


Implementing Intel® Optane™ Persistent Memory BIOS Optimizations for Memory Mode

For workloads using Intel Optane PMem in Memory Mode, Intel recommends using the new Balanced Profile BIOS setting to improve performance



Intel® Optane™ persistent memory 100 series (Intel Optane PMem) coexists with traditional DDR4 DIMMs on the same 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 memory transactions. These profiles are configurable options through a persistent memory BIOS setting. Previously, systems could choose profiles to either optimize DRAM bandwidth or latency. A new BIOS profile called Balanced Profile was developed to optimize Memory Mode performance. The three profiles now available include:

- **Bandwidth Optimized** arbitrates between DRAM and Intel Optane PMem to maximize DRAM bandwidth on the memory bus.
- **Latency Optimized** arbitrates between DRAM and Intel Optane PMem to minimize DRAM latency on the memory bus.
- **Balanced Profile** is optimized for Memory Mode by allowing the controller to switch more often between DRAM and Intel Optane PMem so that eviction transactions in DRAM can execute faster.

Figure 1 shows how the different BIOS profiles affect the memory controller policies for switching between DRAM and Intel Optane PMem requests.

Persistent Memory BIOS Profiles

Bandwidth Optimized Profile

Arbitrates between DRAM and Intel Optane PMem to maximize DRAM bandwidth on the memory bus.



Latency Optimized Profile

Arbitrates between DRAM and Intel Optane PMem to minimize DRAM latency on the memory bus.



Balanced Profile

Optimized for Memory Mode by allowing the controller to switch more often between DRAM and Intel Optane PMem so that eviction transactions in DRAM can execute faster.



■ DDR4
■ Intel® Optane™ Persistent Memory

Figure 1. Depiction of the memory controller alternating between DRAM and Intel Optane PMem transactions.

How Balanced Profile Optimizes for Memory Mode Performance

In Memory Mode, DRAM is used as cache and Intel Optane PMem is used as volatile main memory to deliver DRAM-like performance, depending on the workload. Figure 2 shows the Memory Mode transaction flow.

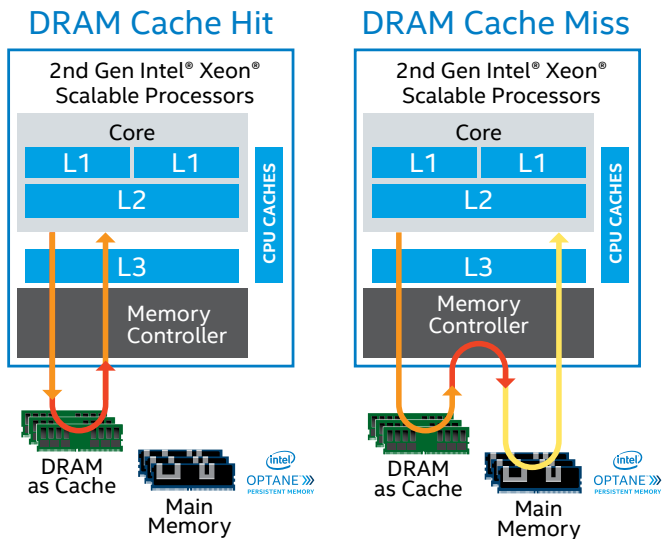


Figure 2. The Memory Mode transaction flow between DRAM and Intel Optane PMem.

The CPU memory controller will first attempt to retrieve data from the DRAM cache. When the data is present, it will return the request from the DRAM cache, similar to the way DRAM access works today. When the data is not present, the request will be 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 when 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 while maintaining the performance of DRAM cache hits.

Intel worked with server OEM partners to release this new persistent memory BIOS setting, as shown in Table 1. Contact your preferred OEM vendor for more information as the setting may not be the default.

Conclusion

Intel recommends the Balanced Profile BIOS setting for workloads using Intel Optane PMem in Memory Mode. We worked with server OEM partners to release this optimization as a new persistent memory BIOS setting. Please contact your preferred OEM vendor for more information.

Additional Resources

- intel.com/optanepersistentmemory

Table 1. OEM BIOS release versions with the new setting to optimize Memory Mode performance

OEM	BIOS Release Information
Cisco	Available under the “NVM Performance Setting” listed as Balanced Profile <ul style="list-style-type: none">• UCS C-Series Rack-Mount UCS-Managed Server Software Release 4.1(2a)• UCS C480 M5 Rack Server Software Release 4.1(2a)• UCS C220 M5 Rack Server Software Release 4.1(2a)• UCS C240 M5 Rack Server Software Release 4.1(2a)
Dell	Available within “System Profile Settings” under “Intel Persistent Memory Performance” and listed as Balanced Profile <ul style="list-style-type: none">• PowerEdge BIOS R740/R740xd/R640/R940/7920R Version 2.7.7
HPE	Available as a new BIOS/Platform Configuration (RBSU) option in the “Intel Persistent DC Performance Settings” listed as Balanced Performance Mode <ul style="list-style-type: none">• ProLiant DL360 Gen10 version 2.30_02-11-2020
Inspur	Available in Inspur NF5280M5 version 4.1.16
Intel	Available under the “NVM Performance Setting” listed as Balanced Profile <ul style="list-style-type: none">• Intel® Server Board S2600WF Release 02.01.0012
Lenovo	Available – Contact your Lenovo representative.
Supermicro	Available – Contact your Supermicro representative.

