



# **Intel® Server S2600(WF)(BP) Family BIOS Settings and Options for Intel® Optane™ Persistent Memory**

***User Guide***

**Rev 1.0**

***October 2020***

<Blank page>

## ***Document Revision History***

| <b>Date</b>  | <b>Revision</b> | <b>Changes</b> |
|--------------|-----------------|----------------|
| October 2020 | 1.0             | First release. |
|              |                 |                |

## ***Disclaimers***

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software, or service activation. Performance varies depending on system configuration. No computer system can be absolutely secure. Check with your system manufacturer or retailer or learn more at intel.com.

You may not use or facilitate the use of this document in connection with any infringement or other legal analysis concerning Intel products described herein. You agree to grant Intel a non-exclusive, royalty-free license to any patent claim thereafter drafted which includes subject matter disclosed herein.

No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document.

The products described may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Intel disclaims all express and implied warranties, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and non-infringement, as well as any warranty arising from course of performance, course of dealing, or usage in trade.

Copies of documents which have an order number and are referenced in this document may be obtained by calling 1-800-548-4725 or by visiting [www.intel.com/design/literature.htm](http://www.intel.com/design/literature.htm).

Intel, the Intel logo, Xeon, and Optane are trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries.

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

© Intel Corporation

# Table of Contents

|   |           |
|---|-----------|
| <b>1. Introduction .....</b>  | <b>7</b>  |
| 1.1 Supplemental Product Information .....                          | 8         |
| 1.2 Target Audience.....  | 8         |
| 1.3 Applicable Intel Server Product.....                            | 8         |
| <b>2. BIOS Configuration Settings Overview.....</b>                 | <b>10</b> |
| 2.1 Test Configuration .....  | 10        |
| 2.2 BIOS Setup Menu .....   | 10        |
| 2.3 BIOS Main Menu Page .....                                       | 12        |
| 2.4 Memory Configuration.....                                       | 12        |
| 2.4.1 Memory Operating Speed Selection.....                         | 12        |
| 2.4.2 Volatile Memory Mode.....                                     | 13        |
| 2.4.3 DCPMM Error Injection .....                                   | 14        |
| 2.4.4 Snoopy Mode for 2LM.....                                      | 14        |
| 2.4.5 Snoopy Mode for AD .....                                      | 15        |
| 2.4.6 NVM Performance Setting .....                                 | 16        |
| 2.4.7 DIMM Information.....   | 16        |
| 2.5 Intel® Optane™ PMem Main Setup Screen locations .....           | 17        |
| 2.6 The Intel® Optane™ Persistent Memory Configuration Options..... | 19        |
| 2.6.1 DIMMs Sub-menu .....  | 19        |
| 2.6.1.1 Monitor Health Sub-menu.....                                | 21        |
| 2.6.1.2 Updating Firmware Sub-menu .....                            | 22        |
| 2.6.1.3 Configuring Security Sub-menu.....                          | 23        |
| 2.6.2 Regions Sub-menu .....  | 25        |
| 2.6.3 Namespace Sub-menu .....                                      | 29        |
| 2.6.4 Total Capacity Sub-menu .....                                 | 29        |
| 2.6.5 Diagnostics Sub-menu.....                                     | 30        |
| 2.6.6 Preferences Sub-menu.....                                     | 30        |
| <b>3. FAQ .....</b>   | <b>31</b> |
| 3.1 Intel® Optane™ PMem Startup FAQ .....                           | 31        |
| 3.2 Intel® Optane™ PMem 100 Series Firmware FAQ .....               | 32        |
| <b>Appendix A. Glossary .....</b>                                   | <b>33</b> |
| <b>Appendix B. Getting Help .....</b>                               | <b>34</b> |

## List of Figures

|  |    |
|--|----|
| Figure 1. Intel® Optane™ Persistent Memory 100 Series .....              | 7  |
| Figure 2. Intel® Optane™ PMem Population configurations per CPU .....    | 10 |
| Figure 3. BIOS Setup Menu Total Effective Memory in Memory Mode .....    | 11 |
| Figure 4. BIOS Setup Menu Total Effective Memory in App-Direct Mode..... | 11 |
| Figure 5. BIOS Main page Total Reported Memory .....                     | 12 |
| Figure 6. Memory Operating Speed Selection .....                         | 13 |
| Figure 7. Volatile Memory Mode.....                                      | 13 |
| Figure 8. Intel® Optane™ PMem Error Injection.....                       | 14 |
| Figure 9. Snoopy mode for 2LM .....                                      | 15 |
| Figure 10. Snoopy mode for AD.....                                       | 15 |
| Figure 11. NVM Performance Settings .....                                | 16 |
| Figure 12. DIMM Information .....  | 17 |
| Figure 13. PCI Configuration .....                                       | 17 |
| Figure 14. UEFI Option ROM Control.....                                  | 18 |
| Figure 15. Memory Controller.....  | 18 |
| Figure 16. The Intel® Optane™ DC Persistent Memory Configuration .....   | 19 |
| Figure 17. DIMM Specific Information.....                                | 19 |
| Figure 18. DIMM Details.....   | 20 |
| Figure 19. DIMM Details (cont.).....                                     | 20 |
| Figure 20. DIMM Details (cont.).....                                     | 21 |
| Figure 21. Monitor Health Screen .....                                   | 21 |
| Figure 22. Monitor Health Screen (cont.) .....                           | 22 |
| Figure 23. Updating Firmware .....                                       | 22 |
| Figure 24. Configuring Security.....                                     | 23 |
| Figure 25. Unlocking the Intel® Optane™ PMem .....                       | 24 |
| Figure 26. Boot Manager OS load .....                                    | 24 |
| Figure 27. Create Goal Config .....                                      | 25 |
| Figure 28. Creating a Goal Config at the Platform level.....             | 26 |
| Figure 29. Create goal config at the socket level.....                   | 26 |
| Figure 30. Memory mode vs. reserved mode .....                           | 27 |
| Figure 31. App-Direct Mode Config Setup .....                            | 27 |
| Figure 32. Region Configuration after reboot .....                       | 28 |
| Figure 33. Region ID Summary .....                                       | 28 |
| Figure 34. Creating Namespace .....                                      | 29 |
| Figure 35. Total Capacities .....  | 29 |
| Figure 36. Diagnostics submenu.....                                      | 30 |
| Figure 37. Preferences .....   | 30 |

# 1. Introduction

---

Intel server systems compatible with the 2<sup>nd</sup> generation Intel® Xeon® processor Scalable Family introduced support for the Intel® Optane™ Persistent Memory 100 Series.



**Figure 1. Intel® Optane™ Persistent Memory 100 Series**

---

Note: Intel BIOS Setup menu screens may still identify “Intel® Optane™ DC Persistent Memory”. In early 2020 Intel changed the naming of these devices to “Intel® Optane™ Persistent Memory 100 Series”.

---

The Intel® Optane™ Persistent Memory has two operating modes: Memory Mode (MM) and App Direct (AP) Mode.

Memory Mode is great for large memory capacity and does not require application changes which makes Intel® Optane™ persistent memory easy to adopt. In Memory Mode, the Intel® Optane™ persistent memory extends the amount of available volatile memory visible to the Operating System. DRAM is used as Cache for the Intel® Optane™ persistent memory. The CPU memory controller uses the DRAM as cache and the Intel® Optane™ persistent memory as addressable main memory.

In App Direct Mode, software and applications can talk directly to the Intel® Optane™ persistent memory, which reduces complexity in the stack. There is the option of having App Direct Mode use legacy storage APIs. This allows it to act like an SSD and can boot an OS. The operating system sees Intel® Optane™ persistent memory and DRAM as two separate pools of memory. It is persistent like storage, byte addressable like memory, cache coherent which extends the usage of persistent memory outside the local node, and consistent low latency supporting larger datasets.

The implementation of Intel® Optane™ Persistent Memory (PMem) modules may require certain system BIOS configuration setting changes depending on the desired operating mode. In Memory Mode, basic functionality runs “out of the box” and requires little if any, BIOS configuration. However, using App Direct mode and for obtaining the best performance from these devices, various optimization settings within the system BIOS must be set.

This document provides guidance for configuration of the system BIOS settings to support the desired operating mode of the Intel® Optane™ Persistent Memory. See section 1.1 for additional information.

## 1.1 Supplemental Product Information

The following list of Intel documents can be downloaded and referenced for additional product and setup information.

- 1.) Intel Corporation. (2020, Mar) Intel® Server Board S2600BP Intel Compute Module HNS2600BP Product Family Technical Product Specification  
<https://www.intel.com/content/www/us/en/support/articles/000024315/server-products/server-boards.html>
- 2.) Intel Corporation. (2020, Aug) Intel® Server Board S2600WF Product Family Technical Product Specification <https://www.intel.com/content/www/us/en/support/articles/000023750/server-products/server-boards.html>
- 3.) Intel Corporation. (2020, May) Intel® Server Board S2600 Family BIOS Setup Utility Specification Rev 1.17 <https://cdrdv2.intel.com/v1/dl/getContent/611025>
- 4.) Intel Corporation. (2020, June) Intel® Server Board S2600 Family BIOS External Product Specification (EPS) Rev 1.16 <https://cdrdv2.intel.com/v1/dl/getContent/611024>
- 5.) [Spec.org CPU2017 Platform Settings for Supermicro Systems](http://spec.cs.miami.edu/cpu2017/flags/Supermicro-Platform-Settings-V1.2-CLX-revF.html)  
<http://spec.cs.miami.edu/cpu2017/flags/Supermicro-Platform-Settings-V1.2-CLX-revF.html>
- 6.) Intel Corporation. (2020, July) Implementing Persistent Memory BIOS Optimizations for Memory Mode. <https://cdrdv2.intel.com/v1/dl/getContent/630736>
- 7.) Intel Corporation. (2020, February) Intel® Optane™ Persistent Memory (PMem) – BIOS settings <https://www.intel.com/content/dam/support/us/en/documents/memory-and-storage/data-center-persistent-mem/Intel-Optane™-DC-Persistent-Memory-BIOS-settings.pdf>
- 8.) Intel Corporation. (2020, June) Intel® Optane™ Data Center Persistent Memory Operations Guide <https://cdrdv2.intel.com/v1/dl/getContent/610782?wapkw=Peristent%20Memory%20Operations%20guide>

## 1.2 Target Audience

This document is intended for system integrators and system users who are responsible for configuring, upgrading, and optimizing their Intel server system when configured with Intel® Optane™ Persistent Memory (PMem) 100 Series modules.

## 1.3 Applicable Intel Server Product

This document applies to the following Intel server product families when configured with Intel® Optane™ Persistent Memory 100 Series:

---

Note: Intel server boards, systems, and compute modules for the defined product families below do not have support for Intel® Optane™ Persistent Memory 200 Series.

---

### Intel® Server S2600WF Product Family:

- All Intel and non-Intel server systems integrated with one of the following Intel server boards:  
Intel® Server Board S2600WF Product Family
  - iPC – S2600WF0R
  - iPC – S2600WFTR
  - iPC – S2600WFQR
- Intel® Server System R1000WF Product Family
  - iPC – R1304WFxxxR
  - iPC – R1208WFxxxR



- Intel® Server System R2000WF Product Family
  - iPC – R2308WFxxxR
  - iPC – R2312WFxxxR
  - iPC – R2208WFxxxR
  - iPC – R2224WFxxxR

#### **Intel® Server S2600BPR product family**

- All Intel and non-Intel server systems integrated with one of the following Intel server boards:  
Intel® Server Board S2600BP Product Family
  - iPC – BBS2600BPBR
  - iPC – BBS2600BPQR
  - iPC – BBS2600BPSR
- Intel® Compute Module HNS2600BP Product Family
  - iPC – HNS2600BPBR
  - iPC – HNS2600BPB24R
  - iPC – HNS2600BPSR
  - iPC – HNS2600BPS24R
  - iPC – HNS2600BPQR
  - iPC – HNS2600BPQ24R
  - iPC – HNS2600BPBLCR
  - iPC – HNS2600BPBLC24R

iPC = Intel Product Code

Use of Intel® Optane™ persistent memory 100 series within the defined Intel server product families is only supported when the following are true:

- Product order code ends with an 'R'
- Installed processor(s) are 2nd Gen Intel® Xeon® processor Scalable family

Intel server boards, systems, and compute modules within the defined product families with non-(R) product codes do NOT have support for Intel® Optane™ persistent memory.

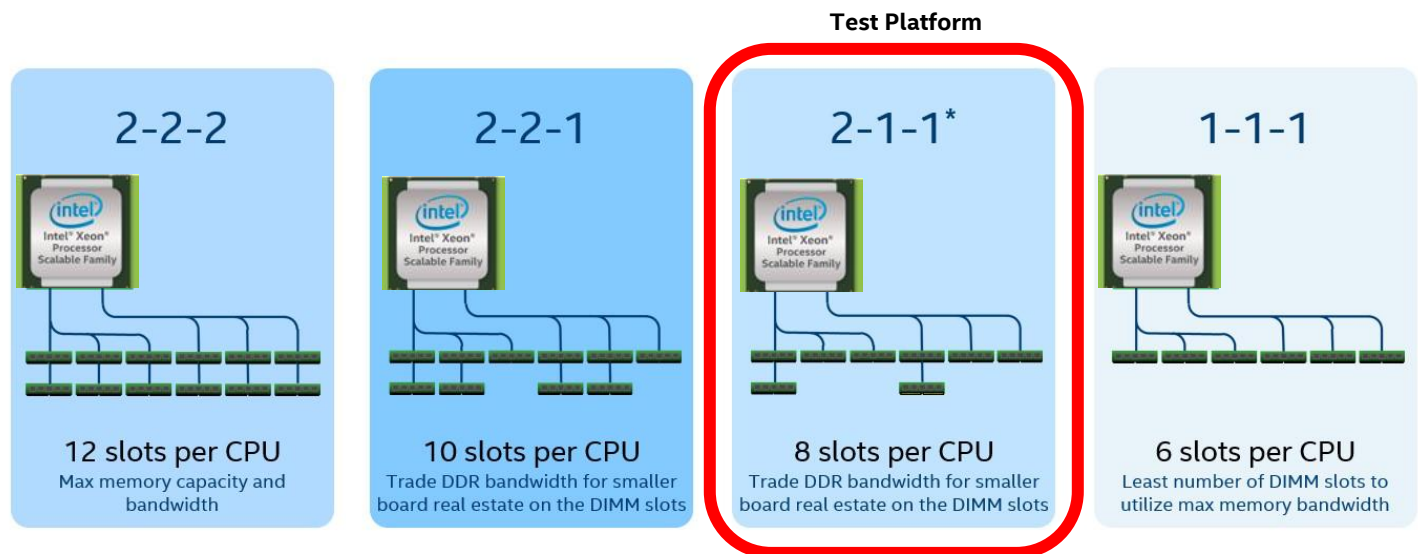
## 2. BIOS Configuration Settings Overview

This chapter provides an overview of BIOS Setup configuration options that can or must be set to provide best performance and operation for Intel servers configured with Intel® Optane™ Persistent Memory 100 series modules.

### 2.1 Test Configuration

The Intel platform used to develop the contents of this document was configured as follows:

- Intel server using an Intel® Server Board S2600BPR
- System BIOS 02.01.0012
- Two 2<sup>nd</sup> Gen Intel® Xeon® processor scalable family
- 192GB DDR4 DRAM DIMMs – 12 x 16GB Memory Modules
- 2TB Intel® Optane™ PMem 100 series – 4 x 512GB in a 2-1-1 configuration (see Figure 2)



**Figure 2. Intel® Optane™ PMem Population configurations per CPU**

In the above figure, a 2-1-1 configuration has all DDR4 DIMMs populated in the top memory row (Slot #1 for each memory channel) and Intel® Optane™ PMem are populated in the bottom memory row (Slot#2 for each memory channel with a 2<sup>nd</sup> memory slot).

Please reference the Intel® Server Board S2600BP or Intel® Server Board S2600WF Technical Product Specifications for full memory population rules

### 2.2 BIOS Setup Menu

On Intel server systems, the BIOS Setup utility is accessible when the system is first powered on or after a system reset during the system POST (Power On Self-Test) and before the Operating System begins to boot. During POST, when prompted, the BIOS Setup utility is accessed by pressing the <F2> key on the keyboard.

The BIOS setup Main page will display the total effective memory size. The memory size reported will be dependent upon the Intel® Optane™ PMem operating mode. In Figure 3, the Intel® Optane™ PMem is configured in memory mode and all DDR memory is configured as cache, therefore the total effective memory is only the amount of Intel® Optane™ PMem (2TB). In Figure 4, the Intel® Optane™ PMem is

configured in App-Direct mode, and the total effective memory is larger, with the DDR seen as system memory instead of cache.

Note: BIOS Setup screen shots provided in this document were captured using a system within the Intel® Server System S2600BP product family and System BIOS 02.01.0012. BIOS Setup screens may differ slightly when using another supported Intel server system or if an earlier revision of system BIOS is installed.

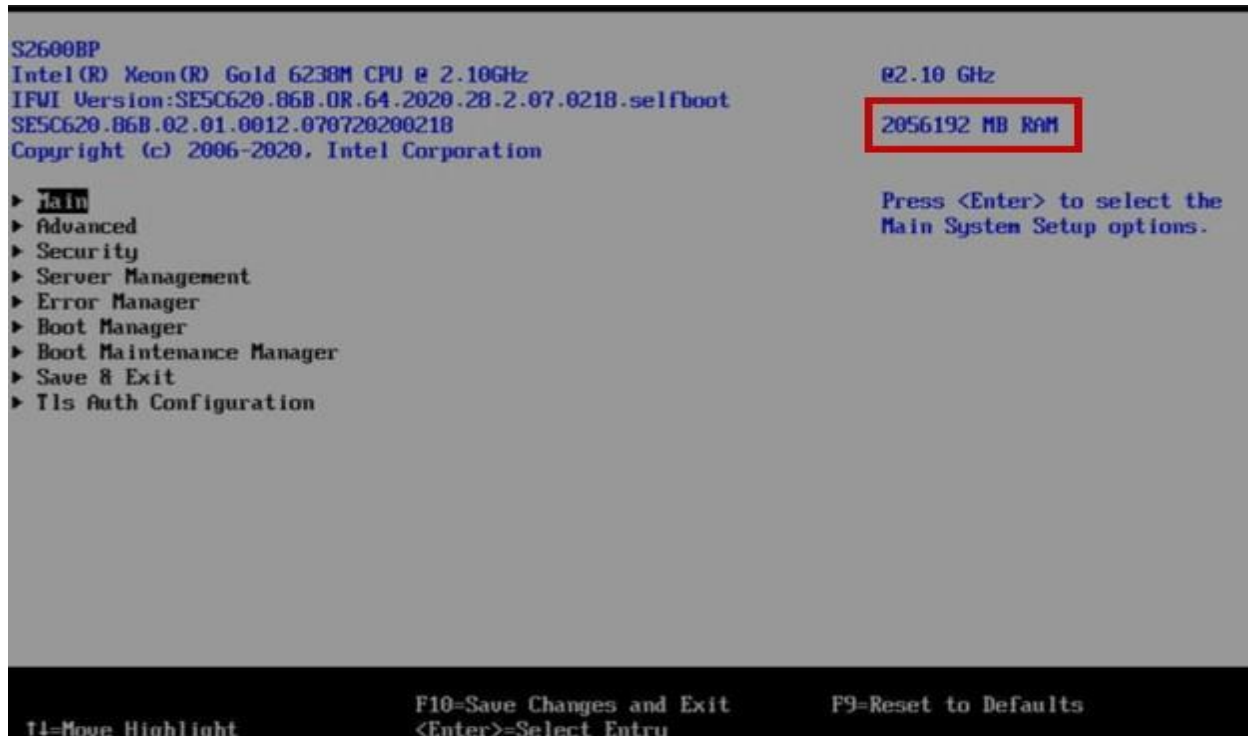


Figure 3. BIOS Setup Menu Total Effective Memory in Memory Mode

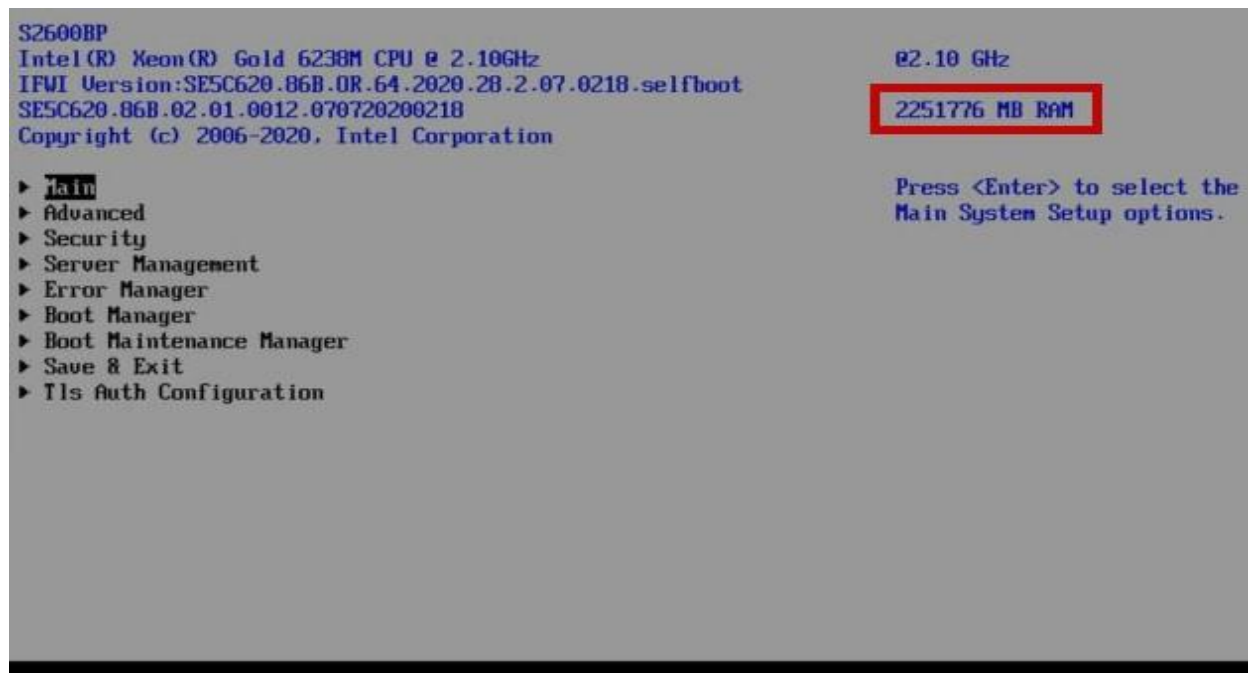


Figure 4. BIOS Setup Menu Total Effective Memory in App-Direct Mode

## 2.3 BIOS Main Menu Page

The BIOS main page will report the total memory configuration. It includes the total amount of DDR memory and the total Intel® Optane™ PMem. This is all memory reported to BIOS during boot and is independent of the mode the Intel® Optane™ PMem may be operating in.

Figure 5 shows the *Memory Configuration* field reporting four numbers: Total DDR4 memory installed, total Intel® Optane™ PMem memory installed, and memory available for each of the two PMem operating modes: Memory Mode and App-Direct mode.

The Intel® Optane™ PMem (DCPMM) status 2008GB-2008GB-0GB shown in Figure 5 is defined as follows:

- 1<sup>st</sup> number - 2008GB is total Intel® Optane™ PMem installed
- 2<sup>nd</sup> number - 2008GB is available in Memory Mode
- 3<sup>rd</sup> number - 0GB is available in App-Direct mode.



Figure 5. BIOS Main page Total Reported Memory

## 2.4 Memory Configuration

From the **Main Page**, navigate to the **Advanced Page**, then the **Memory Configuration**, which identifies various options for configuring Intel® Optane™ PMem in its various modes. The options need to be configured for operation and may optimize performance of the Intel® Optane™ PMem.

### 2.4.1 Memory Operating Speed Selection

The *Memory Operating Speed Selection* (see Figure 6) allows the user to select a specific speed at which memory operates. Only speeds that are supported are available; that is, the user can only specify speeds less than or equal to the auto selected memory speed. The Intel® Optane™ PMem 100 series operates at 2666Mhz or lower. If higher frequency of the DDR4 DIMMS is implemented, all speeds will reduce to the slowest speed. By default, this setting should be set to **Auto**.

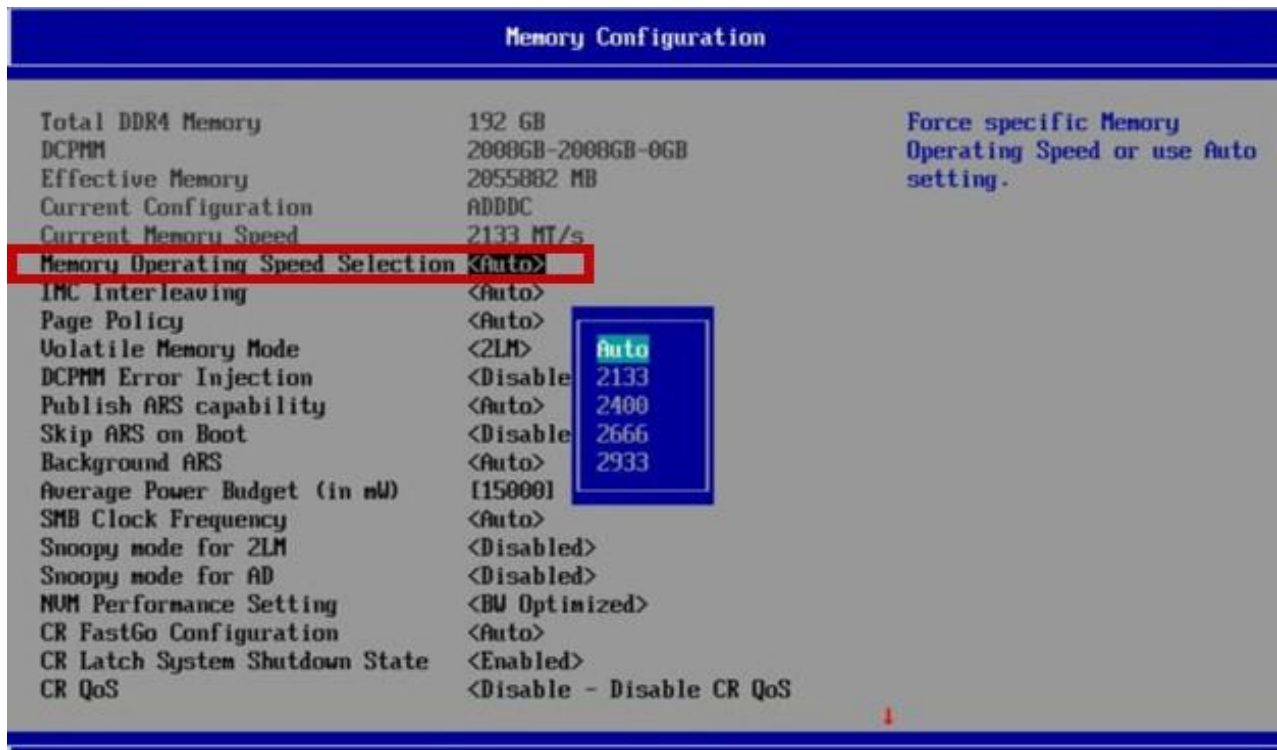


Figure 6. Memory Operating Speed Selection

## 2.4.2 Volatile Memory Mode

The *Volatile Memory Mode* setting (see Figure 7) is displayed when the BIOS supports Intel® Optane™ PMem. When *Auto* is selected, the BIOS will check for the 2LM (2 Layer Memory) configuration and if not found, will auto set to 1LM. The 2LM configuration is configured by the bios when set to Auto and when operating Intel® Optane™ PMem in Memory Mode. This allows 2LM to facilitate the DDR4 memory operating as cache. When Intel® Optane™ PMem is operating in App-Direct Mode, the bios will set this option to 1LM. The default is **Auto**.



Figure 7. Volatile Memory Mode



### 2.4.3 DCPMM Error Injection

The *DCPMM Error Injection* setting (see Figure 8) is an option that allows users to utilize a set of tools to simulate different kinds of errors and verify proper response paths to monitoring systems. Intel® Optane™ PMem related tools like IPMCTL and NDCTL have options to inject errors. For those tools to function with error injection, this setting needs to be set to Enabled. By Default, this option is set to **DISABLED**. Reference the IPMCTL help file and the NDCTL user guide for examples and settings. See section 1.1

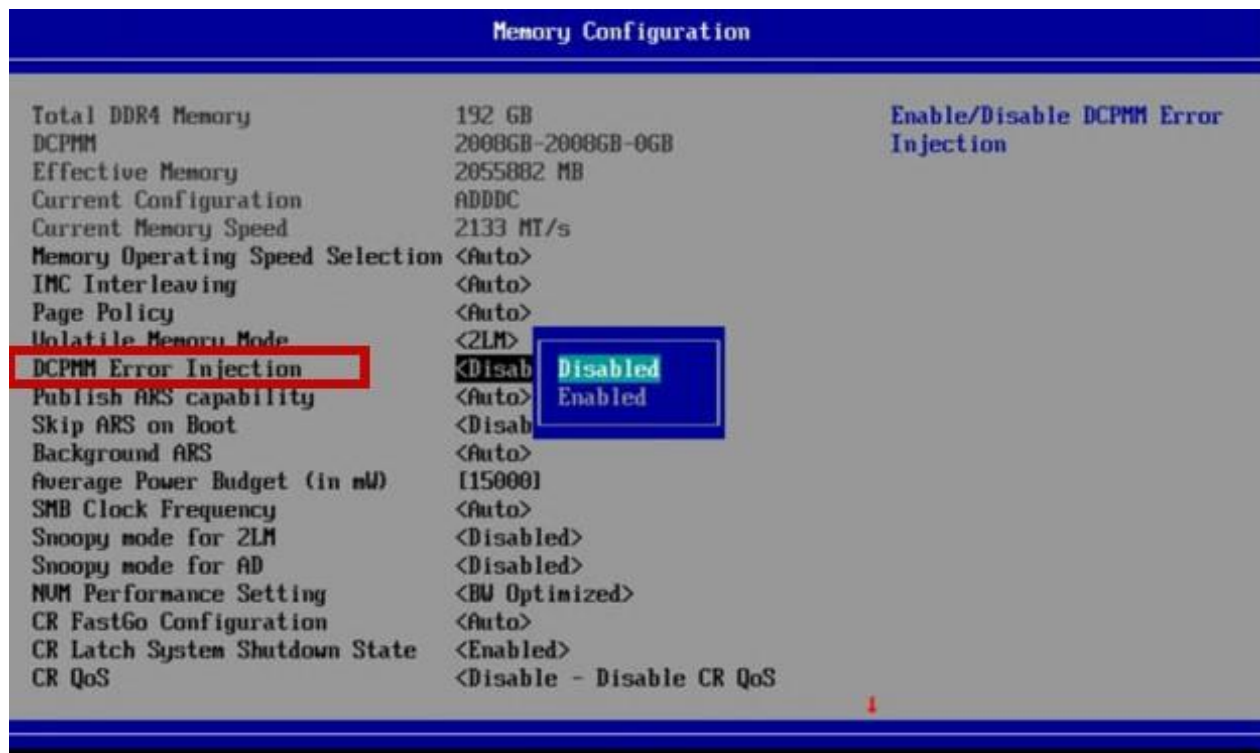


Figure 8. Intel® Optane™ PMem Error Injection

### 2.4.4 Snoopy Mode for 2LM

The *Snoopy Mode for 2LM* setting (see Figure 9) can be enabled for 2LM configurations to avoid directory updates to far-memory from non-NUMA optimized workloads. The default setting for Snoopy Mode for 2LM is **Disabled**. With this option Enabled best performance is obtained by Disabling the *Stale AtoS* setting discussed below.

The *Stale AtoS* setting is found under Advanced/UIP Configuration BIOS screen pages. This BIOS setting has two options “Enabled” and “Disabled” and is **Disabled** by default. This feature is related to in-memory directories. The in-memory directory has three states: **I**-Invalid, **A**-Snoop All, and **S**-Shared. The **I** (Invalid) state means the data is clean and does not exist in any other socket’s cache. The **A** (snoop All) state means the data may exist in another socket in an exclusive or modified state. The **S** (Shared) state means the data is clean and may be shared across one or more sockets caches.

In situations where a directory line is in an A state, all other sockets must be snooped because another socket may have the line in a modified state. If this is the case, the snoop will return the modified data. However, if the line is in an A state, it’s possible all snoops come back as a miss. This can happen when another socket reads the directory line but drops it from its own cache without modifying it.

With *Stale AtoS* enabled, in the situation where a line in A state returns only snoop misses, the line will transition to S state. This will cause subsequent reads to that directory line to encounter it in the S state and not have to snoop, saving latency and snoop bandwidth.

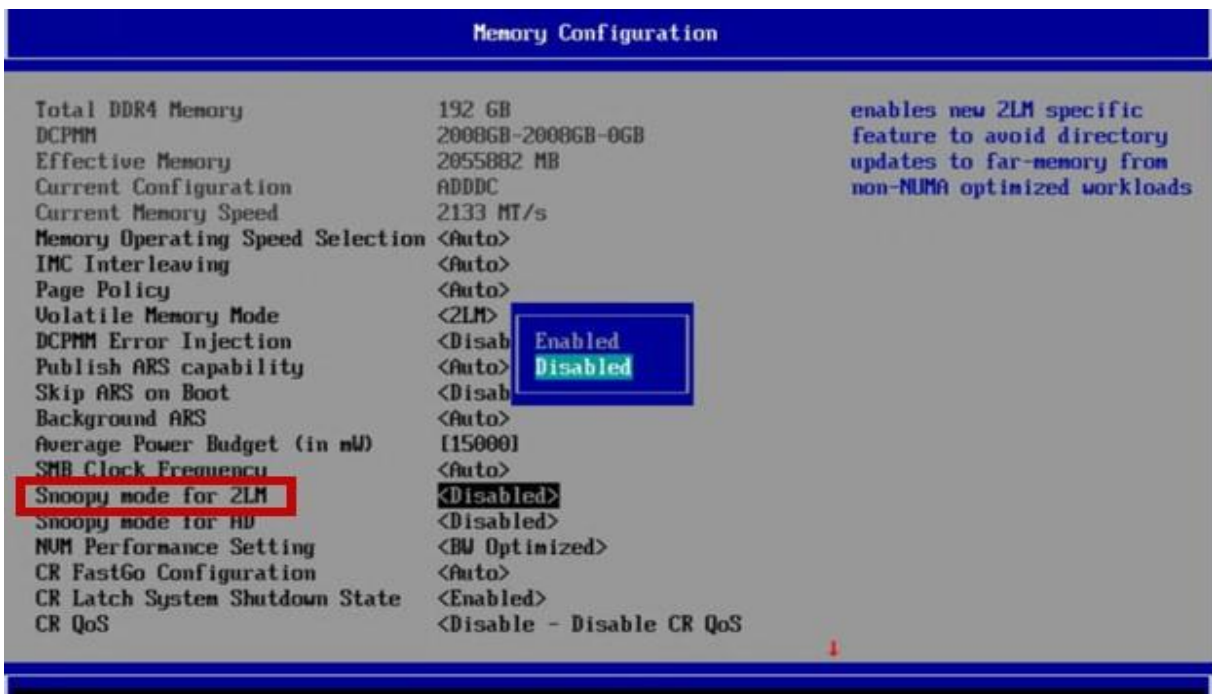


Figure 9. Snoopy mode for 2LM

#### 2.4.5 Snoopy Mode for AD

The *Snoopy Mode for AD* setting (see Figure 10) can be enabled for an App-Direct configuration to avoid directory updates to DDRT memory from non-NUMA optimized workloads. The default setting for *Snoopy Mode of AD* is **Disabled**. When enabling this option, best performance can be attained by having the *Stale AtoS* setting set to **Disabled**.

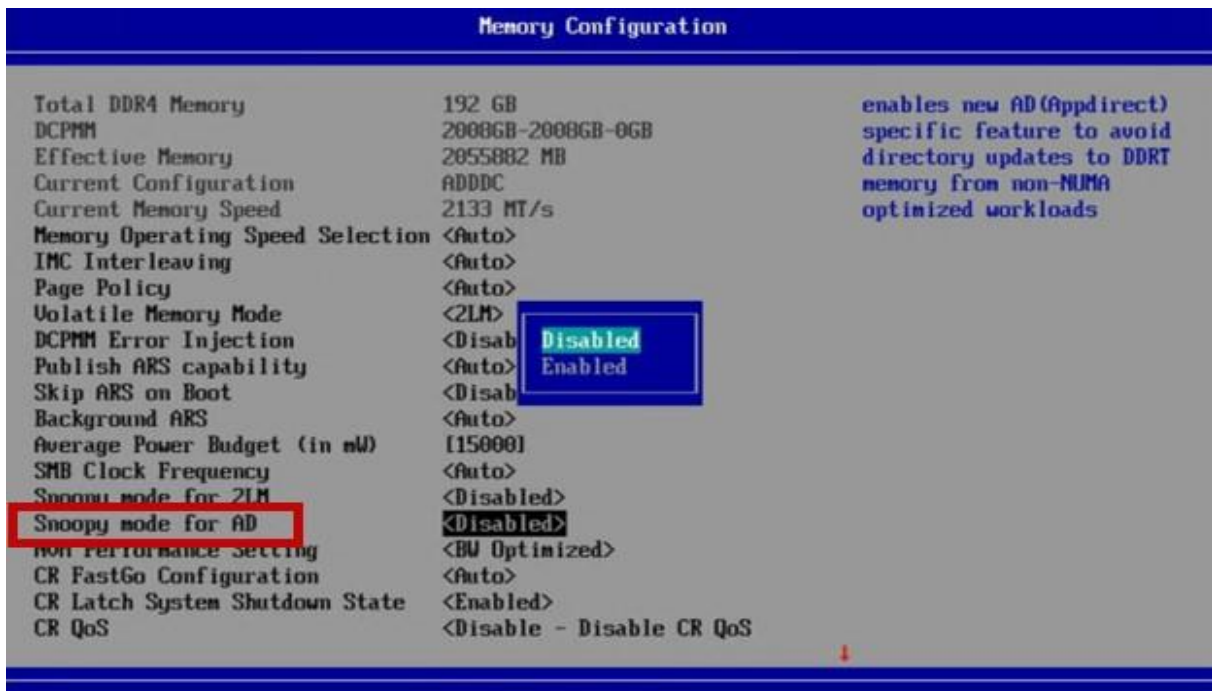


Figure 10. Snoopy mode for AD

## 2.4.6 NVM Performance Setting

The *NVM Performance Setting* (see Figure 11) is a baseline performance setting depending on the workload behavior. There are three options available for this setting as listed below:

- 1.) **Balanced Profile (Default):** 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.
- 2.) **BW Optimized:** Arbitrates between DRAM and Intel Optane™ PMem to maximize DRAM bandwidth on the memory bus.
- 3.) **Latency Optimized:** Arbitrates between DRAM and Intel Optane™ PMem to minimize DRAM latency on the memory bus.

---

Intel Optane™ PMem coexists with traditional DDR4 DIMMs on the same bus. Different arbitration profiles have been defined to determine the algorithm for when the memory controller switches between DRAM and the Intel® Optane™ Pmem memory transactions. For more information, please refer to the technical brief: "Implementing Persistent memory BIOS Optimizations for Memory Mode" (see section 1.1)

---

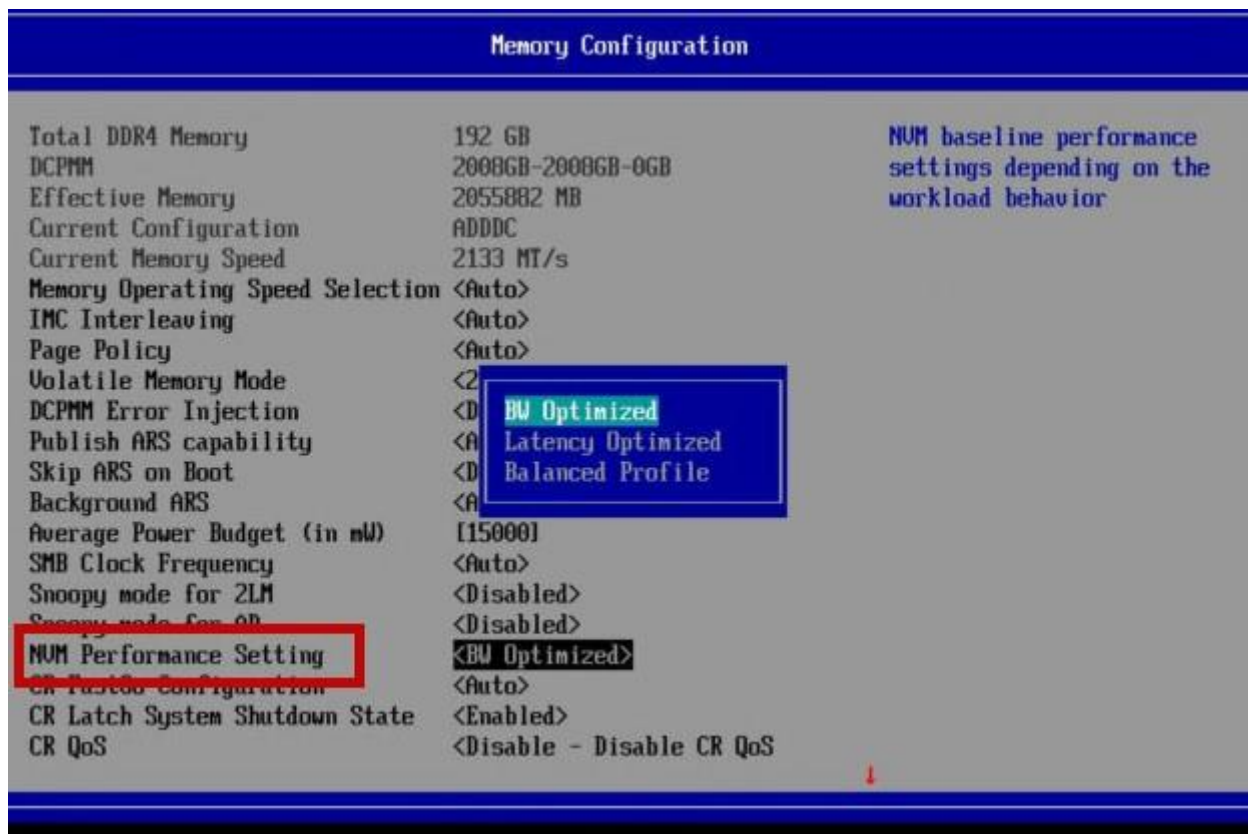


Figure 11. NVM Performance Settings

## 2.4.7 DIMM Information

The *DIMM Information* area of the Memory Configuration screen (see Figure 12) provides a view of the physical population of DDR memory and Intel® Optane™ PMem within the system. In Figure 12, the system is populated in a 2-1-1 configuration of Optane™ Pmem vs. DDR4 memory.



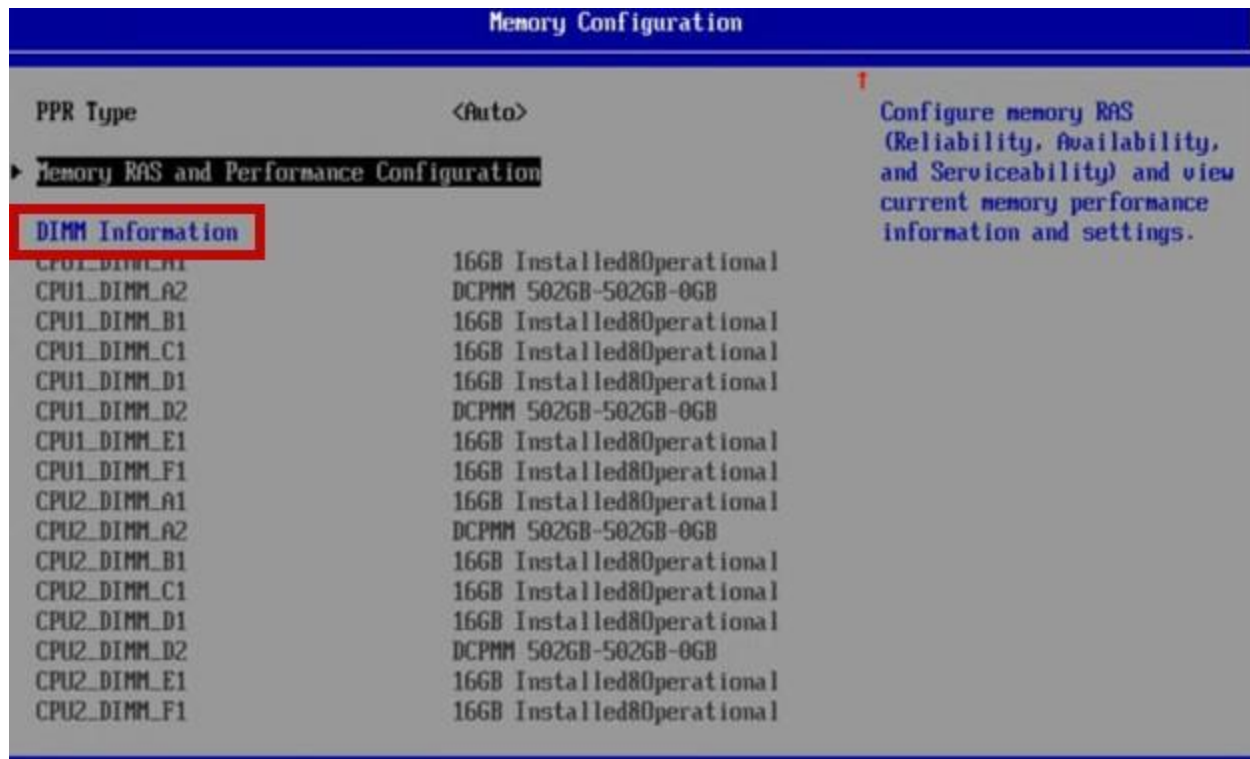


Figure 12. DIMM Information

## 2.5 Intel® Optane™ PMem Main Setup Screen locations

The Intel® Optane™ PMem main setup options are found in the **Intel® Optane™ DC Persistent Memory Configuration** Screen which is accessed by selecting the following BIOS Setup menu options: **Advanced** -> **PCI Configuration** -> **UEFI Option ROM Control** -> **Intel® Optane™ DC Persistent Memory Configuration** sub menus. (see Figures 13, 14, 15)

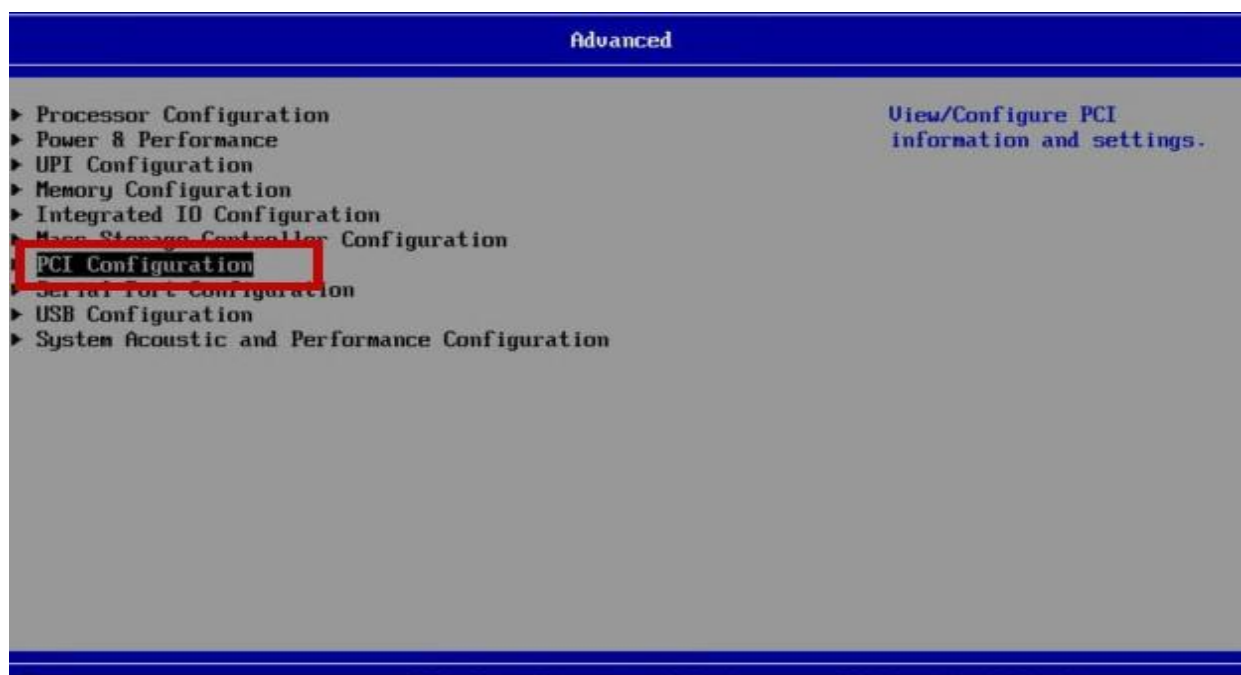


Figure 13. PCI Configuration

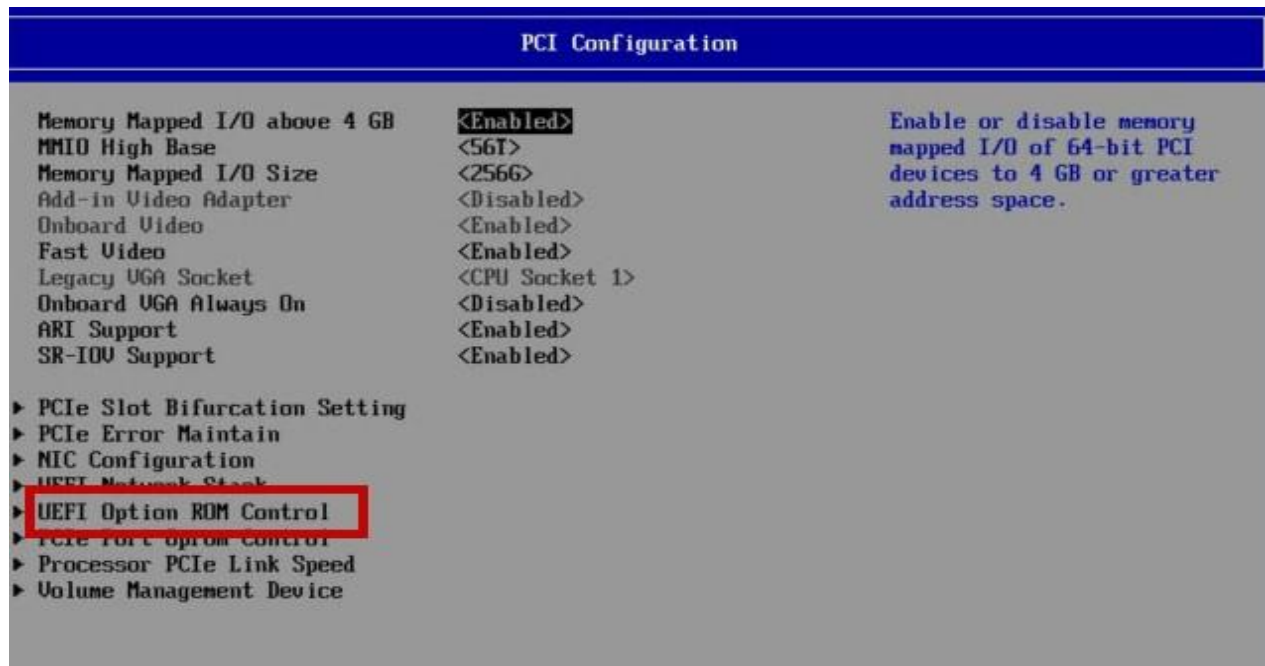


Figure 14. UEFI Option ROM Control

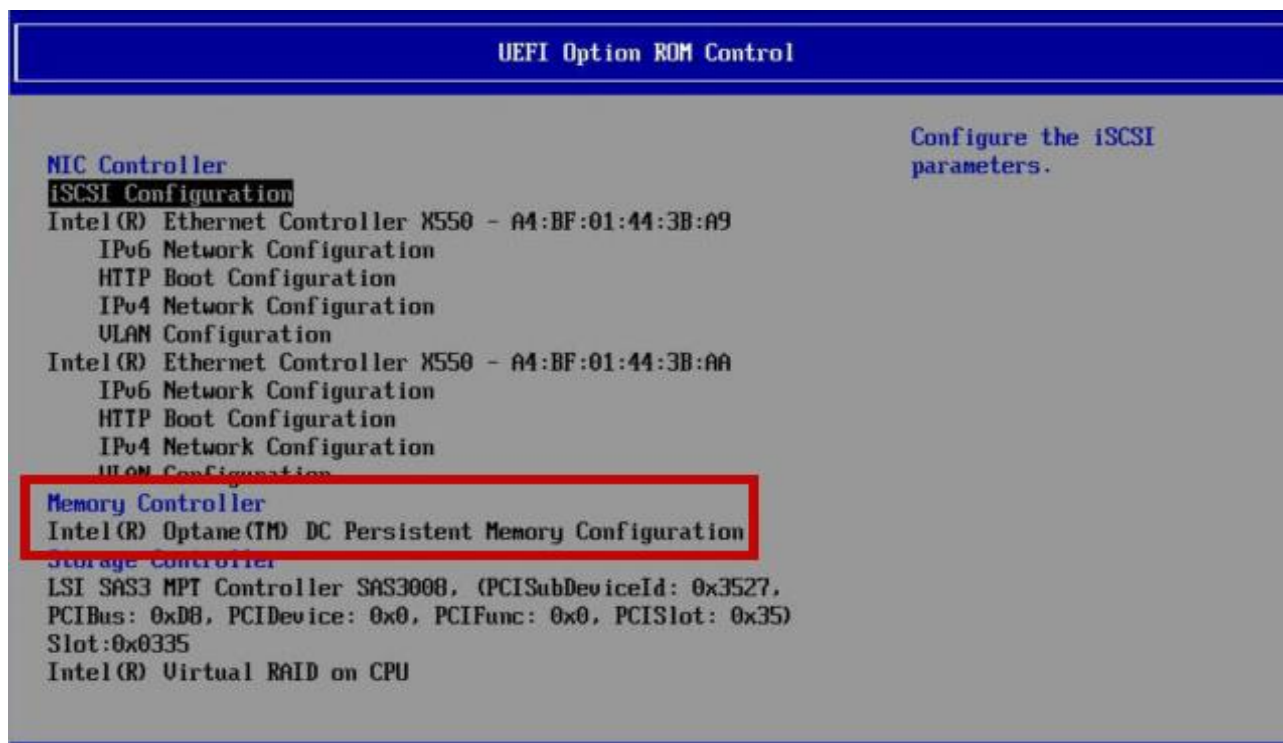


Figure 15. Memory Controller

## 2.6 The Intel® Optane™ Persistent Memory Configuration Options

This section describes the *Intel® Optane™ Persistent Memory Configuration* sub-menu options (see Figure 16). For additional information download the reference document *Intel® Optane™ DC persistent Memory BIOS settings* (see section 1.1).

The PMem Configuration screen displays the number of Intel® Optane™ PMem modules installed, provides a quick check on their operational health, and provides access points to several sub-menu options that provide various PMem information and configuration options.




Figure 16. The Intel® Optane™ DC Persistent Memory Configuration

### 2.6.1 DIMMs Sub-menu

The *DIMMs* sub-menu (see Figure 17) displays the ID # and related CPU identifier for each installed PMem device. In addition, the screen provides options to Monitor Pmem Health, Update Pmem Firmware, and Configure Pmem Security.

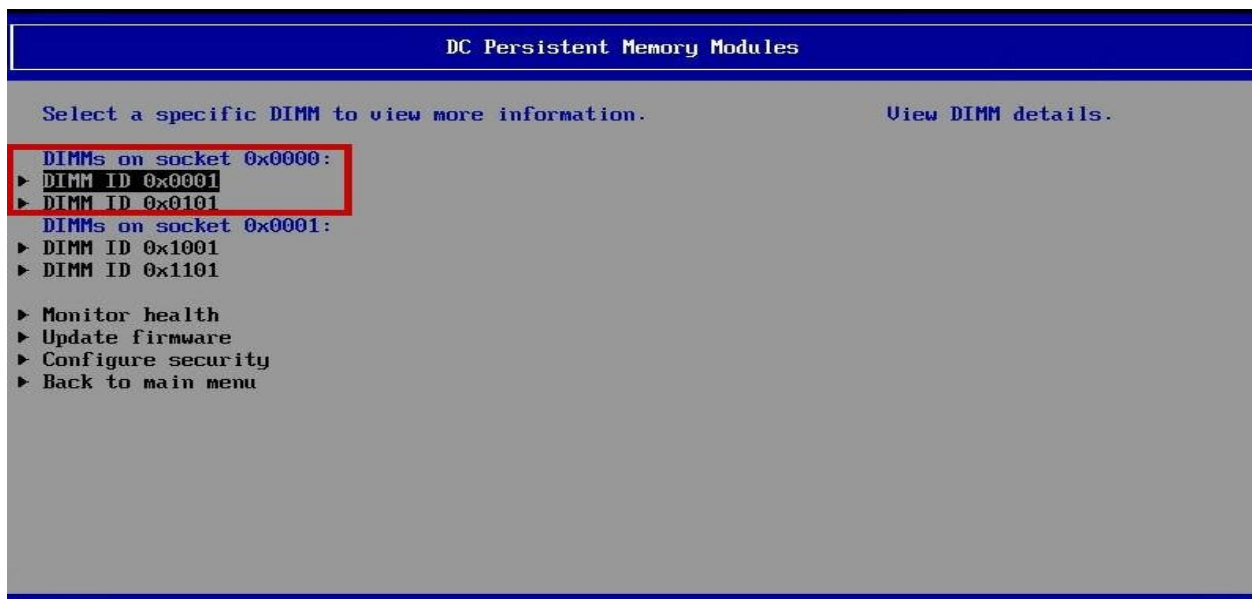


Figure 17. DIMM Specific Information

Selecting a specific PMem device (identified as a DIMM) from those listed will provide a screen showing information about the device (see Figure 18).

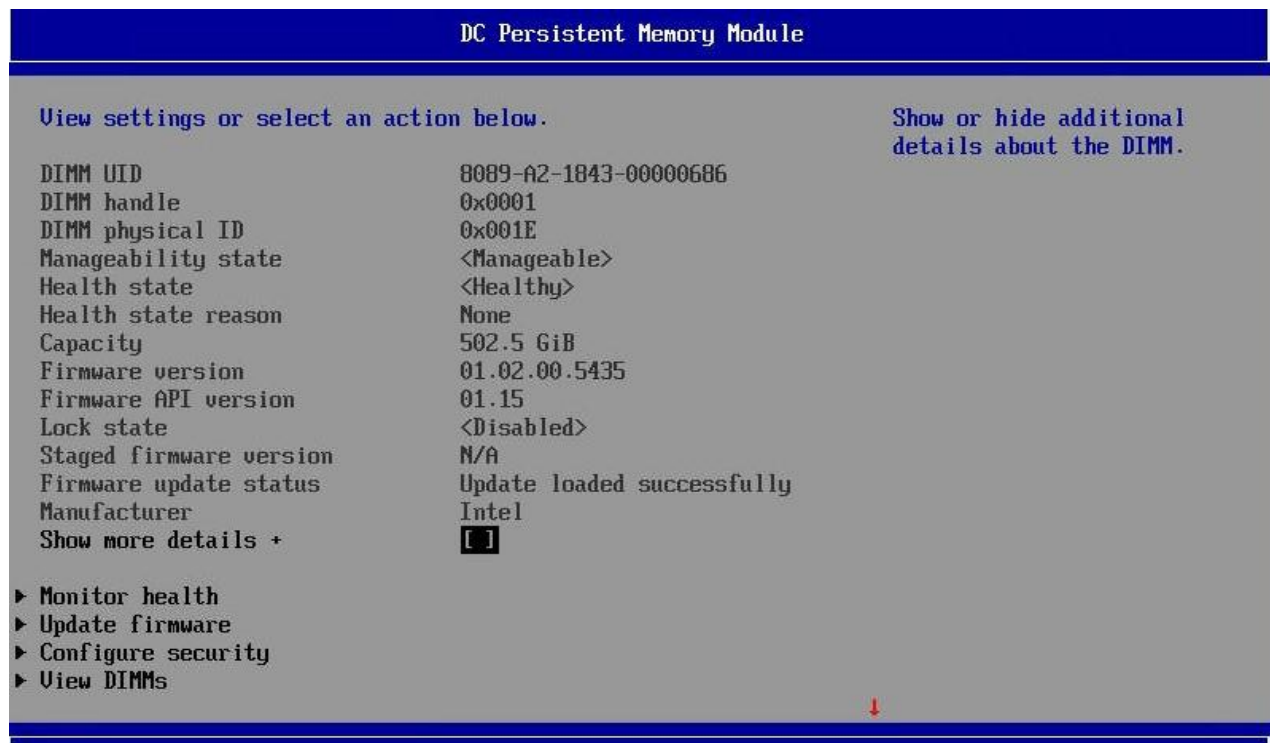


Figure 18. DIMM Details

Additional information related to the selected PMem device can be viewed by selecting the *Show more details* option on the screen (see figure 19).

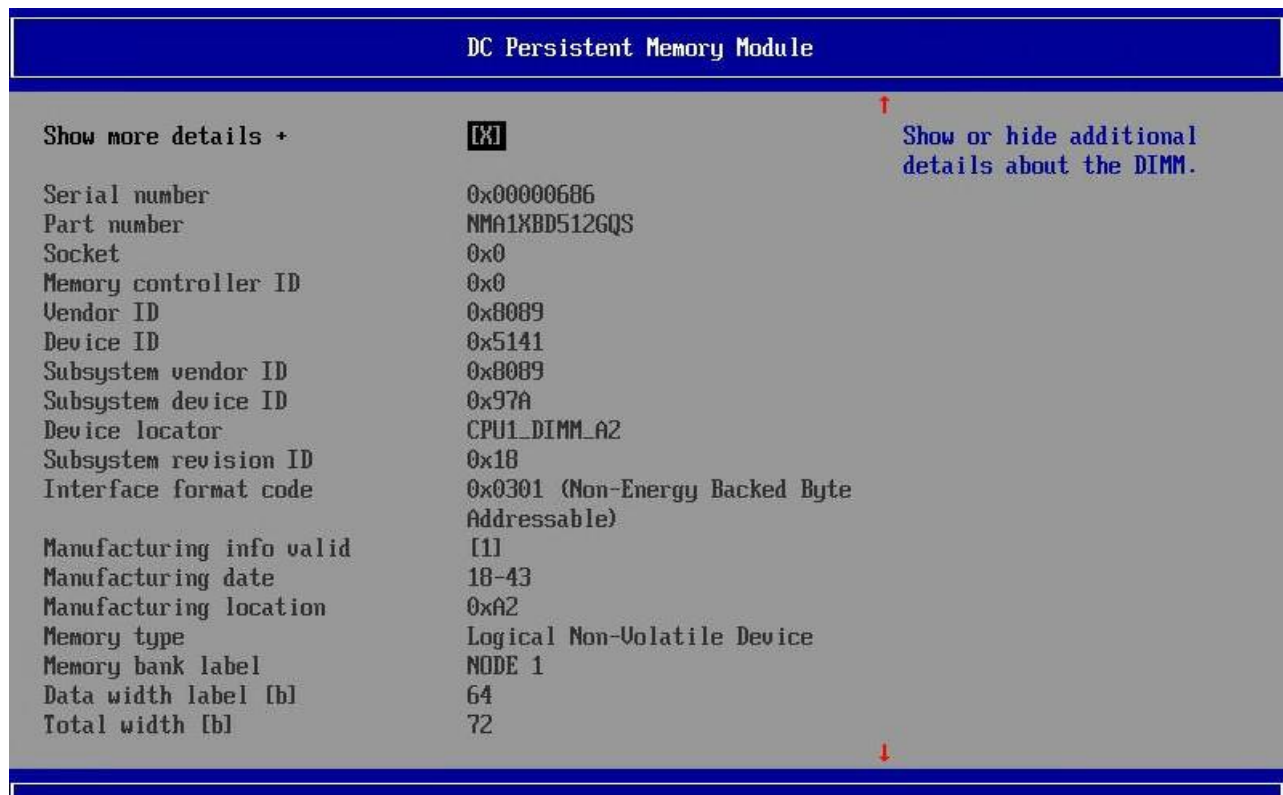


Figure 19. DIMM Details (cont.)

| DC Persistent Memory Module          |                         |  |
|--------------------------------------|-------------------------|--|
| Modes supported                      | Memory Mode, App Direct | ↑ Display DIMM health status and thresholds. |
| Boot status                          | Unknown                 |  |
| AIT DRAM enabled                     | <1>                     |  |
| Error injection enabled              | <0>                     |  |
| Media temperature injection enabled  | <0>                     |  |
| Software triggers enabled            | <0>                     |  |
| Software triggers enabled details    | None                    |  |
| Poison error injections counter      | [0]                     |  |
| Poison error clear counter           | [0]                     |  |
| Media temperature injections counter | [0]                     |  |
| Software triggers counter            | [0]                     |  |
| Master Passphrase Enabled            | [0]                     |  |
| ▶ Monitor health                     |                         |  |
| ▶ Update firmware                    |                         |  |
| ▶ Configure security                 |                         |  |
| ▶ View DIMMs                         |                         |  |
| ▶ Back to main menu                  |                         |  |

Figure 20. DIMM Details (cont.)

### 2.6.1.1 Monitor Health Sub-menu

The *Monitor Health* sub-menu is accessible from selection options located on both the *DIMMs* and *DIMM Details* sub-menu screens.

Upon selecting the *Monitor Health* option, a screen showing PMem health information is presented (see Figures 21 & 22). Each monitored sensor type displayed reports the current sensor reading along with any applicable threshold values.

| Monitor Health           |                          |                              |
|--------------------------|--------------------------|------------------------------|
| Sensor Type              | <Health>                 | Non-critical threshold alarm |
| Value                    | <Healthy>                |                              |
| Sensor Type              | <Controller temperature> |                              |
| Value                    | <57 C>                   |                              |
| Non-critical threshold   | [98]                     |                              |
| Critical lower threshold | [99]                     |                              |
| Critical upper threshold | [99]                     |                              |
| Fatal threshold          | [102]                    |                              |
| State                    | <Normal>                 |                              |
| Alarm enabled state      | <1>                      |                              |
| Sensor Type              | <Media temperature>      |                              |
| Value                    | <54 C>                   |                              |
| Non-critical threshold   | [82]                     |                              |
| Critical lower threshold | [83]                     |                              |
| Critical upper threshold | [83]                     |                              |
| Fatal threshold          | [85]                     |                              |
| State                    | <Normal>                 |                              |
| Alarm enabled state      | <1>                      |                              |

Figure 21. Monitor Health Screen



Scrolling to the bottom of the *Monitor Health* screen presents the option to *Modify Non-critical Thresholds* for several PMem sensors (see Figure 22). Care should be taken when choosing to change these settings from their default values. Setting them too aggressively may impact PMem functionality. Setting them too conservatively may impact the long-term reliability of the PMem modules.



Figure 22. Monitor Health Screen (cont.)

### 2.6.1.2 Updating Firmware Sub-menu

Periodically, Intel may release firmware updates for its Intel® Optane™ Persistent Memory. To update the PMem firmware, select the option *Updating Firmware* from either the *DIMMs* or *DIMM Details* sub-menu screens.

Upon selecting this option, a screen is displayed showing the current firmware version and an input field where the path for the file containing the firmware update can be found (see Figure 23). Once the file path is entered, select the *Update* option.




Figure 23. Updating Firmware

### 2.6.1.3 Configuring Security Sub-menu

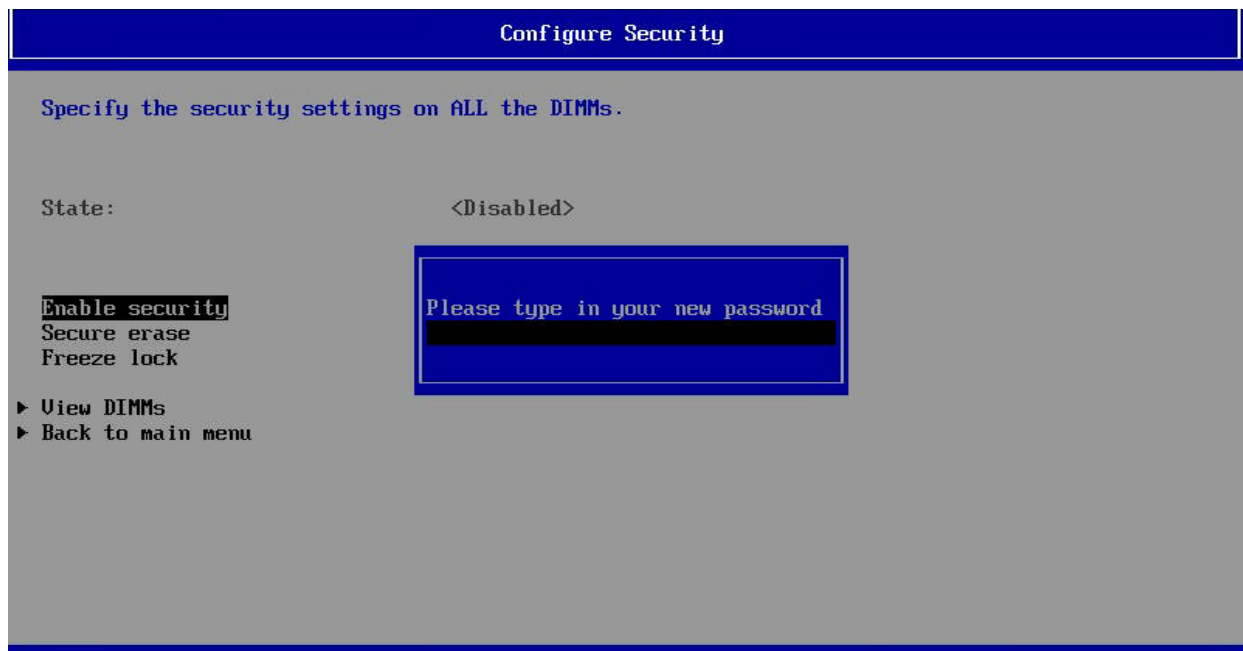
Intel® Optane™ Persistent Memory supports several security options. The setting in BIOS Setup to configure these options can be found by accessing the *Configuring Security* sub-menu option found on the *DIMMs* menu screen (see Figure 17).

The *Configure Security* sub-menu provides options to: enable/disable security, perform a secure erase, and freeze lock the Intel® Optane™ PMem modules.

- The Secure Erase option is used to erase all encrypted data on the PMem modules.
- The Freeze Lock option is used to lock the PMem security settings.

The *Secure Erase* and *Freeze Lock* options are only available when PMem Security is Enabled with a Password.

To enable the PMem security options and secure the modules, select **Enable security** (see Figure 24) and enter a password. With confirmation that Security is enabled, press F10 and “Y” when prompted to save the configuration and reboot the server.



**Figure 24. Configuring Security**

The PMem modules will be locked on every system boot after saving the configuration and restarting the server. During POST (Power-on self-test), the system will display a prompt to enter the password. Entering the correct password will unlock the PMem devices.

---

**CAUTION:** Do NOT forget or lose the password used to enable the PMem security features. Failing to enter the correct password when prompted will render the devices non-functional. The devices will be locked and **render any data they contain non-recoverable!**

There is no PMem Security password reset or Password clear mechanism available without entering the password originally set in BIOS Setup.

**Notice:** Enabling PMem Security within BIOS is **NOT** the same as enabling Security through the operating system or with the Intel® Optane™ PMem related tools (IPMCTL or NDCTL). **These security options are not interchangeable, and only one should be set, either BIOS or the OS tools**

---

To unlock and/or disable PMem security, re-enter bios setup and navigate to the DC Persistent Memory Module menus and select the **Configure Security** option.

The state of Intel® Optane™ PMem will be **<Locked>** (see Figure 25). The modules size will be seen and configured (either App Direct or Memory Mode) but not accessible by the OS, until the modules are unlocked.

To unlock, select **Unlock** and enter the password when prompted.

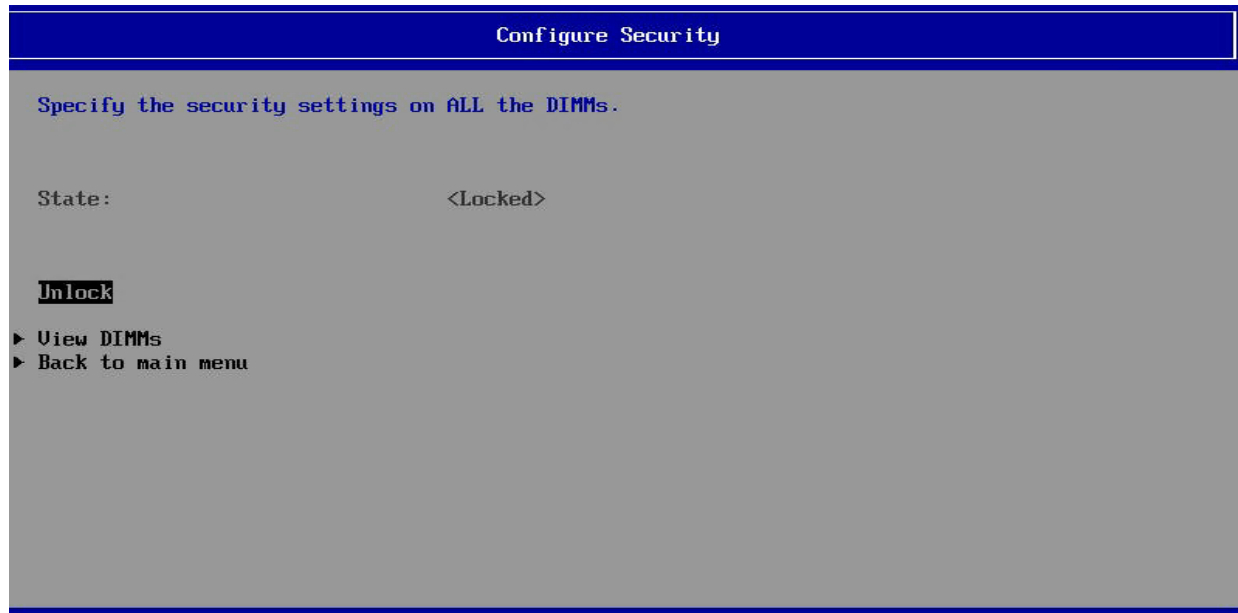


Figure 25. Unlocking the Intel® Optane™ PMem

Do **NOT** save and reboot the server. Navigate to the BIOS main menu, enter the **Boot** Manager, and select the boot device (Figure. 26).

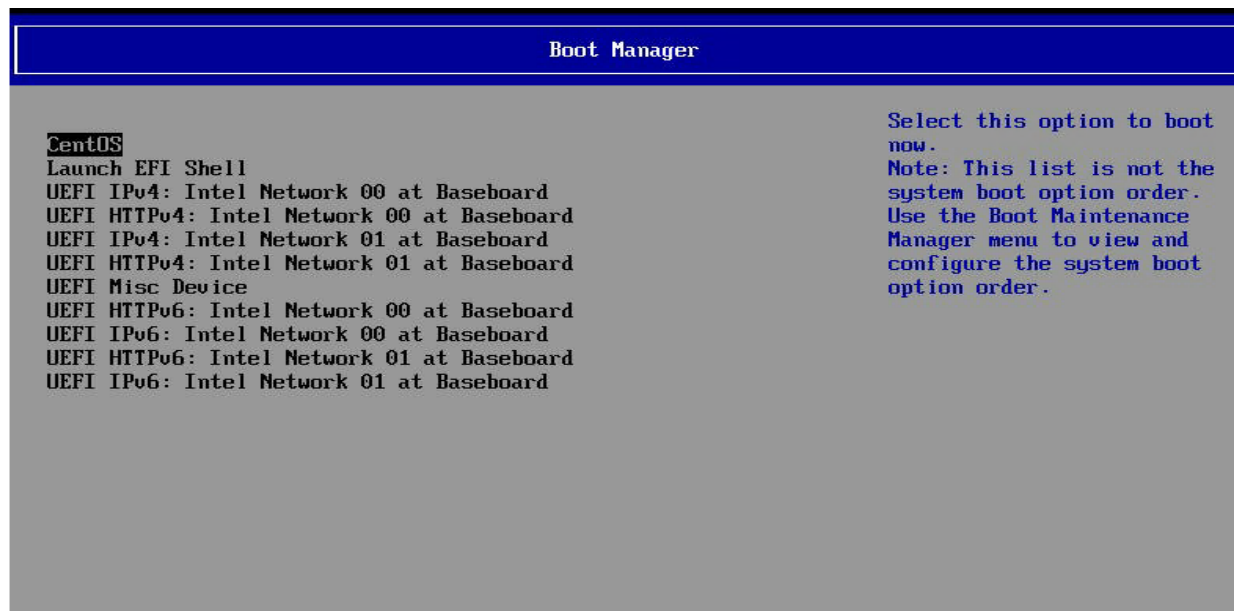


Figure 26. Boot Manager OS load



## 2.6.2 Regions Sub-menu

Both of the following sections, **Regions** and **Namespaces**, are used specifically to changing PMem modes of operation between Memory Mode, and App-Direct Mode.

To use Intel® Optane™ PMem devices in App-Direct mode, several steps need to be performed, including:

- creating a goal which identifies how much memory to use
- creating regions which are a group of one or more Intel® Optane™ PMem modules
- creating namespaces which define a contiguously addressed range of non-volatile memory, conceptually similar to a hard disk partition.

All these steps can be performed from the **Intel® Optane™ DC Persistent Memory Configuration** setup page (see Figure 16). Note that creating a goal is part of the Namespaces creation process.

From the Intel® Optane™ DC Persistent Memory Configuration setup page, select the **Regions** submenu option (see Figure 16), then select the **Create Goal Config** option (see Figure 27)



**Figure 27. Create Goal Config**

When creating a goal, there are several options presented on the Create Goal Config screen (see Figure 28).

- *Create goal config for*
  - Platform (**Default**) (see Figure 28) – Setting a common goal for all PMem devices
  - Socket – Setting a goal for specific PMem devices tied to a specific CPU Socket – CPU 0 (0x0000) or CPU 1 (0x0001) (see Figure 29)

Note: It is important to note configuring one socket for App Direct and the other socket for Memory Mode is not a supported configuration (Figure 26).

| Create Goal Config  |              |
|---|--------------|
| Select the scope of the new region then set the desired sizes.        |              |
| Create goal config for:   | <Platform>   |
| Reserved [%]:   | [0]          |
| Memory Mode [%]:  | [0]          |
| Persistent memory type:   | <App Direct> |
| Namespace Label version:  | <1.2>        |
| ▶ Create goal config<br>▶ Back to Regions menu<br>▶ Back to main menu |              |

Figure 28. Creating a Goal Config at the Platform level

| Create Goal Config  |              |
|---|--------------|
| Select the scope of the new region then set the desired sizes.        |              |
| Create goal config for:   | <Socket>     |
| Socket 0x0000   | [ ]          |
| Socket 0x0001   | [ ]          |
| Reserved [%]:   | [0]          |
| Memory Mode [%]:  | [0]          |
| Persistent memory type:   | <App Direct> |
| Namespace Label version:  | <1.2>        |
| ▶ Create goal config<br>▶ Back to Regions menu<br>▶ Back to main menu |              |

Figure 29. Create goal config at the socket level

- **Reserved [%]** - The reserved space is any amount of capacity (0% to 100%) that will not be mapped into the system physical address space. (see Figure 30)
- **Memory Mode [%]** - The memory mode percentage has a twofold usage. When converting from App-Direct back to Memory Mode, the memory mode percentage is converted from 0% to 100% which changes the mode to 100% memory mode. Alternatively, if a mixed mode configuration is required, the percentage of memory mode would be assigned. (see Figure 30).

For example, in a 70% App-Direct mode, 30% Memory mode configuration, 30% of the Intel® Optane™ PMem modules would be in memory mode (all DDR4 modules will be assigned as Cache) and 70% of the modules will be converted to persistent storage.

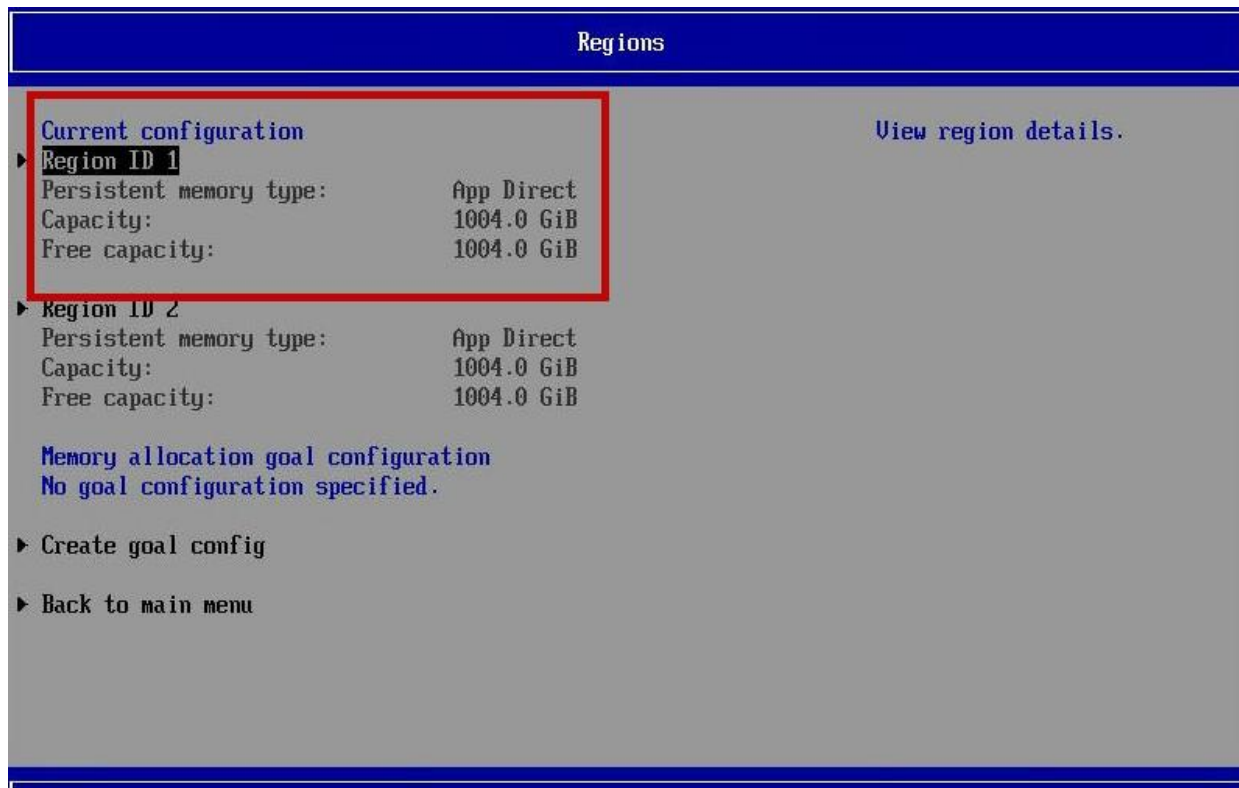
Figure 30. Memory mode vs. reserved mode

Another option is determining if the modules will be standard App Direct which is interleaved, or if the modules will be non-interleaved. This can be selected under the **Persistent Memory Type** option. Once the selections are made, the final step is to select **Create goal config** option, which will create the goal configuration and show on the goal config menu as a summary referenced in Figure 31.

**Notice:** For the changes to take effect, BIOS options must be saved <F10> and the system reset.

Figure 31. App-Direct Mode Config Setup

After the system reboot, go back into BIOS Setup, and navigate to the *Regions* sub menu. There the screen will show a summary of the current configuration (see Figure 32).



**Figure 32. Region Configuration after reboot**

Selecting a *Region ID* will provide details for that region (see Figure 33) including: region ID, modules listed as part of the Region ID, total capacity, persistency type, Health and CPU Socket related to the region.



**Figure 33. Region ID Summary**

### 2.6.3 Namespace Sub-menu

From the *Intel® Optane™ DC Persistent Memory Configuration* main setup page, select the *Namespace* sub-menu option (Figure 16). This option is used to assign a name or label to a specified region and assign a size to the namespace (see Figure 34). After setting your options, select **Create Namespace**.

| Create Namespace  |                   |
|---|-------------------|
| Create a namespace.   | Name of namespace |
| Name  |                   |
| Region ID   | <0x0001>          |
| Mode  | <None>            |
| Capacity input  | <Remaining>       |
| Units   | <GiB>             |
| Capacity  | 1004.0            |
| ▶ Create namespace<br>▶ Back to Namespaces<br>▶ Back to main menu |                   |

Figure 34. Creating Namespace

### 2.6.4 Total Capacity Sub-menu

From the *Intel® Optane™ DC Persistent Memory Configuration* main setup page, select the *Total Capacity* sub-menu (see Figure 16). In the *Total Capacity* menu, the capacity of the Intel® Optane™ PMem modules can be viewed, along with its configuration in Memory capacity, App Direct capacity, Reserved capacity and unconfigured capacity. The Inaccessible capacity value shown is a combination of rounding of allocation blocks, where the free space could not be allocated because it was smaller than the allocation size (1GB for example) and meta-data which is stored on the modules. (see Figure 35)

| Total Capacity  |                          |
|---|--------------------------|
| Total DCPMM resource allocation across the host server. | Return to the main menu. |
| Raw capacity:   | 1.9 TiB                  |
| App Direct capacity:                                    | 1.9 TiB                  |
| Memory capacity:  | 0 B                      |
| Unconfigured capacity:                                  | 0 B                      |
| Inaccessible capacity:                                  | 2.3 GiB                  |
| Reserved capacity:                                      | 0 B                      |
| ▶ Back to main menu                                     |                          |

Figure 35. Total Capacities

## 2.6.5 Diagnostics Sub-menu

From the *Intel® Optane™ DC Persistent Memory Configuration* main setup page, select the **Diagnostics** sub-menu (Figure 16), to select and run diagnostics on individual PMem modules, or all modules, as well as validating the configuration parameters, and checking the firmware and security settings. (see Figure 36).

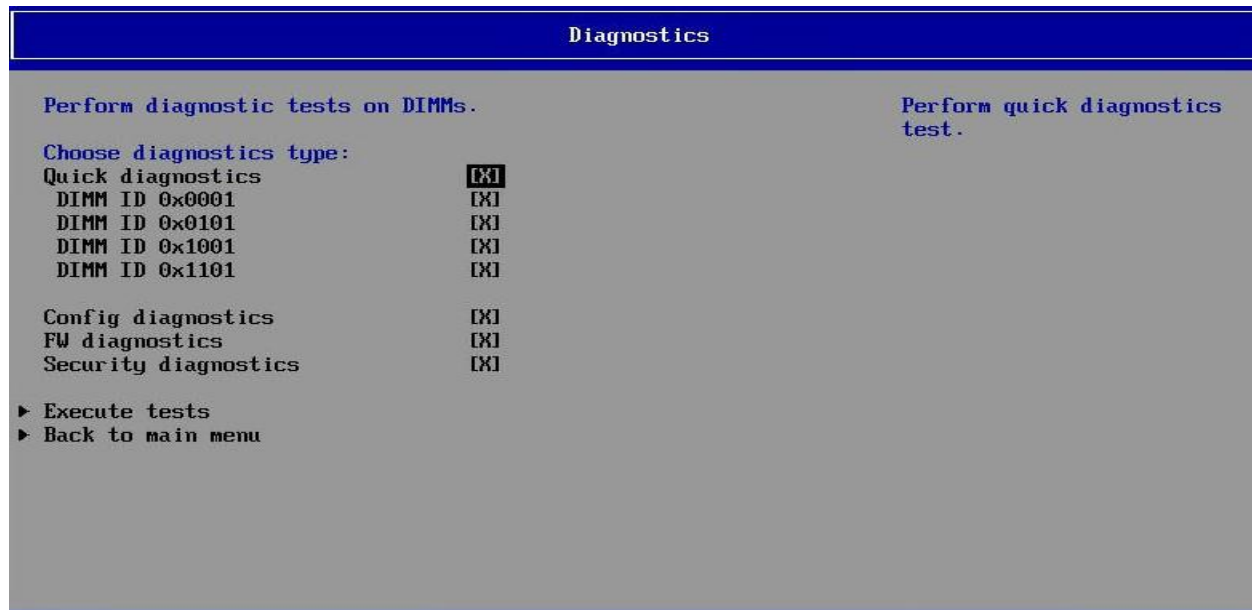


Figure 36. Diagnostics submenu

## 2.6.6 Preferences Sub-menu

From the *Intel® Optane™ DC Persistent Memory Configuration* main setup page, select the *Preferences* submenu (see Figure 16). The Preferences submenu allows the user to view or change the Default DIMM ID (by handle or UID), Capacity Units, and App Direct granularity (see Figure 37). App Direct granularity controls the size of the allocation blocks in either 32GB or 1GB chunks. The default value of 32GB is the recommended value.

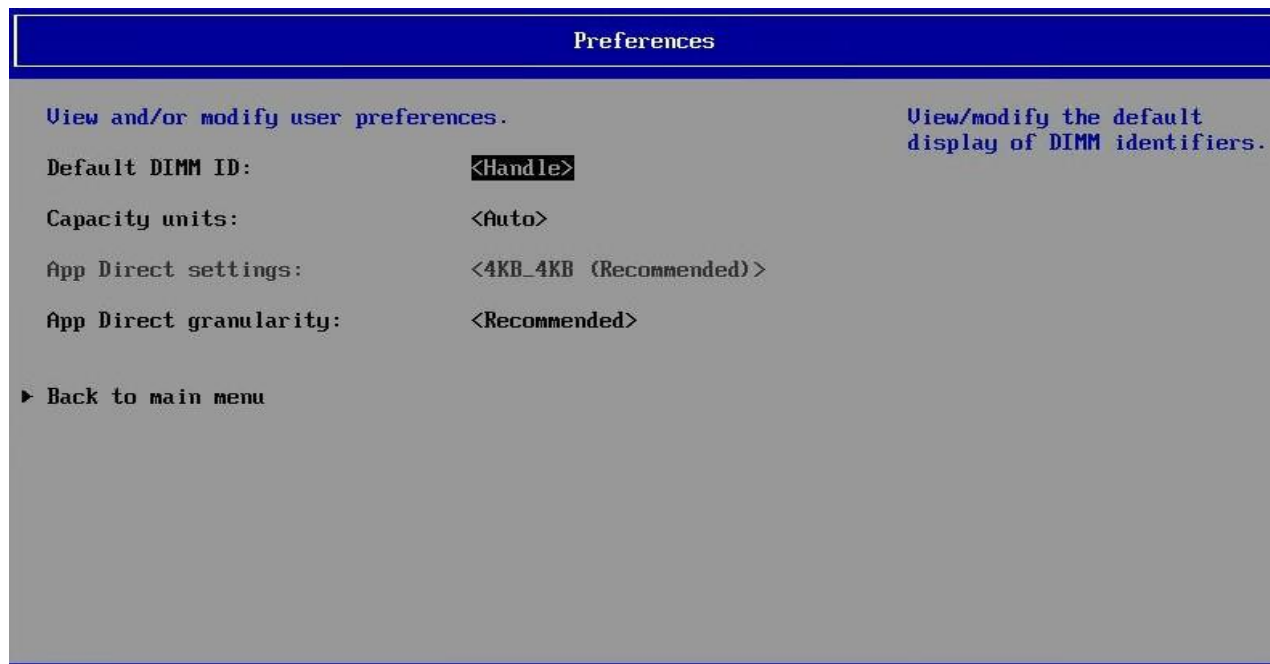


Figure 37. Preferences



## 3. FAQ

---

The following sections cover common questions and answers associated with Intel® Optane™ Persistent Memory and related components.

This list will evolve over time with additional Q&A periodically added. A more comprehensive FAQ can be found in the Intel® Optane™ Datacenter Persistent Memory Module Operations Guide. See section 1.1

### 3.1 Intel® Optane™ PMem Startup FAQ

This Q & A section is specific to new Intel® Optane™ PMem devices received from the factory. Common topics are firmware versions, default configuration, default settings, and visual traits such as labels, color, etc.

#### **Q1 - Why do I have to populate Intel® Optane™ PMem devices in specific DIMM slots?**

There are multiple reasons for this requirement.

Server platforms with two memory slots per channel require the memory slot farthest from the processor be populated before the near slot on the same channel to provide electrical signal termination of both ends of the bus, maintaining signal integrity. Installing a second DRAM on the same channel in the near slot is allowed as long it has equal or lesser electrical load than the DRAM in the far slot. The memory controller in the processor terminates one end of the DDR bus while the DRAM terminates the far end.

Electrically, Intel® Optane™ PMem devices appear as an LRDIMM to the platform and has reduced electrical load characteristics compared with an RDIMM. To maintain signal integrity on the DDR bus, the highest load device must occupy the farthest slot on the channel, while Intel® Optane™ PMem devices occupy the near slot.

UEFI FW memory initialization routines have rules to check for DIMM population and to simplify some operations. On systems where two DRAM and DCPMM are discovered, the near slot is expected to contain the Intel® Optane™ PMem device.

When a channel is populated by only a Intel® Optane™ PMem device, it must occupy the farthest slot on the channel to maintain electrical signal integrity.

When LRDIMM's are installed in place of RDIMM's, due to the UEFI FW requirement, they must be populated in the far slot on a two slot per channel configuration, while the Intel® Optane™ PMem device occupies the near slot.

UEFI FW also has rules regarding which processor socket and memory controller is required to be populated first, again to simplify the rule set. On systems with more than one processor, the first processor (boot service processor), must be populated before other processors.

The first processor memory controller must be populated before the second memory controller.

The expectation is that all memory controllers and sockets are populated symmetrically to reduce the number of permutations that must be validated.

With that concept, the populate notation has been abbreviated to terms like 2-2-2, 2-2-1, 2-1-1, and 1-1-1, which represents the number of DIMM's populated per channel. Generally, the quantity of DRAM devices is greater than or equal to the number of Intel® Optane™ PMem devices in this notation.

## **Q2 - What is the FW version on a production Intel® Optane™ PMem device?**

[See answer in Firmware FAQ below.](#)

## **Q3 - When Intel updates the Intel® Optane™ PMem firmware in its factory, how do you recognize parts with the new firmware?**

Intel's plan is to manufacture Intel® Optane™ PMem devices with the original Production Version (PV) firmware installed. There is no plan to change firmware version in our manufacturing line.

[See also answer in Firmware FAQ below.](#)

## **Q4 - What is the default power limit on new Intel® Optane™ PMem devices?**

The default is 15 watts. If the chassis supports a higher power setting the OEM may enable a capability to change that setting to a higher value.

## **3.2 Intel® Optane™ PMem 100 Series Firmware FAQ**

The Firmware FAQ covers common questions about Intel® Optane™ PMem firmware functionality and maintenance.

### **Q1 - What is the FW version on production Intel® Optane™ PMem devices?**

The firmware version on production Intel® Optane™ PMem 100 series devices is version 1.2.0.5355, which is the first production version release.

### **Q2 - When Intel updates the firmware on new Intel® Optane™ PMem devices, how does someone identify them from the product label?**

The label does not indicate firmware version, however there is a factory date code embedded in the serial number that identifies the year and week the device was manufactured.

Intel's plan is to manufacture Intel® Optane™ PMem devices with the Production Version (PV) firmware installed. There is no plan to change firmware version in our manufacturing line.

Should the need arise to change the factory firmware version, the change will be announced in the Intel® Optane™ PMem Message of the Week, as well as common external channels. The best indication would be the factory date code, which can be correlated with the announcement date for the change.

### **Q3 - Can Intel® Optane™ PMem devices with different FW versions work together?**

Yes, they should be able to work together, however Intel strongly recommends upgrading them all to the latest available firmware version from the hardware provider.

### **Q4 - Can Intel® Optane™ PMem Firmware be downgraded?**

The short answer is yes, and there are both options and limitations.

Intel® Optane™ PMem firmware version numbers have a field for called Security Version Number (SVN) that gets incremented each time a security feature is enabled or revised. The behavior is that once you install firmware with a specific SVN, you cannot downgrade to a lower version to expose a possible security vulnerability that may have been present in the previous firmware version.

Intel provides an SVN bypass 'Opt In' mechanism through UEFI firmware that an OEM can choose to enable downgrade of firmware past the SVN field to enable downgrading any firmware version. The opt in election must be performed on the first boot of a new Intel® Optane™ PMem device prior to POST complete to enable this feature. Therefore, the UEFI firmware must be coded to support this feature as it is not end user selectable.

The ability to downgrade past the SVN is an OEM platform specific feature.



## Appendix A. Glossary

| Word/Acronym   | Definition                            |
|----------------|---------------------------------------|
| <b>1LM</b>     | 1 Layer Memory                        |
| <b>2LM</b>     | 2 Layer Memory                        |
| <b>AD Mode</b> | App-Direct Mode                       |
| <b>BIOS</b>    | Basic Input Output System             |
| <b>DCPMM</b>   | Data Center Persistent Memory Module  |
| <b>DIMM</b>    | Dual inline memory Module             |
| <b>MM mode</b> | Memory Mode                           |
| <b>PCI</b>     | Peripheral Component interconnect     |
| <b>PMEM</b>    | Persistent Memory                     |
| <b>UEFI</b>    | Unified Extensible Firmware Interface |

## Appendix B. Getting Help

To obtain support for an issue with the server system, follow these steps:

1. Visit the following Intel® support web page: <http://www.intel.com/support/>

This web page provides 24x7 support for the latest and most complete technical support information on all Intel Enterprise Server and Storage Platforms. Information available at the support site includes:

- Latest BIOS, firmware, drivers and utilities
  - Product documentation, setup, and service guides
  - Full product specifications, technical advisories and errata
  - Compatibility documentation for memory, hardware add-in cards, and operating systems
  - Server and chassis accessory parts list for ordering upgrades or spare parts
  - A searchable knowledgebase to search for product information throughout the support site
2. If a solution cannot be found at Intel's support site, send an email to Intel's support center using the online form available at: [http://www.intel.com/p/en\\_US/support/contactsupport](http://www.intel.com/p/en_US/support/contactsupport).
  3. Lastly, contact an Intel support representative using one of the support phone numbers available at: <http://www.intel.com/support/feedback.htm?group=server> (charges may apply).

Intel also offers Channel Program members around-the-clock, 24x7, technical phone support on Intel server boards, server chassis, server RAID controller cards, and Intel Server Management at: <http://www.intel.com/reseller/>.

---

**Note:** Access to the 24x7 number requires a login to the reseller site.

---

### Warranty Information

To obtain warranty information, visit [http://www.intel.com/p/en\\_US/support/warranty](http://www.intel.com/p/en_US/support/warranty).