

TECHNICAL ARTICLE

WHAT IS INTEL® OPTANE™ PERSISTENT MEMORY (PMem) 200 SERIES?

| CONTENTS

| | |
|---|-----------|
| WHAT IS INTEL OPTANE™ PERSISTENT MEMORY? | 03 |
| WHERE DID IT COME FROM? | 03 |
| WHY DO I NEED IT? | 04 |
| HOW DOES IT WORK? | 05 |
| WHAT IS THE 2-TIER MEMORY/STORAGE CONFIGURATION? | 06 |
| HARDWARE CONSIDERATIONS | 08 |
| SUPPORTED OPERATING SYSTEMS | 10 |
| CAPACITY/MODELS AVAILABLE | 11 |



| WHAT IS INTEL® OPTANE™ PERSISTENT MEMORY (PMem)?

Intel® Optane™ Persistent Memory (PMem) 200 is the second generation of Intel's DC Persistent Memory (DCPMM) which is an emerging technology where non-volatile media is placed onto a Dual In-Line Memory Module (DIMM) and installed on the memory bus, traditionally used only for volatile memory.

This persistent memory is designed to exist on the same bus alongside volatile memory such as DRAM and often works in conjunction with it to achieve higher overall system memory capacity or to allow it to achieve better performance through DRAM caching.

The key thing that differentiates persistent memory DIMMS from DRAM DIMMs is that the data stored on them can be retained when the system is shut down or loses power. This allows the technology to be used as a form of permanent storage like Hard Disk Drives (HDDs) or Solid-State Drives (SSDs), but with memory-like latencies. Intel® Optane™ PMem is based on the Intel® Optane™ Memory technology and provides the ability to keep more data closer to the CPU for faster processing making the setup “warmer”.

DCPMM which is now known as Intel® Optane™ PMem 100 Series is designed for use with the Second-Generation Intel® Xeon® Scalable Platform processors, whereas Intel® Optane™ PMem 200 is designed for the Third-Generation Intel® Xeon® Scalable Platform processors code-named Ice Lake.

| WHERE DID IT COME FROM?

Intel® Optane™ PMem has had a long route to market. For those who have been following this may not be new but equally, it may tie up some previously heard of technologies or terms that have been in the press over the last 4-5 years. DRAM has been on the market in various generational changes since the mid to late 1940s and while progress on speed/capacity marches on, the underlying technology was largely the same to this day. NAND technology however, has had a shorter history but still has been around since the mid-1980s.

Intel® and Micron® has been working closely together to introduce a new type of memory technology, coined as 3D Xpoint in 2015. This is a new technology that can function in a similar way to DRAM and NAND depending on the application but has been stated to be, in theory, 1000x faster and 1000x more enduring than NAND flash memory.

3D Xpoint was then marketed and packaged into an Intel® storage/NAND replacement product for desktop/enthusiast use called Optane™ memory - which should not be confused with persistent memory. Optane™ memory originally got released in the form of two M.2 modules in 16GB and 32GB capacities, and these were placed in the market as a HDD performance accelerator, not as a direct memory extension or replacement.

Additionally, a data centre/enterprise, ultra-high performance, and endurance NVME based SSD was brought to market using Optane™ memory under the guise of the Intel® Optane™ SSD DC P4800X series.

Initially available in capacities up to 1.5TB with performance of more than half a million IOPS, and 2GB/s throughput the Optane™ series was the absolute definition of a flagship storage product.

| WHY DO I NEED IT?

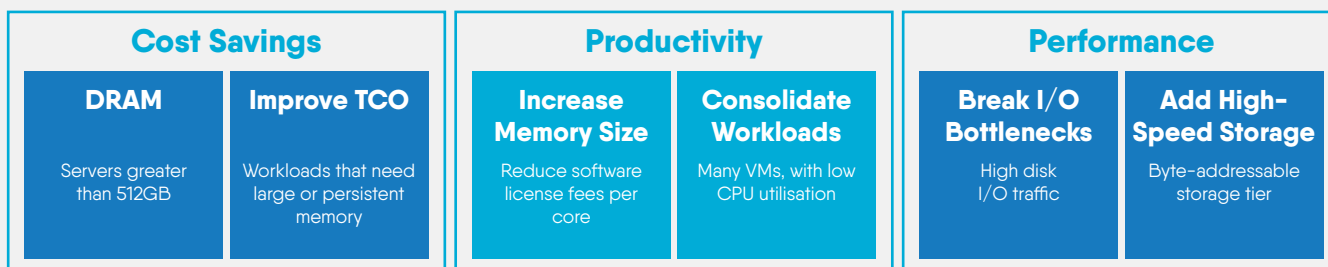
Over 90 percent of enterprises are in the midst of their digital transformation journey to become data-centric businesses. Thus, they will need to capture, analyse, and secure increasing amounts of data. Data growth is accelerating the demand for high performance, data-intensive computing. As the demand for compute grows, memory capacity in the system typically need to scale along with it. Large pools of DRAM help accelerate computing with low latency, but DRAM is limited in capacity and is becoming one of the most expensive components in modern server 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. This new memory 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 workloads can optimise performance and cost by creating a 2-tier hierarchy in memory and storage. Supported for these processors on 1, 2, and 4-socket platforms. If you have 2nd Gen Intel® Xeon® Scalable processors, then Intel® Optane™ PMem 100 series is for you.

CURRENT PROBLEMS



USE INTEL® OPTANE™ PERSISTENT MEMORY FOR...



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

Native persistence – the ability to store data even when powered off.

Large capacities, at lower prices than DRAM, mean larger datasets can exist closer to the CPU for faster processing. Higher capacities of Intel® Optane™ persistent memory is accelerating the industry-wide trend

towards real-time data processing. Delivered on the 2nd and 3rd Generation Intel® Xeon® scalable processors, large memory-bound workloads will have significant performance increase for rapid data processing.

Regarding hardware encryption, Intel® Optane™ persistent memory has industry standard AES-256 hardware encryption so you can rest easy knowing your data is secure. The encryption key is stored in a security metadata region on the module and is only accessible by the Intel® Optane™ persistent memory controller. When the module is repurposed or discarded, a secure cryptographic erase and DIMM over-write is utilised to keep data from being accessed.

The Intel® Optane™ persistent memory has two operating modes: Memory Mode – great for large memory capacity and does not require application changes as the memory is treated as volatile. The second mode is App Direct Mode which provides large memory capacity and allows enabled applications to talk directly to PMem as a 2nd tier memory. With distinct operating modes, you have the flexibility to take advantage of Intel® Optane™ persistent memory benefits across multiple workloads.

| HOW DOES IT WORK?

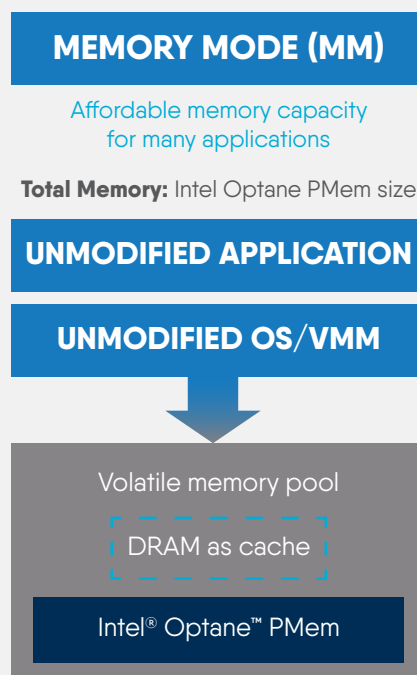
OPERATIONAL MODES

Intel® Optane™ PMem 200 series has multiple operating modes

Memory Mode is ideal for large memory capacity and doesn't need application changes, making Intel Optane™ persistent memory simple to adopt. When in Memory Mode, Optane™ PMem extends the amount of available volatile memory visible to the OS. The CPU uses DRAM as a fast cache to the Intel® Optane™ PMem. The CPU memory controller uses the DRAM as cache while the Intel® Optane™ persistent memory is used as addressable main memory.

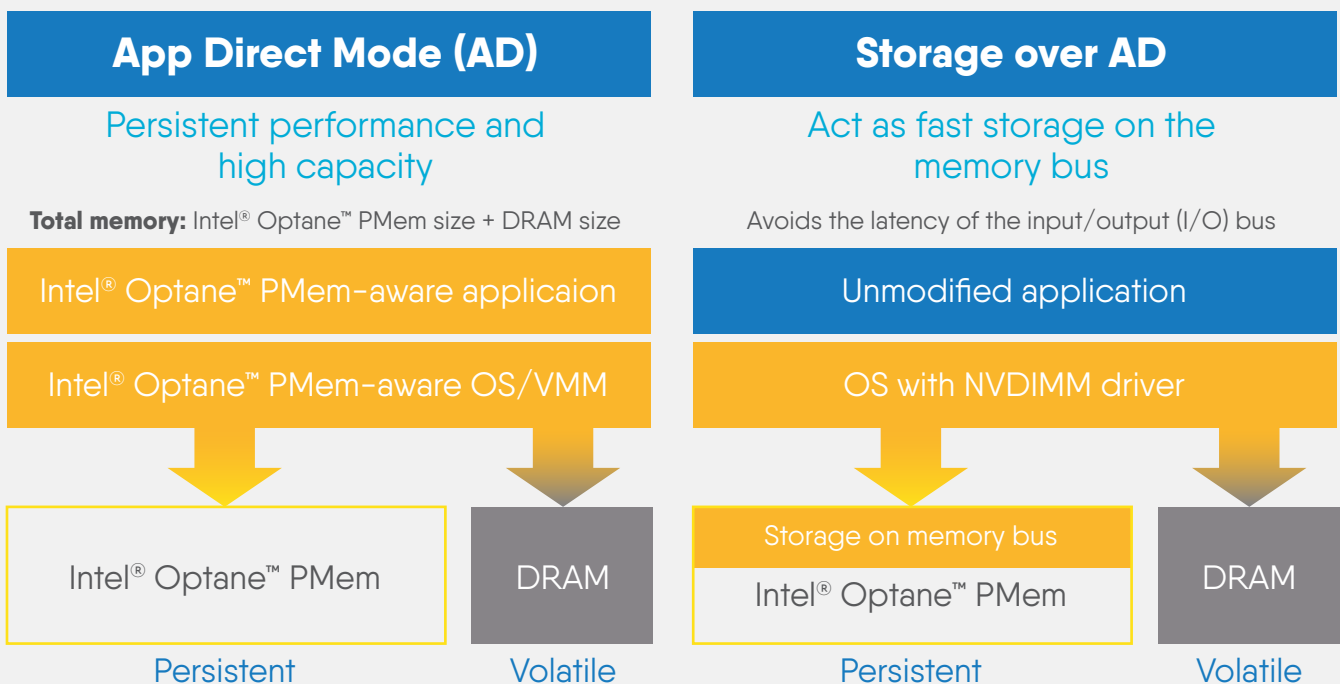
Intel® Optane™ persistent memory in Memory Mode can be an advantage for virtualisation as there is a bigger memory capacity which delivers increased VMs and higher memory per VM at a reduced cost in comparison to DRAM. I/O bound workloads can also benefit from utilising Memory Mode as Intel® Optane™ persistent memory because it delivers larger memory capacity which supports larger databases and is more cost-efficient than DRAM. With increased capacity, there is greater VM, and application density, which increases the utilisation of 2nd and 3rd Gen Intel® Xeon® Scalable processors.

When used in Memory Mode, the data in the Intel® Optane™ persistent memory is volatile. It is protected with a single encryption key that is discarded upon power down, making the data inaccessible.



App Direct Mode enables large memory capacity and data persistence for software to 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 have the ability to talk directly to PMem. Direct access reduces the complexity in the stack and takes full advantage of byte-addressable persistent memory with cache coherence, which 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 the same persistent memory address space (called Storage over App Direct) without any 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.



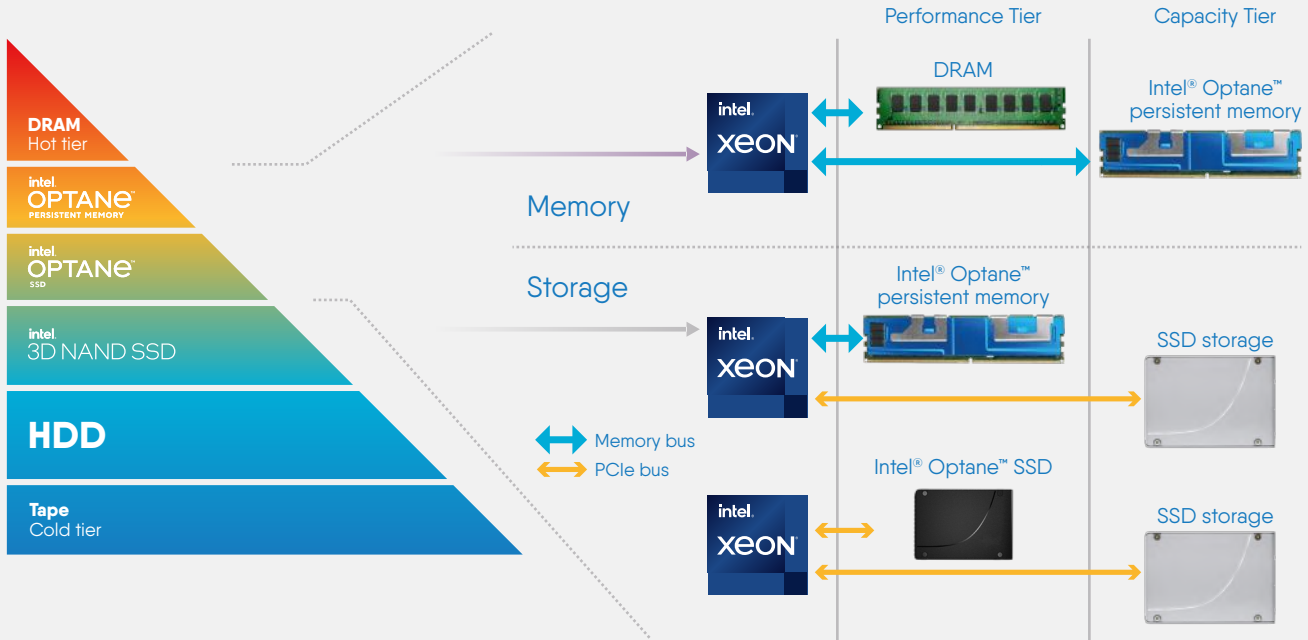
In App Direct Mode, data is encrypted using a key stored on the module in a security metadata region, which can only be accessed by the Intel® Optane™ PMem 200 series controller. The modules are locked at power loss, and a passphrase is needed to unlock and access the data. If a module is repurposed or discarded, a secure cryptographic erase and DIMM over-write operation keeps data from being accessed.

| WHAT IS THE 2-TIER MEMORY/STORAGE CONFIGURATION?

Intel® Optane™ PMem offers system architects and application developers options to create hierarchical memory and storage tiers to address data performance and capacity challenges. A hierarchical approach allows application developers to optimise 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 capacity available in the system.

For a two-tier memory model, low-latency DRAM offers exceptionally fast performance, while Intel® Optane™ PMem 200 series creates large memory capacity to store and protect data with DRAM-like speed. Depending on the application, PMem used as the large memory capacity can be persistent or volatile.

VISION OF 2-TIER MEMORY/STORAGE FOR OPTANE



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.

With a two-tier storage model, Intel® Optane™ PMem 200 series, utilised as a performance tier, delivers fast, byte-addressable access to most frequently accessed data. Other technologies, such as SSDs, offer slower access to warm data storage in a capacity tier.

These hierarchical memory and storage architectures made possible by Intel® Optane™ technology help optimise speed, latency, capacity, and cost. Choosing the right combinations for each application can help optimise systems and their workloads, from cloud to databases, in-memory analytics, virtualised infrastructure, content delivery networks, and more.

HARDWARE CONSIDERATIONS

While 2nd Generation of Intel® Xeon® Scalable Processors is a minimum requirement to run Intel® Optane™ Persistent Memory, not all the Xeon® SKUs are made equal. Platinum and Gold SKUs, all support Intel® Optane™ PMem 100 modules, but Silver and Bronze do not. With one exception being the Intel® Xeon® Silver 4215, which can support the modules making it the lowest spec SKU that is PMem 100 capable. X11 Generation Supermicro systems are the only generation of platform that can support Intel® Optane™ PMem 100 Series assuming that the latest BIOS/IPMI is used.

PROCESSOR PLATFORM SUPPORT

Intel Optane™ PMem 200 Series is supported by the following processors:

- 3rd Gen Intel® Xeon® Scalable (83xx/63xx/53xx/4314 series) processors

Supermicro Motherboard Platform Support for Intel® Optane™ PMem 200 Series:

- Intel® PMem 200 Series supports Supermicro X12SPx/X12Dxx/X12Qxx motherboards based on the Xeon® 3rd Gen processors.

Support for Multi Socket systems is per below:

| | | |
|--|---|--|
| AES 256-BIT encryption Secure Erase Up to 512 GB modules | UP TO 32% higher average memory bandwidth over the previous generation* AES 256-BIT encryption Secure Erase Up to 512 GB modules | |
| Intel® Optane™ PMem 100 series 2nd Generation Intel Xeon Scalable processors on 2S/4S/8S platform | Intel® Optane™ PMem 200 series 3rd Generation Intel Xeon Scalable processors on 4S platform | Intel® Optane™ PMem 200 series 3rd Generation Intel Xeon Scalable processors on 4S platform |
| 8-28 cores 6 channels memory | 18-28 cores 6 channels memory | 18-28 cores 6 channels memory |
| 3TB Intel® Optane™ PMem per socket* 4.5TB Total system memory per socket* 2,666 MT/s DDR 4+ Intel® Optane™ PMem | 3TB Intel® Optane™ PMem per socket* 4.5TB Total system memory per socket* 2,666 MT/s DDR 4+ Intel® Optane™ PMem | 4TB Intel® Optane™ PMem per socket* 6TB Total system memory per socket* 3,200 MT/s DDR 4+ Intel® Optane™ PMem |
| 18W Max thermal design power | eADR 15W Max thermal design power | eADR 15W Max thermal design power |

* Based on testing by Intel as of April 27 2020 (Baseline) and March 31 2020 (New).
 * 3TB Intel Optane PMem = 6 x 512GB Intel Optane PMem per socket, 4.5TB System Memory = 6 x 512GB Intel Optane PMem per socket + 6 x 256GB
 ** 4TB Intel Optane PMem = 8 x 512GB Intel Optane PMem per socket, 6TB System Memory = 8 x 512GB Intel Optane PMem per socket + 8 x 256GB

The following Supermicro Systems have been tested and certified to use Intel Optane™ PMem 200 Series:

| SKU | FORM FACTOR | # OF CPU | CPU TYPE | SYSTEM MEMORY |
|----------------------|-------------|----------|---|---|
| SYS-110P -WTR | 1U | 1 | <ul style="list-style-type: none"> Support Intel® 3rd Gen Xeon® Scalable Processors Single Socket LGA 4189 (Socket P+) supported TDP up to 270W | <ul style="list-style-type: none"> 8 DIMM slots Intel® DCPMM, DDR4-3200MHz ECC LRDIMM, DDR4-3200MHz ECC RDIMM, DDR4-3200MHz |
| SYS-120C -TN10R | 1U | 2 | <ul style="list-style-type: none"> 3rd Gen Intel® Xeon® Scalable processors Dual Socket LGA-4189 (Socket P+) supported TDP up to 270W; 3 UPI | <ul style="list-style-type: none"> 16 DIMM slots Up to 4 TB Intel® DCPMM, DDR4-3200MHz Up to 4 TB ECC LRDIMM, DDR4-3200MHz, Up to 4 TB ECC RDIMM, DDR4-3200MHz |
| SYS-120H -TNR | 1U | 2 | <ul style="list-style-type: none"> 3rd Gen Intel® Xeon® Scalable processors Dual Socket LGA-4189 (Socket P+) supported TDP up to 270W; 3 UPI up to 11.2GT/s | <ul style="list-style-type: none"> 32 DIMM slots Up to 8TB Intel® DCPMM, DDR4-3200MHz Up to 8TB ECC LRDIMM, DDR4-3200MHz, Up to 8TB ECC RDIMM, DDR4-3200MHz |
| SYS-220H -TN24R | 2U | 2 | <ul style="list-style-type: none"> 3rd Gen Intel® Xeon® Scalable processors Dual Socket LGA-4189 (Socket P+) supported TDP up to 270W; 3 UPI up to 11.2GT/s | <ul style="list-style-type: none"> 32 DIMM slots Up to 8TB Intel® DCPMM, DDR4-3200MHz Up to 8TB ECC LRDIMM, DDR4-3200MHz, Up to 8TB ECC RDIMM, DDR4-3200MHz |
| SYS-220H -TNR | 2U | 2 | <ul style="list-style-type: none"> 3rd Gen Intel® Xeon® Scalable processors Dual Socket LGA-4189 (Socket P+) supported TDP up to 270W; 3 UPI up to 11.2GT/s | <ul style="list-style-type: none"> 32 DIMM slots Up to 8TB Intel® DCPMM, DDR4-3200MHz Up to 8TB ECC LRDIMM, DDR4-3200MHz, Up to 8TB ECC RDIMM, DDR4-3200MHz |
| SYS-220HE -FTNRD | 2U | 2 | <ul style="list-style-type: none"> 3rd Gen Intel® Xeon® Scalable processors Dual Socket LGA-4189 (Socket P+) supported TDP up to 270W; 3 UPI up to 11.2GT/s | <ul style="list-style-type: none"> 32 DIMM slots Up to 8TB Intel® DCPMM, DDR4-3200MHz Up to 8TB ECC LRDIMM, DDR4-3200MHz, Up to 8TB ECC RDIMM, DDR4-3200MHz |
| SYS-510P -MR | 1U | 1 | <ul style="list-style-type: none"> 3rd Gen Intel® Xeon® Scalable processors Single Socket LGA-4189 (Socket P+) supported TDP up to 220W; | <ul style="list-style-type: none"> 8 DIMM slots Intel® DCPMM, DDR4-3200MHz ECC LRDIMM, DDR4-3200MHz, ECC RDIMM, DDR4-3200MHz |
| SYS-510P -WTR | 1U | 1 | <ul style="list-style-type: none"> 3rd Gen Intel® Xeon® Scalable processors Single Socket LGA-4189 (Socket P+) supported TDP up to 270W; | <ul style="list-style-type: none"> 8 DIMM slots Intel® DCPMM, DDR4-3200MHz ECC LRDIMM, DDR4-3200MHz, ECC RDIMM, DDR4-3200MHz |
| SYS-620C -TN12R | 2U | 2 | <ul style="list-style-type: none"> 3rd Gen Intel® Xeon® Scalable processors Dual Socket LGA-4189 (Socket P+) supported TDP up to 270W; 3 UPI | <ul style="list-style-type: none"> 16 DIMM slots Up to 4 TB Intel® DCPMM, DDR4-3200MHz Up to 4 TB ECC LRDIMM, DDR4-3200MHz, Up to 4 TB ECC RDIMM, DDR4-3200MHz |
| SSG-620P -ACR12H | 2U | 2 | <ul style="list-style-type: none"> 3rd Gen Intel® Xeon® Scalable processors Dual Socket LG-4189 (Socket P+) supported TDP up to 270W; up to 11.2GT/s | <ul style="list-style-type: none"> 16 DIMM slots Up to 18TB Intel® DCPMM, DDR4-2666MHz Up to 4TB ECC RDIMM/LRDIMM, DDR4-3200MHz |
| SSG-640P -E1CR36H | 4U | 2 | <ul style="list-style-type: none"> 3rd Gen Intel® Xeon® Scalable processors Dual Socket LG-4189 (Socket P+) supported TDP up to 270W; up to 11.2GT/s | <ul style="list-style-type: none"> 16 DIMM slots Up to 18TB Intel® DCPMM, DDR4-2666MHz Up to 4TB ECC RDIMM/LRDIMM, DDR4-3200MHz |
| SYS-6029TP -HTR | 2U | 2 | <ul style="list-style-type: none"> Dual Socket P (LGA 3647) 2nd Gen Intel® Xeon® Scalable Processors (Cascade Lake/Skylake), Dual UPI up to 10.4GT/s | <ul style="list-style-type: none"> 16 DIMM slots Up to 4TB 3DS ECC DDR4-2933MHz RDIMM/LRDIMM Supports Intel® Optane™ DCPMM |

With new systems continually being added to the supported and certified list, it is advisable to contact us if you have specific requirements or questions about current supported hardware.

Each system has different configuration and PMem installation layout (different slots and capacities) based upon the operating mode being used and should refer to the server manual for configuration process.



| SUPPORTED OPERATING SYSTEMS

Intel® Optane™ PMem is supported by the following operating systems:

INTEL® OPTANE™ PERSISTENT MEMORY 100 SERIES


- Microsoft Windows Server 2019
- Red Hat Enterprise Linux 7.6
- Red Hat Enterprise Linux 8.0
- SUSE Linux Enterprise Server 12 SP4
- SUSE Linux Enterprise Server 15 SP1
- SUSE Linux Enterprise Server 15 (with SUSE-SU-2019:0224-1 kernel update)
- VMware 6.7 U1

INTEL® OPTANE™ PERSISTENT MEMORY 200 SERIES

- Microsoft Windows Server 2019
- Red Hat Enterprise Linux 7.9
- Red Hat Enterprise Linux 8.2
- Red Hat Enterprise Linux 8.3
- SUSE Linux Enterprise Server 12 SP5


- SUSE Linux Enterprise Server 15 SP2
- VMware 7.0 U 2

| CAPACITY/MODELS AVAILABLE



INTEL® OPTANE™ PERSISTENT MEMORY 200 SERIES (128GB PMem) MODULE

- 128GB Capacity
- DDR-T Interface
- Hardware Encryption: AES 256 bit
- Boston/Intel SKU: NMB1XXD128GPSU4




INTEL® OPTANE™ PERSISTENT MEMORY 200 SERIES (256GB PMem) MODULE

- 256GB Capacity
- DDR-T Interface
- Hardware Encryption: AES 256 bit
- Boston/Intel SKU: NMB1XXD256GPSU4




INTEL® OPTANE™ PERSISTENT MEMORY 200 SERIES (512GB PMem) MODULE

- 512GB Capacity
- DDR-T Interface
- Hardware Encryption: AES 256 bit
- Boston/Intel SKU: NMB1XXD512GPSU4




INTEL® OPTANE™ PERSISTENT MEMORY 100 SERIES (128GB PMem) MODULE

- 128GB Capacity
- DDR-T Interface
- Hardware Encryption: AES 256 bit
- Boston/Intel SKU: NMA1XXD128GPSU4



INTEL® OPTANE™ PERSISTENT MEMORY 100 SERIES (256GB PMem) MODULE

- 256GB Capacity
- DDR-T Interface
- Hardware Encryption: AES 256 bit
- Boston/Intel SKU: NMA1XXD256GPSU4



INTEL® OPTANE™ PERSISTENT MEMORY 100 SERIES (512GB PMem) MODULE

- 512GB Capacity
- DDR-T Interface
- Hardware Encryption: AES 256 bit
- Boston/Intel SKU: NMA1XXD512GPSU4

If you would like more information or design and architecture help around Intel's Optane™ DC Persistent Memory, then we'd be keen to hear from you. You can get in touch via phone **+44 (0) 1727 876 100** or email **sales@boston.co.uk**

Further information on Intel® DCPMM and compatible products can be found on our website: www.boston.co.uk.

Written By Waseem Pervez

BOSTON
Servers | Storage | Solutions

intel®

UK (Headquarters)

+44 (0) 1727 876 100
sales@boston.co.uk
www.boston.co.uk

FRANCE

+33 (0)1 72 81 06 76
sales@boston-it.fr
www.boston-it.fr

INDIA

+91 22 5002 3262 (Mumbai office)
+91 80 4308 4000 (Bangalore office)
sales@bostonindia.in
www.bostonindia.in

AFRICA

+27 826 486 485
sales@boston.co.uk
www.boston.co.uk

GERMANY

+49 (0) 89 9090 1993
sales@boston-it.de
www.boston-it.de

SWITZERLAND

+47 (0) 71 554 22 75
sales@boston-it.ch
www.boston-it.ch