a high-performance, large-capacity persistent memory tier with Intel® Optane™ PMem 200 series to improve key data center metrics such as:**

- Increased CPU utilization and utility of each server
- Increased application throughput with more memory capacity
- Increased VM density or support for more services and users
- Support for larger datasets without prohibitive price tags
- Improve business resilience for systems with critical data
- Consolidate server footprint, reduce software licensing costs, and optimize return on your enterprise investment

## Data Utilization

Intel Optane PMem 200 series is the next-generation of a groundbreaking memory technology innovation. Deployed with 3rd Gen Intel Xeon Scalable processors, this technology can transform critical data workloads—from cloud to databases, in-memory analytics, virtualized infrastructure, content delivery networks, and more.

### Intel® Optane™ Persistent Memory 200 Series Data Sheet

Product Family	Intel® Optane™ Persistent Memory 200 Series					
Compatible Processor	3rd Gen Intel® Xeon® Scalable processors on 2-socket and 4-socket platforms					
Form Factor	Persistent Memory Module					
SKU+	128 GB		256 GB		512 GB	
User Capacity+	126.7 GB		253.7 GB		507.7 GB	
Platform Capacities	4S systems: 3 TB PMem + 1.5 TB DRAM per socket (4.5 TB total) per socket 2S systems: 4 TB PMem + 2 TB DRAM per socket (6 TB total) per socket					
MOQ	4	50	4	50	4	50
MM#	999HGR	999HGZ	999HH0	999HH1	999HH2	999HH3
Product Code	NMB1XXD128GPSU4	NMB1XXD128GPSUF	NMB1XXD256GPSU4	NMB1XXD256GPSUF	NMB1XXD512GPSU4	NMB1XXD512GPSUF
Model String	NMB1XXD128GPS		NMB1XXD256GPS		NMB1XXD512GPS	
Technology	Intel® Optane™ Technology					
Limited Warranty	5 years					
AFR	≤ 0.44					
Endurance 100% Writes 15W 256B	292 PBW		497 PBW		410 PBW	
Endurance 67% READ; 33% Write 15W 256B	224 PBW		297 PBW		242 PBW	
Endurance 100% Write 15W 64B	73 PBW		125 PBW		103 PBW	
Endurance 67% READ; 33% Write 15W 64B	56 PBW		74 PBW		60 PBW	
Bandwidth 100% Read 15W 256B	7.45 GB/s		8.10 GB/s		7.45 GB/s	
Bandwidth 67% Read; 33% Write 15W 256B	4.25 GB/s		5.65 GB/s		4.60 GB/s	
Bandwidth 100% Write 15W 256B	2.25 GB/s		3.15 GB/s		2.60 GB/s	
Bandwidth 100% Read 15W 64B	1.86 GB/s		2.03 GB/s		1.86 GB/s	
Bandwidth 67% Read; 33% Write 15W 64B	1.06 GB/s		1.41 GB/s		1.15 GB/s	
Bandwidth 100% Write 15W 64B	0.56 GB/s		0.79 GB/s		0.65 GB/s	
DDR Frequency	Up to 2666 MT/s (4-socket systems); Up to 3,200 MT/s (2-socket systems)					
Operating Modes	App Direct only (4-socket systems); App Direct & Memory Mode (2-socket systems)					
Max TDP	15W					
Temperature (Tjmax)	≤ 83°C (85°C shutdown, 83°C default) media temperature					
Temperature (Tambient)	48°C @ 2.4m/s for 12W					
Temperature (Tambient)	43°C @ 2.7m/s for 15W					

+GB = 10<sup>9</sup>

Bandwidths are +/- 3%

Learn more at  
[intel.com/optanepersistentmemory](https://intel.com/optanepersistentmemory)



#### Notices & Disclaimers

<sup>1</sup> [www.intel.com/PerformanceIndex](https://www.intel.com/PerformanceIndex)

<sup>2</sup> Up to 6 TB on 2-socket systems. Up to 4.5 TB on 4-socket systems.

<sup>3</sup> Source: IDC 2020 MaturityScape Digital Transformation.

<sup>4</sup> Source: Aerospike 2.5X: Citation #5

<https://edc.intel.com/content/www/us/en/products/performance/benchmarks/intel-optane-persistent-memory-200-series/>

<sup>5</sup> Source: Katana Graph 2X: Citation #4

<https://edc.intel.com/content/www/us/en/products/performance/benchmarks/intel-optane-persistent-memory-200-series/>

<sup>6</sup> Source: VMware 25%: Citation #3

<https://edc.intel.com/content/www/us/en/products/performance/benchmarks/intel-optane-persistent-memory-200-series/>

<sup>7</sup> Source: CDN 63%: Citation #2

<https://edc.intel.com/content/www/us/en/products/performance/benchmarks/intel-optane-persistent-memory-200-series/>

<sup>8</sup> Requires 3rd Gen Intel Xeon Scalable processors to deliver Intel Optane persistent memory 200 series capabilities.

<sup>9</sup> Lower overall total cost of ownership (TCO).\*

\*Supports for more than 1.2x more VMs than the previous generation. Based on testing by Intel as of April 27, 2020 (Baseline) and March 31, 2020 (New).  
Baseline: 1-node, 2x Intel® Xeon® Platinum 8270 processor on 2S platform with 768GB (DDR 24 slots / 32GB / 2666) total memory, ucode 0x4000014 running Windows Server 2019 RS5-17763, and OLTP Cloud benchmark, storage is 7x Samsung PM963 M.2 960GB, 4x Intel SSDs S4600 (1.92TB), 1x Intel X520 SR2 (10Gb). Tested by Intel, on 31 Jan 2019.

New: 1-node, 2x Intel® Xeon® Platinum 8270 processor on 2S platform with 192, 1024GB (DDR 12 slots / 16 GB / 2666 + PMem 8 slots / 128GB / 2666) total memory, ucode 0x4000014 running Windows Server 2019 RS5-17763, and OLTP Cloud benchmark, storage is 7x Samsung PM963 M.2 960GB, 4x Intel SSDs S4600 (1.92TB), 1x Intel X520 SR2 (10Gb). Tested by Intel, on 31 Jan 2019.

Intel® Optane™ persistent memory pricing & DRAM pricing as of March 2021. Pricing referenced in TCO calculations is provided for guidance and planning purposes only and does not constitute a final offer. Pricing guidance is subject to change and may revise up or down based on market dynamics. Contact your OEM/distributor for actual pricing.

Performance varies by use, configuration and other factors. Learn more at [www.intel.com/PerformanceIndex](https://www.intel.com/PerformanceIndex).

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.

Your costs and results may vary.

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

© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries.

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