

Intel® Optane™ Persistent Memory with VMware Configuration Guide

Evaluating performance requirements for Intel® Optane™ persistent memory configurations in Memory Mode on VMware vSphere platforms



Intel® Optane™ persistent memory (Intel® Optane™ PMem) is an innovative technology that delivers a unique combination of affordable, large-capacity memory with support for data persistence. In virtualized environments, Intel Optane PMem can provide more memory per VM and improve performance of memory-consuming workloads that cause swapping. More VMs per server can increase the total throughput and improve processor utilization, leading to decreased total costs through server consolidation or reduced software licensing costs.

VMware vSphere [versions 6.7 EP10 and later](#) support Intel Optane PMem 100 series with 2nd Generation Intel® Xeon® Scalable processors in App Direct and Memory Mode. This document provides guidelines assessing current deployments using metrics in VMware vCenter and determining the platform configuration to use with Intel Optane PMem in Memory Mode for VMware vSphere.

Determining Platform Configuration Needs

The goal of the following steps is to assess the workload profile of the current virtualized environment and determine the optimal platform configuration with Intel Optane PMem for host servers.

1. **Characterize the current deployment** (DRAM-only systems) using VMware vCenter metrics.
2. **Assess the workload fit** for Intel Optane PMem.
3. **Identify opportunities** to consolidate servers or replace DRAM with Intel Optane PMem.
4. **Determine the best Intel Optane PMem and DRAM configuration.**

Step 1. Characterize the Current Deployment

Identifying trends, peaks, and troughs on currently deployed workloads will help determine the opportunities to replace DRAM or consolidate servers with Intel Optane PMem. For best results, collect at least one week of data for representative workloads on the host systems for the following vCenter metrics (for instruction on collecting metrics refer to the [vSphere Performance and Monitoring](#) guide):

- **CPU Usage (%)**. Represents CPU usage for the host.
- **Memory Usage (%)**. Represents memory usage for the host.
- **Memory Active (MB)**. Represents the total active guest physical memory of all (powered on) VMs on the host.
- **Memory Swap Used (MB)**. Represents the total memory swapped by all (powered on) VMs on the host.
- **Memory Consumed (MB)**. Represents the amount of memory used on the host.

Step 2. Assess the Workload Fit for Intel® Optane™ PMem

Using the metrics from Step 1, analyze the data for opportunities where Intel Optane PMem can provide additional memory capacity, relieving constraints on applications and VMs while meeting performance needs. The following scenarios indicate that workloads and VMs on the host server are generally well-suited for Intel Optane PMem:

- The most important criteria is that hosts have a **memory active (MB)** that is less than 25 percent of **total physical memory of the host (MB)**. For simplicity throughout this document, this calculation of **memory active (MB)/total physical memory of the host (MB)** will be referred to as the **% of active memory**.
- Hosts that are running out of memory, which is indicated by high **memory usage (%)**. The threshold at which the memory usage is considered “high” for the host will be unique to your deployment environment.
- Hosts that experience memory swap, which is indicated by **memory swap used (MB)**. A value greater than 0 indicates that memory swap is used on the host.

Table 1 shows an assessment of six hosts and their suitability for Intel Optane PMem that looks at the three indicators listed above and SLA thresholds that will be unique to your deployment environment.

Table 1. Example Assessment of Workload Fit for Intel® Optane™ Persistent Memory Suitability

	Is Memory Swap Used (MB) > 0?	Is Memory Active ÷ Total Physical Memory ≤ 25%?	What is the Memory Usage %?	SLA	Is the Workload Suitable for Intel Optane PMem?
DRAM-Host1	Yes	Yes	High	Strict	Yes
DRAM-Host2	No	Yes	Medium	Relaxed	Yes
DRAM-Host3	No	No	Medium	Strict	No
DRAM-Host4	No	Yes	High	Strict	Yes
DRAM-Host5	Yes	Yes	High	Relaxed	Yes
DRAM-Host6	Yes	Yes	Medium	Relaxed	Yes

The following recommendations determine which hosts can benefit from Intel Optane PMem configuration and which should remain on DRAM-only:

- **Host1** has % of active memory less than 25%, a high memory usage, and memory swap is used, which indicates the host can benefit from more memory capacity, making it a good fit for Intel Optane PMem despite a strict SLA.
- **Host2** has % of active memory less than 25% and is good fit for Intel Optane PMem because it has medium amount of memory usage and a relaxed SLA.
- **Host3** has % of active memory greater than 25%, and the strict SLA makes it best suited to remain as a DRAM only host.
- **Host4** has % of active memory less than 25%, and high memory usage, indicating a good fit for Intel Optane PMem despite a strict SLA.
- **Host5** has % of active memory less than 25%, high memory usage, and a relaxed SLA, making it a good fit for Intel Optane PMem.
- **Host6** has % of active memory less than 25%, medium memory usage, and a relaxed SLA, making it a good fit for Intel Optane PMem.

Step 3. Identify Server Consolidation or Memory Replacement Opportunities

VMs and applications typically request significantly more memory allocation than what is actually used, leading to servers that are not fully utilized. Using the analysis from Step 2, now look at the host CPU utilization to identify opportunities to replace the DRAM with Intel Optane PMem to lower memory costs or add more memory capacity and consolidate servers to help reduce the cost of maintenance, hardware, and licensing (see Table 2).

Table 2. Example Analysis of Six Hosts for Consolidation or DRAM Replacement Opportunities

	CPU Usage %	Memory Usage %	Recommendation from Step 2	Consolidate or Replace
DRAM-Host1	20%	70%	PMem	Consolidate hosts 1 and 2
DRAM-Host2	40%	60%	PMem	Consolidate hosts 1 and 2
DRAM-Host3	40%	20%	DRAM	Remain DRAM only
DRAM-Host4	50%	80%	PMem	DRAM Replacement
DRAM-Host5	25%	75%	PMem	Consolidate hosts 5 and 6
DRAM-Host6	35%	55%	PMem	Consolidate hosts 5 and 6

The CPU and memory usage thresholds will be unique to your deployment requirements. In this example, the thresholds are 60 percent for CPU utilization and 75 percent for memory usage. Using the 6-host example in Table 2, consider the following options:

- **Host1** and **Host2** can be consolidated into one Intel Optane PMem-based server because the total combined CPU utilization and memory capacity requirements are within the defined thresholds.
- **Host3** is not evaluated, as the Step 2 analysis showed it was not a suitable fit for Intel Optane PMem.
- **Host4** is a candidate for Intel Optane PMem to replace a portion of the DRAM for lower memory costs.
- **Host5** and **Host6** can be consolidated into one Intel Optane PMem-based server.

Step 4. Determine the System Memory Configuration for Intel Optane PMem and DRAM

Servers deployed with Intel Optane PMem in Memory Mode utilize the DRAM installed in the system as a cache to deliver DRAM-like performance, which is used as high-capacity main memory. To determine the optimal system configuration, consider the following:

- **To determine the Intel Optane PMem configuration**, add the **memory consumed** and **memory swap used** of the hosts that will be consolidated or where DRAM will be replaced. Optionally, for capacity planning, add 10-25 percent of headroom for high-availability or other failover scenarios.

$$\text{Intel Optane PMem Needed For Host} = \text{Memory Consumed} + \text{Memory Swap Used} + \text{Optional Headroom Buffer}$$

- **To determine the DRAM configuration**, add **memory active** of the hosts that will be consolidated or where DRAM will be replaced and multiply by α , where α is a multiplier that can range from 1-2. VMware recommends that a host steady-state active memory should not exceed 50 percent of the available DRAM, which represents a multiplier of 2. Depending on the specific VM performance, SLA, or total cost constraints, a smaller DRAM configuration may meet deployment needs, which could be a multiplier of 1-2.

$$\text{DRAM Needed For Host} = \text{Active Memory} * \alpha$$

Table 3 summarizes the common Intel Optane PMem-to-DRAM configurations per 2-socket host (for 4-socket platforms, scale the configuration accordingly). The rows highlighted in green are memory configurations supported by VMware for production deployment when used along with the Balanced Profile PMem setting in the BIOS (see the knowledge base article [vSphere Support for Intel's Optane Persistent Memory \(PMEM\) \(67645\)](#)). The Balanced Profile setting in the BIOS (see [Appendix B](#) for additional information) is recommended for Intel Optane PMem in Memory Mode for all configurations and workloads.

Table 3. Common Intel® Optane™ Persistent Memory-to-DRAM Configurations per 2-Socket Host

PMem Needed for Host	DRAM Needed for Host	Intel Optane PMem per 2-Socket Host	DRAM per 2-Socket Host
384 GB or less	384 GB or less	N/A	384 GB (12x32 GB DRAM)
512 GB or less	192 GB or less	512 GB (4x128 GB PMem)	128 GB (8x16 GB DRAM)
		512 GB (4x128 GB PMem)	192 GB (12x16 GB DRAM)
Between 512 GB – 1,024 GB	384 GB or less	1,024 GB (8x128 GB PMem)	192 GB (12x16 GB DRAM)
		1,024 GB (8x128 GB PMem)	384 GB (12x32 GB DRAM)
Between 1,024 GB – 1,536 GB	384 GB or less	1,536 GB (12x128 GB PMem)	192 GB (12x16 GB DRAM)
		1,536 GB (12x128 GB PMem)	384 GB (12x32 GB DRAM)
1,536 GB or More	384 GB or more	3 TB (12x256 GB PMem)	768 GB (12x64 GB DRAM)
		6 TB (12x512 GB PMem)	1.5 TB (12x128 GB DRAM)

■ VMware-supported configurations when deployed with Balanced Profile BIOS setting

The non-highlighted configurations are still valid and can be deployed with production support with VMware approval through the Request for Product Qualification (RPQ) process. Contact your VMware representative for more information.

In the 6-host example, the following memory configurations are recommended based on the previous analysis and considering the SLA requirements:

Table 4. Optimal Intel® Optane™ Persistent Memory-to-DRAM Configurations

	PMem Needed	DRAM Needed	SLA	Recommendation from Step 3	Configuration
DRAM-Host1	768 GB	60 GB	Strict	Consolidate hosts 1 and 2	1.5 TB Intel Optane PMem + 384 GB DRAM
DRAM-Host2	768 GB	100 GB	Relaxed	Consolidate hosts 1 and 2	
DRAM-Host3	768 GB	150 GB	Strict	DRAM	768 GB DRAM
DRAM-Host4	1.5 TB	120 GB	Strict	DRAM Replacement	1.5 TB Intel Optane PMem + 384 GB DRAM
DRAM-Host5	512 GB	80 GB	Relaxed	Consolidate hosts 5 and 6	1 TB Intel Optane PMem + 192 GB DRAM
DRAM-Host6	512 GB	70 GB	Relaxed	Consolidate hosts 5 and 6	

Recommended configurations:

- **Host1** and **Host2** can be consolidated on a single server with Intel Optane PMem, using a supported configuration with 384 GB of DRAM to meet the strict SLA required for host 1.
- **Host3** should remain unchanged as a DRAM-only system.
- **Host4** memory capacity will not change once Intel Optane PMem is added.
- **Host5** and **Host6** can be consolidated on a single server with Intel Optane PMem. Based on the relaxed SLA requirements and optimal cost savings, a 1 TB Intel Optane PMem with 192 GB DRAM configuration is optimal. This configuration requires VMware approval for production support in deployment.

Conclusion

Intel Optane PMem is an innovative memory technology that delivers a unique combination of affordable, large-capacity memory with support for data persistence. For virtualization, large memory capacity significantly increases VM density, server utilization, and throughput by avoiding swapping. By providing more memory per VM, Intel Optane PMem can improve performance, and more VMs per server increases throughput and processor utilization, leading to decreased total costs.

Find the solution that is right for your organization. Contact your Intel representative or visit [intel.com/optanepersistentmemory](https://www.intel.com/optanepersistentmemory).

Learn More

- [VMware vSphere Performance with Intel Optane Persistent Memory in Memory Mode – Balanced Profile](#)
- [Intel Optane Persistent Memory Quick Start Guide](#)
- [VMware vSphere Central - Persistent Memory](#)

Appendix A: System Configuration Considerations

Find servers compatible with Intel Optane persistent memory by using the [VMware Compatibility Guide](#) and searching for the “Persistent Memory” feature. For additional information, refer to the VMware knowledge base article [vSphere Support for Intel's Optane Persistent Memory \(67645\)](#). If you do not see the server of your choice, contact your server vendor for the availability.

Return to [Determining Platform Configuration Needs](#).

Appendix B: Balanced Profile Persistent Memory BIOS Setting

Intel Optane PMem coexists with traditional DDR4 DIMMs on the same memory bus. The memory controller in the 2nd Generation Intel Xeon Scalable processor arbitrates between the memory transactions coming from DRAM and Intel Optane PMem. Different arbitration profiles have been defined to determine the algorithm for when the memory controller switches between DRAM and Intel Optane PMem transactions. These profiles are configurable through a persistent memory BIOS setting. Originally, there were two profiles: Bandwidth Optimized and Latency Optimized. A new profile, Balanced Profile, was added specifically to optimize Memory Mode performance.

- **Bandwidth Optimized** arbitrates between DRAM and Intel Optane PMem to maximize DRAM bandwidth.
- **Latency Optimized** arbitrates between DRAM and Intel Optane PMem to minimize DRAM latency.
- **Balanced Profile** is optimized for Memory Mode, allowing the controller to switch more often between DRAM and Intel Optane PMem. This allows eviction transactions in DRAM to execute faster.

Figure B1 illustrates the memory bus switching between DRAM and Intel Optane PMem traffic based on the BIOS profiles.

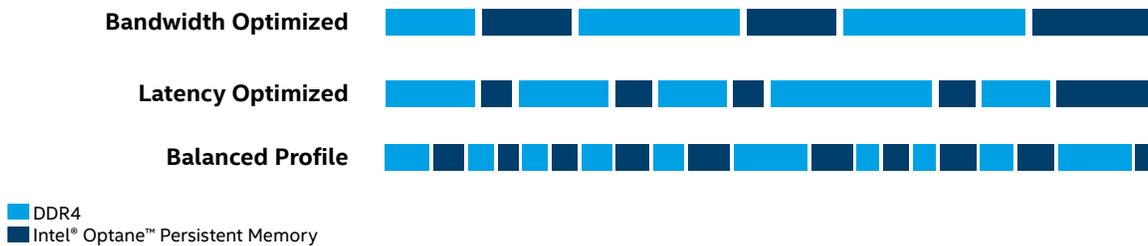


Figure B1. BIOS profiles manage DRAM and Intel® Optane™ Persistent Memory, switching for specific needs.

In Memory Mode, DRAM is used as a cache while the Intel Optane PMem is used as volatile main memory to deliver DRAM-like performance (depending on the workload). The CPU memory controller first attempts to retrieve data from the DRAM cache, and when the data is present it returns the request, similar to how DRAM access works today. When the data is not present, the request is sent to Intel Optane PMem. The request is returned to the CPU and in parallel is sent to the DRAM cache. This extra request, resulting from the cache miss, in addition to the marginally higher latency of Intel Optane PMem compared to DRAM, can negatively impact performance. Balanced Profile optimizes Memory Mode performance by balancing the thresholds between DRAM and Intel Optane PMem to improve the bandwidth for all DRAM cache misses and maintaining the performance of DRAM cache hits.

Intel **recommends the Balanced Profile BIOS setting** for workloads using Intel Optane PMem in Memory Mode. Contact your preferred OEM vendor for more information, as the Balanced Profile may not be the default persistent memory BIOS setting.



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](https://www.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.

© Intel Corporation. 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. 0820/CWAN/KC/PDF 342405-001US