

Introducing the Balanced Profile UEFI setting for Intel Optane Persistent Memory

Planning / Implementation

The Intel Optane Persistent Memory (PMem or DCPMM) provides higher memory capacity support than traditional DRAM which can benefit various applications with massive memory requirements such as virtualization, high-performance computing (HPC), big-data analytics, in-memory databases, Video on Demand (VoD) and SAP HANA.

PMem can be configured in Memory Mode, App Direct Mode, and Mixed Mode (combine Memory and App Direct Mode) to adapt to different user scenarios. Focusing on Memory Mode, this paper elaborates on the low-level access mechanism between DRAM and PMem to explain the difference between the Balanced, Bandwidth Optimized and Latency Optimized profiles for PMem.

Note: The Balanced Profile setting is currently only available on systems with Intel Optane Persistent Memory 100 Series and second-generation Intel Xeon Scalable processors.

Two-Level Memory Architecture

Managed by the memory controller inside Intel Xeon Scalable Processor, PMem adopts a Two-Level Memory (2LM) Architecture in Memory Mode as shown in the following figure. Different from a legacy memory hierarchy, the DRAM under the 2LM topology is treated as an L4 cache to store frequently used data, alongside the PMem working as system memory to provide a larger memory capacity compared to DRAM alone for memory hungry applications.

In 2LM mode, DRAM acts as the next level cache (L4), and PMem acts as system main memory. When data is found in DRAM, it is an L4 hit and access latency is the same as the latency to DRAM. However, when data is not found at DRAM, it is an L4 miss, and memory controller has to access data from PMem. This access from PMem, when it happens, incurs higher memory latency.

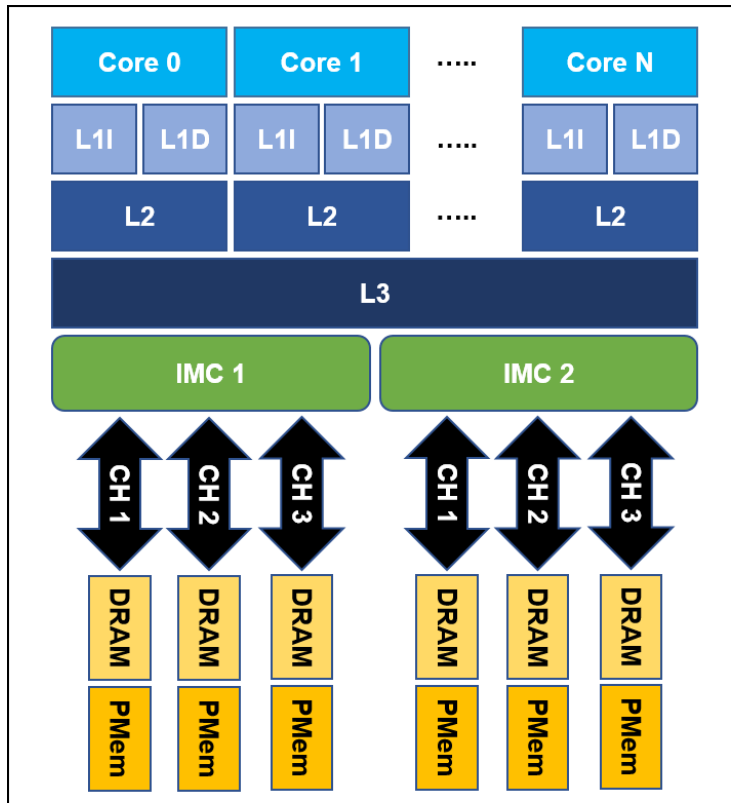


Figure 1. 2LM Architecture

UEFI Preset Modes for PMem

Based on the 2LM topology, the memory controller can only request data from either DRAM or the PMem at a time, such that the dispatch ratio between DRAM and PMem would be a crucial factor of the system performance. Fine tuning numerous thresholds along with dispatch ratio of memory controller and related parameters in the BIOS, Lenovo introduces a new Balanced Profile for PMem memory mode. Based on internal test results, the performance of both memory bandwidth and latency improved significantly when the Balanced Profile was applied.

With the introduction of Balanced Profile, Lenovo now offer three different PMem preset modes in the BIOS to adapt various kind of workloads, as listed in the following table.

Table 1. PMem preset modes

Profile	Purpose
Balanced (default)	Optimize PMem Profile for Memory Mode
BW Optimized	Optimize Bandwidth on both L4 cache and PMem
Latency	Optimize L4 cache Latency in the presence of PMem Bandwidth

While the Balanced Profile is recommended as the default preset mode for PMem memory mode, users can change preset mode as needed by clicking through the following menu selections in System Setup in UEFI: **System Settings > Socket Configuration > Memory Configuration > NGN Configuration > NVM Performance Setting**. Then select from these choices as shown in the following figure:

- BW Optimized
- Latency Optimized
- Balanced Profile

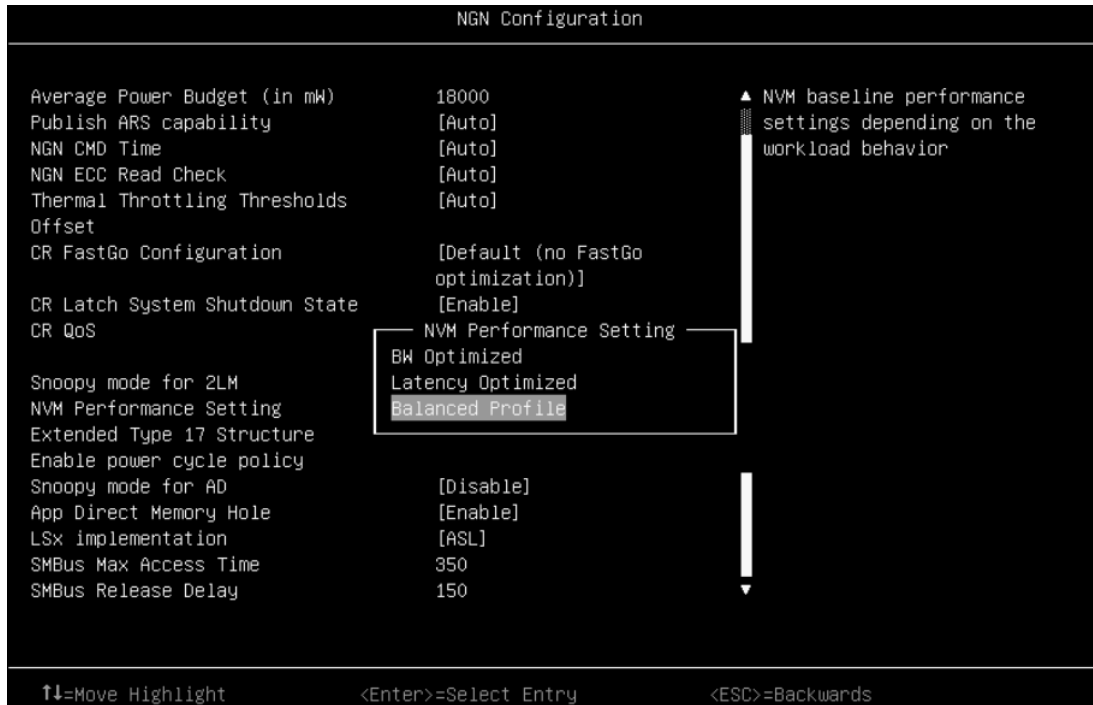


Figure 2. PMem performance mode settings (click the image for a larger version)

OneCLI V2.5.0 or later supports Intel PMem and can also be used to setup the PMem performance mode by using the OneCLI command as shown below:

```
./OneCli config set IntelOptaneDCPMM.DCPMMPerformanceSetting "Latency Optimized"
./OneCli config set IntelOptaneDCPMM.DCPMMPerformanceSetting "BW Optimized"
./OneCli config set IntelOptaneDCPMM.DCPMMPerformanceSetting "Balanced Profile"
```

Example out of the OneCLI config set command is shown in the following figure.

```
[root@aep lnvgy-utl-lxce-onecli]# ./OneCli config set IntelOptaneDCPMM.DCPMMPerformanceSetting "Latency Optimized"

Lenovo XClarity Essentials OneCLI lxce_onecli01o-3.0.0
(C) Lenovo 2013-2020 All Rights Reserved

OneCLI License Agreement and OneCLI Legal Information can be found at the following location:
  "/opt/lenovo/lnvgy-utl-lxce-onecli/Lic"

Invoking SET command ...
Connected to the node(1) by KCS
IntelOptaneDCPMM.DCPMMPerformanceSetting=Latency Optimized
Configure successfully, please reboot system.
Succeed.
```

Figure 3. OneCLI config set command (click the image for a larger version)

About the author

Jimmy Cheng is a performance engineer in the Lenovo Infrastructure Solutions Group (ISG) Performance Laboratory in Taipei, Taiwan. Jimmy joined Lenovo in December 2016. Prior to this, he worked on IBM POWER system assurance and validation, system integration, automation development as well as network performance. Jimmy holds a master's degree in Electronic and Computer Engineering from National Taiwan University of Science and Technology in Taiwan, and a bachelor's Degree in Computer Science and Engineering from Yuan-Ze University, Taiwan.

Thanks to the following people for their contributions to this project:

- Joe Jakubowski, Principal Engineer and Performance Architect in the Lenovo ISG Performance Laboratory in Morrisville, NC
- Nathan Pham, Sr. Performance Engineer and Architect in the Lenovo ISG Performance Laboratory in Morrisville, NC
- Arunkumar Krishnamoorthy, Server Performance Development Engineer in the Lenovo ISG Performance Laboratory in Morrisville, NC
- David Watts, Lenovo Press

Related product families

Product families related to this document are the following:

- [Memory](#)

Notices

Lenovo may not offer the products, services, or features discussed in this document in all countries. Consult your local Lenovo representative for information on the products and services currently available in your area. Any reference to a Lenovo product, program, or service is not intended to state or imply that only that Lenovo product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any Lenovo intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any other product, program, or service. Lenovo may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

Lenovo (United States), Inc.
8001 Development Drive
Morrisville, NC 27560
U.S.A.
Attention: Lenovo Director of Licensing

LENOVO PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. Lenovo may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

The products described in this document are not intended for use in implantation or other life support applications where malfunction may result in injury or death to persons. The information contained in this document does not affect or change Lenovo product specifications or warranties. Nothing in this document shall operate as an express or implied license or indemnity under the intellectual property rights of Lenovo or third parties. All information contained in this document was obtained in specific environments and is presented as an illustration. The result obtained in other operating environments may vary. Lenovo may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any references in this publication to non-Lenovo Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this Lenovo product, and use of those Web sites is at your own risk. Any performance data contained herein was determined in a controlled environment. Therefore, the result obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

© Copyright Lenovo 2024. All rights reserved.

This document, LP1349, was created or updated on August 25, 2020.

Send us your comments in one of the following ways:

- Use the online Contact us review form found at:
<https://lenovopress.lenovo.com/LP1349>
- Send your comments in an e-mail to:
comments@lenovopress.com

This document is available online at <https://lenovopress.lenovo.com/LP1349>.

Trademarks

Lenovo and the Lenovo logo are trademarks or registered trademarks of Lenovo in the United States, other countries, or both. A current list of Lenovo trademarks is available on the Web at <https://www.lenovo.com/us/en/legal/copytrade/>.

The following terms are trademarks of Lenovo in the United States, other countries, or both:

Lenovo®

ThinkSystem®

The following terms are trademarks of other companies:

Intel®, Intel Optane™, and Xeon® are trademarks of Intel Corporation or its subsidiaries.

Other company, product, or service names may be trademarks or service marks of others.