

Dell EMC DCPMM

User's Guide

Notes, cautions, and warnings

 **NOTE:** A NOTE indicates important information that helps you make better use of your product.

 **CAUTION:** A CAUTION indicates either potential damage to hardware or loss of data and tells you how to avoid the problem.

 **WARNING:** A WARNING indicates a potential for property damage, personal injury, or death.

© 2019 Dell Inc. or its subsidiaries. All rights reserved. Dell, EMC, and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners.

1 Introduction.....	5
System requirements.....	5
Terminology.....	6
2 Change list.....	7
3 Hardware.....	8
Server hardware configuration.....	8
DIMM installation and removal.....	8
DCPMM hardware configuration.....	8
DCPMM Recommended Topologies.....	8
CPU type and maximum memory limits.....	10
DCPMM mixing and population rules.....	10
4 BIOS.....	12
BIOS configuration setting for Intel DCPMM.....	12
DIMM discovery.....	12
App-direct mode configuration.....	13
Create goal.....	13
Region information.....	14
Memory mode configuration.....	16
Create goal.....	16
5 DCPMM event reporting.....	17
Events during runtime.....	17
Events during boot time.....	17
6 iDRAC Intel DCPMM management.....	20
iDRAC GUI.....	20
DCPMM firmware version.....	20
DCPMM hardware status.....	20
DCPMM goal configuration using iDRAC GUI.....	21
DCPMM remaining rated write endurance.....	22
7 DCPMM security	23
Memory mode.....	23
App-direct.....	23
Cryptographic erase and DCPMM sanitize.....	24
Crypto erase.....	25
Sanitize.....	25
8 DIMM configuration changes.....	27

9 Windows.....	28
PMEM in App-direct mode.....	28
PMEM Disk management.....	29
List PMEM physical disks and check their health status.....	30
Create PMEM Disks.....	30
Remove PMEM disks.....	31
PMEM disk with interleave sets.....	31
PMEM disk creation with interleave sets.....	31
PMEM in memory Mode	32
Windows troubleshooting and event monitoring.....	32
10 Linux.....	34
Identify and configure persistent memory device.....	34
Listing DCPMM devices.....	34
Create namespace.....	34
Mount file system on namespace device.....	34
Delete namespaces.....	35
Management utility.....	35
Check DCPMM health status.....	35
Linux errata.....	35
11 VMware ESXi.....	36
PMEM in App-direct mode.....	36
PMEM in Memory mode.....	37
PMEM health status.....	37
ESXi troubleshooting and maintenance.....	37
12 System diagnostics.....	39
13 Firmware update.....	40
Dell DUP update.....	40

Introduction

Dell EMC now offers Intel's Data Center Persistent Memory Modules (DCPMMs), a nonvolatile memory that has the same form factor as a standard DDR4 DIMM.

DCPMMs come in three capacities: 128 GB, 256 GB, and 512 GB.

DCPMMs are installed in the memory slots of the server and are compatible with RDIMMs and LRDIMMs.

DCPMM can be configured in two modes:

- Memory mode
- App-direct mode

In Memory mode: DCPMMs act as volatile system memory and any RDIMMs or LRDIMMs installed operate as cache for the DCPMMs.

In App-direct mode: DCPMMs operate as byte addressable memory mapped persistent memory. DCPMMs and DRAM act as independent memory resources. RDIMMs or LRDIMMs acts as volatile system memory. Applications can access persistent memory by using memory load or store commands. Legacy applications that access storage as block devices can access persistent memory (PMEM) through the PMEM block driver.

DCPMMs do not require an extra power supply or battery and are inherently persistent.

This document is intended to help customers understand the basics of Intel's DCPMM technology integrated in Dell PowerEdge systems. It covers the fundamentals of installation and configuration of key settings for both the operating modes.

Topics:

- [System requirements](#)
- [Terminology](#)

System requirements

Table 1. System Requirements

Component	Minimum version required
System	R640, R740, R740xd, R840, R940, R940xa, MX740c and MX840c (2 and 4 socket systems)
Processor	2 nd Generation Intel Xeon Platinum or Gold processors
BIOS	R640 - 2.3.10 R740/R740xd/R940 - 2.2.10 R840 / R940xa - 2.3.10 MX740c / MX840c - 2.3.10
CPLD	R640 - 1.0.6 R740/R740xd - 1.0.8 R840 - 1.0.6 R940 - 1.0.4 R940xa - 1.0.6 MX740c / MX840c - 1.0.6
iDRAC	3.34.34.34
DCPMM FW	Build # 5375

Component	Minimum version required
OS	Microsoft Windows 2019 VMware ESXi 6.7 with EP10 (Build #13981272) Red Hat Enterprise Linux 7.6 SUSE Linux Enterprise Server 15

Terminology

Table 2. Terminology

Terminology	Description
App-direct mode (AD)	Persistent memory is accessed directly by applications as byte-addressable memory.
CLI	Intel command-line interface in UEFI or Linux environment
DCPMM	Intel Data Center Persistent Memory Module
GiB	Gibibyte 1 GiB = 1024 MiB
GB	Gigabyte 1 GB = 1000 MB
iMC	Integrated Memory Controller
Interleave Set	Contiguous App-direct mode capacity that is created by interleaving the persistent capacity of one or more DCPMMs.
x1 (by one) Interleave	Interleave set that includes App-direct capacity from only one DCPMM pronounced "by one." This is essentially "Noninterleaved"
PM, PMEM	Persistent Memory
Memory Mode (MM)	RDIMM or LRDIMM becomes cache for nonpersistent DCPMMs.
MiB	Mibibyte 1 MiB = 1024 KB = 1048576 bytes
MB	Megabyte 1 MB = 1000 KB = 1000000 bytes
Namespace	A persistent memory device that is made available in the file system.
TiB	Tebibyte 1 TiB = 1024 GiB
TB	Terabyte 1 TB = 1000 GB

Change list

Table 3. Change list

Version	Changes
A01	Original Version
A02	Topics with new updates: <ul style="list-style-type: none">• System Requirement• Server hardware configuration• DCPMM recommended topologies• Management utility
A03	Topics with new updates: <ul style="list-style-type: none">• System Requirement• Server hardware configuration

Server hardware configuration

DCPMM is supported in R640, R740/R740XD, R840, R940, R940xa, MX740c and MX840c PowerEdge servers with second-generation Intel Xeon Scalable Processors of Gold and Platinum grade. Refer DCPMM configurations for a list of fully supported and validated DCPMM configurations in two-socket servers. Four-socket configurations are a direct scale up of two-socket configurations.

NOTE: The maximum number of DCPMMs for R640, R740/R740XD, MX740c is 12 and for a R840, R940/R940xa, MX840c is 24. For PowerEdge R940/R940xa, 2400 W PSUs are required when Intel DCPMMs are present on the system. When required PSUs are not installed, users see messages in the Lifecycle Controller logs about CPLD Arm failure errors and warnings about invalid PSU configurations. These messages occur during the boot time. PWR message occurs only once when AC is powered on and will not re-occur.

DIMM installation and removal

Industry standard DIMM practices and procedures must be followed while handling, installing, or removing DCPMM memories.

For more details on procedures for installing/removing modules, see the "Standard Practices and Procedures - Module Insertion Procedure for DIMM and miniDIMM Connectors" documentation by JEDEC.

JEDEC Standards (www.jedec.org): Document Number SPP-023B.

DCPMM hardware configuration

DCPMM Recommended Topologies

This section contains the high-level introduction on DCPMM configuration and provisioning concepts.

The following topologies are recommended per CPU socket. For multiple socket systems, each socket should be populated identically.

See the Installation and Service Manual for Memory installation guidelines of respective servers.

NOTE: The following image and table is for reference which displays the R740/R740XD CPU and DIMM slot locations.




Figure 1. Memory Layout for R740/R740XD

Table 4. DCPMM configurations

DCPMMs Populated	DCPMM Capacity	LR/RDIMM Capacity	Population Locations	Memory Mode Supported?	DDR:DCPMM Ratio in Memory Mode
6x	128 GB	32 GB	RDIMM: Slot 1-6 DCPMM: Slot 7-12	Y	1:4
6x	128 GB	64 GB	RDIMM: Slot 1-6 DCPMM: Slot 7-12	N	N/A
6x	128 GB	128 GB	RDIMM: Slot 1-6 DCPMM: Slot 7-12	N	N/A
6x	256 GB	32 GB	RDIMM: Slot 1-6 DCPMM: Slot 7-12	Y	1:8
6x	256 GB	64 GB	RDIMM: Slot 1-6 DCPMM: Slot 7-12	Y	1:4
6x	256 GB	128 GB	RDIMM: Slot 1-6 DCPMM: Slot 7-12	N	N/A
6x	512 GB	32 GB	RDIMM: Slot 1-6 DCPMM: Slot 7-12	N	N/A
6x	512 GB	64 GB	RDIMM: Slot 1-6 DCPMM: Slot 7-12	Y	1:8
6x	512 GB	128 GB	RDIMM: Slot 1-6 DCPMM: Slot 7-12	Y	1:4
4x	512 GB	32 GB	RDIMM: Slot 1-6	N	N/A

DCPMMs Populated	DCPMM Capacity	LR/RDIMM Capacity	Population Locations	Memory Mode Supported?	DDR:DCPMM Ratio in Memory Mode
			DCPMM: Slot 7,8,10,11		
4x	512 GB	64 GB	RDIMM: Slot 1-6 DCPMM: Slot 7,8,10,11	N	N/A
4x	256 GB	32 GB	RDIMM: Slot 1-6 DCPMM: Slot 7,8,10,11	N	N/A
4x	256 GB	64 GB	RDIMM: Slot 1-6 DCPMM: Slot 7,8,10,11	N	N/A
1x	128 GB	32 GB	RDIMM: Slot 1-6 DCPMM: Slot 7	N	N/A
1x	128 GB	64 GB	RDIMM: Slot 1-6 DCPMM: Slot 7	N	N/A

i NOTE: Servers with one DCPMM can be configured on a single socket platform only.

CPU type and maximum memory limits

Table 5. CPU type and maximum memory limits

CPU type	Maximum memory supported (Includes volatile and persistent memory capacity)
All CPU SKUs	1 TB per CPU socket
M SKUs	2 TB per CPU socket
L SKUs	4.5 TB per CPU socket

DCPMM mixing and population rules

This section has general rules for DIMM mixing and population.

Each system must contain only one capacity of DCPMM. If you mix DCPMM capacities, an F1/F2 warning message is displayed. This is not a supported configuration and must not be populated. The table DCPMM configurations replace the following rules:

Mixing rules

- DCPMM can be mixed with RDIMM, LRDIMM, and 3DS LRDIMM.
- Mixing DDR4 DIMM types (RDIMM, LRDIMM, 3DS LRDIMM), within a channel, iMC, socket, or across sockets are not supported.
- x4 and x8 DDR4 DIMMs can be mixed within a channel.
- Mixing DCPMM operating modes (App-direct, Memory mode) is not supported.

Population rules

- Maximum of one DCPMM per channel.
- If only one DIMM is populated on a channel, it should always go to the first slot in that channel (white slot).
- If a DCPMM and a DDR4 DIMM are populated on the same channel, always plug DCPMM on the second slot (black slot).

- If the DCPMM is configured in Memory Mode, the recommended DDR4 to DCPMM capacity ratio is 1:4 to 1:16 per iMC.

BIOS configuration setting for Intel DCPMM

DIMM discovery

All installed DCPMMs that the BIOS has discovered during system inventory is displayed in the BIOS Intel Persistent Memory tab:

Memory Settings > Persistent Memory > Intel Persistent Memory > Persistent Memory DIMM Configuration.

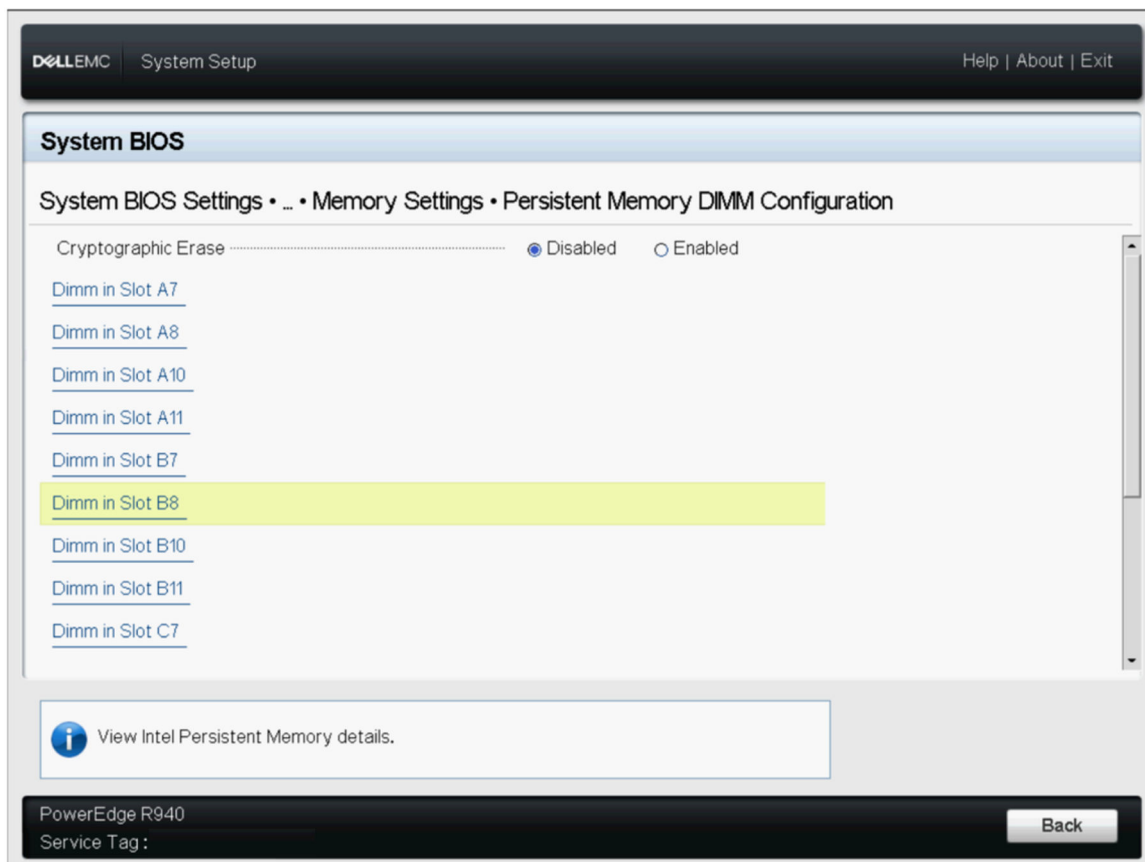


Figure 2. Persistent Memory screen

NOTE: DCPMMs are shown as DIMMs.

There is one entry for each DCPMM that is installed and the current health and status information for each DCPMM is displayed as:

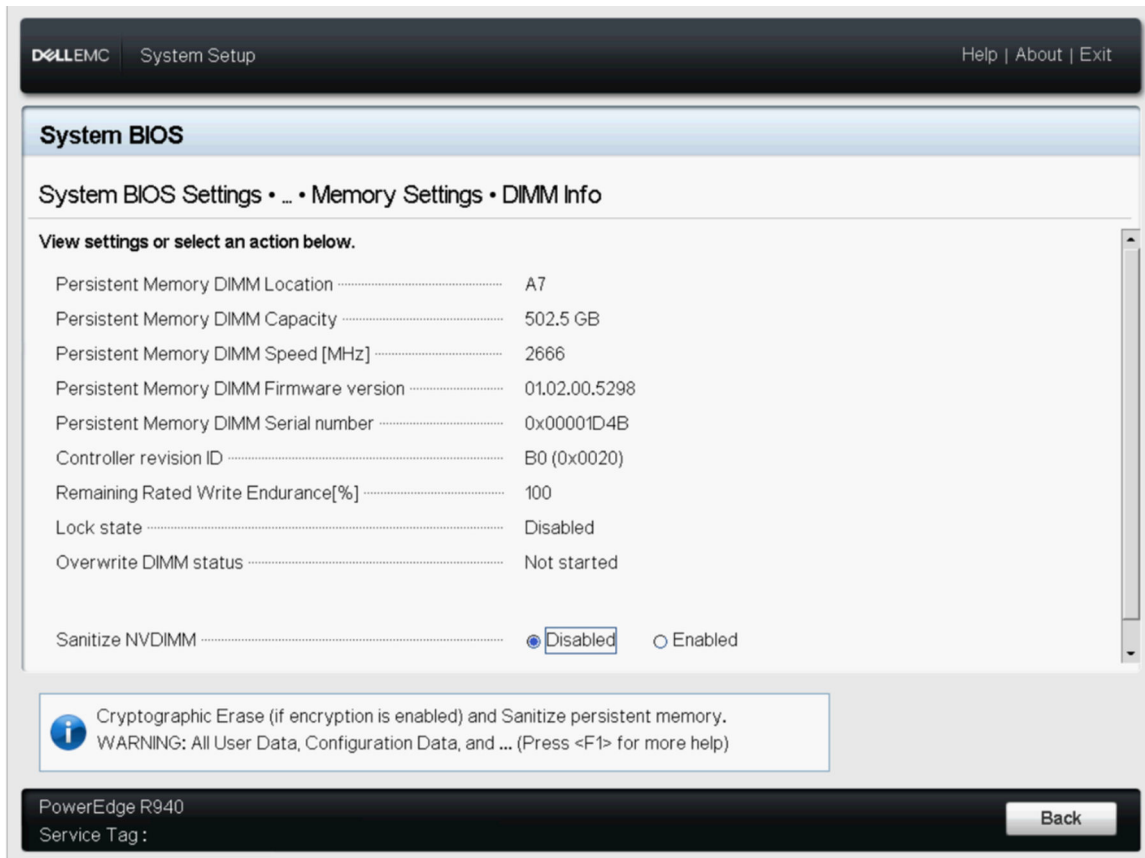


Figure 3. Memory info

NOTE: Data is always assumed to be in units of MiB/GiB/TiB even if labeled MB/GB/TB. User capacity overhead is up to 2% of capacity (GiB). Another overhead may be required for Regions, Namespace, and Filesystems.

App-direct mode configuration

Create goal

Goal is created in BIOS.

To create a goal in BIOS, go to: **Memory Settings > Persistent Memory > Intel Persistent Memory > Region Configuration > Create Goal Config.**




Figure 4. Goal configuration

The BIOS options determine how the goal is created and the DCPMMs are configured:

Persistent [%]:

- No Change - Does not apply any changes to the current goal.
- 100 - Creates a goal of 100% Persistent memory across the selected DCPMMs.
- 0 - Creates a goal of 0% Persistent memory across the selected DCPMMs. This operation configures all the DCPMM as Memory mode.

Persistent memory type:

- App-direct Interleaved - Persistent mode interleave across the DCPMMs in a socket. The DCPMMs are displayed one PMEM device per socket in the operating system.
- App-direct Not Interleaved - Persistent mode is applied to DCPMM individually. Each DCPMM is displayed as an individual PMEM device in the operating system.

After the goal is configured, and exited from BIOS, the goal will be created across the DCPMMs using the user specified settings during the next boot.

Region information

Information about each region that is created during the **Create Goal Config** process can be accessed in the **Region Configuration** tab in BIOS after a system reset:

Memory Settings > Persistent Memory > Intel Persistent Memory > Region Configuration.




Figure 5. Region configuration

The number of regions that are displayed depends on the number of processors in the system and not on the DCPMMs interleaved. If the DCPMMs are configured as interleaved, one Persistent Memory Region is listed per socket in the system that has DCPMMs installed. If the DCPMMs are configured as non-interleaved, one Persistent Memory Region is listed per DCPMM installed in the system.

Region information can be accessed by clicking each Persistent Memory Region link in the BIOS. Following is an example:

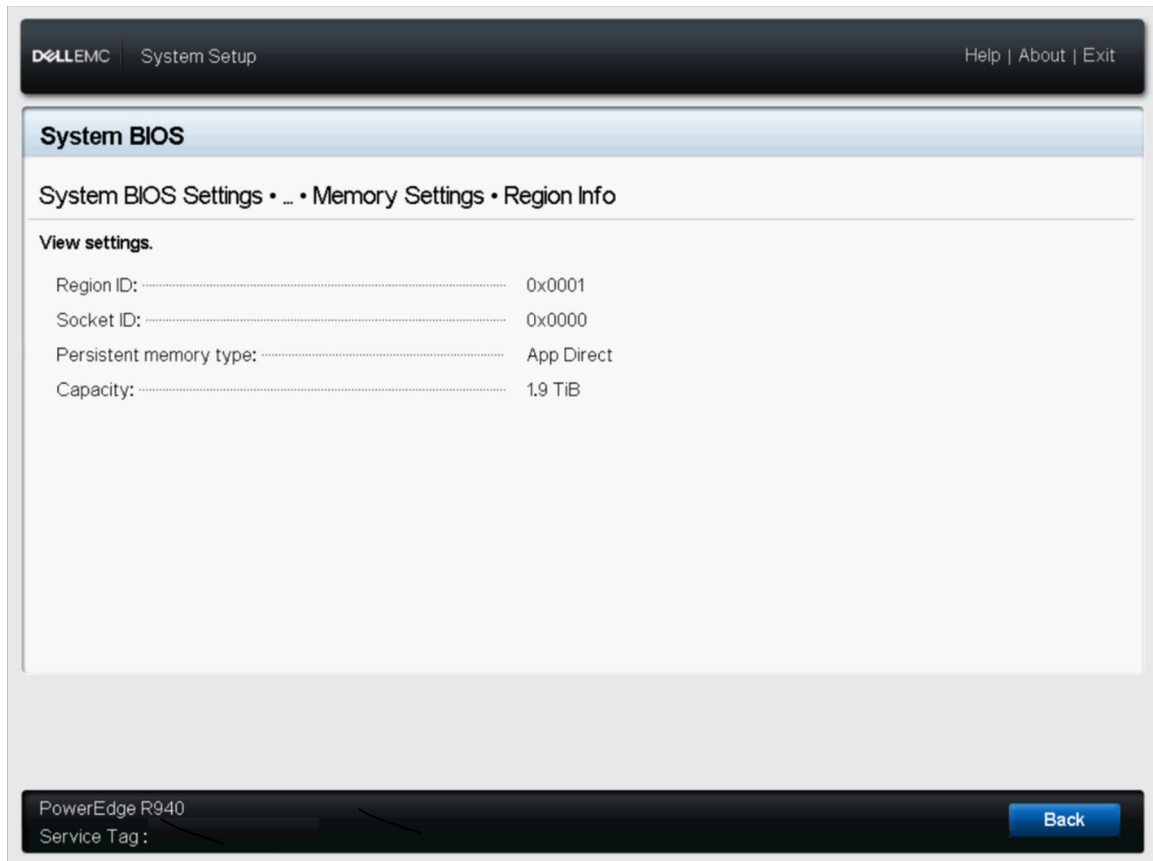


Figure 6. Region info

Memory mode configuration

Create goal

Goal is created in the BIOS.

To create a goal in BIOS, go to: **Memory Settings > Persistent Memory > Intel Persistent Memory > Region Configuration > Create Goal Config.**

The BIOS options determine how the goal is created and the DCPMMs are configured:

Operation Target:

- Platform - Applies the goal to all the DIMMs in the system (recommended).

Persistent [%]:

- No Change - Does not apply any changes to the current goal.
- 100 - Creates a goal of 100% Persistent memory across the selected DCPMMs.
- 0 - Creates a goal of 0% Persistent memory across the selected DCPMMs. This operation configures all the DCPMM as Memory mode.

DCPMM event reporting

When system detects DCPMM-related event either during runtime or POST, system will log the events in Server System Event Log(SEL) and Life-Cycle Log(LCL). If an event is detected during boot time, system halts during the POST and user needs to press F1 to continue the boot process.

NOTE: NVDIMM is frequently used in these messages. The term NVDIMM is generic to several different families of persistent memory including DCPMM and not meant to indicate NVDIMM-N modules.

Topics:

- [Events during runtime](#)
- [Events during boot time](#)

Events during runtime

- **MEM0001** : Multibit memory errors detected on a memory device at location <location>.

Recommended Action : Reinstall the memory component. If the problem persists, contact technical support.
- **MEM0701** : Correctable memory error rate exceeded for <location>.

Recommended Action : Reinstall the memory component. If the problem continues, contact support.
- **MEM9022** : A Non-Critical event was detected on the Non-Volatile Dual In-line memory module (NVDIMM) device in the slot <location>.

Recommended Action : If the issue persists, contact the service provider.
- **MEM9040** : A Critical event was detected on the Non-Volatile Dual In-line memory module (NVDIMM) device in the slot <location>.

Recommended Action : Remove and reinstall the Non-Volatile Dual In-line memory module (NVDIMM) device. If the issue persists, contact the service provider. For information about removing and reinstalling the NVDIMM, see the system owner's manual on the support site.
- **MEM9061**: An Information-Only event was detected on the Non-Volatile Dual In-line memory module (NVDIMM) device in the slot <location>. The NVDIMM is operating normally.

Recommended Action: No response action is required.
- **MEM9073** : Unable to update firmware for the NVDIMM identified in the message.

Recommended Action: Retry the operation. If the issue persists, replace the NVDIMM, or contact your service provider. It is also recommended that the system is upgraded to the latest BIOS.
- **MEM9072** : The system memory has faced an uncorrectable multibit memory errors in the nonexecution path of a memory device at the location <location>.

Recommended Action: Reinstall the memory device. If the issue persists, contact your service provider. To select a convenient contact method, see the product documentation available on the support site.

Events during boot time

- **UEFI0337** : Unable to update the firmware of NVDIMM located in the memory slot <slot number>.

Recommended Action : Disconnect the input power to the system, wait for 30 seconds, and reconnect the power, Power on the server and retry the operation. If the issue persists, replace the NVDIMM. For more information, see the product's owner's manual on the support site.
- **UEFI0338** : The firmware of NVDIMM located in the memory slot <slot number> is successfully updated.

Recommended Action : N/A
- **UEFI0345** : The erase operation on the nonvolatile DIMM with serial number <serial number> in slot <slot number> is successfully completed.

Recommended Action: N/A

- **UEFI0347** : Unable to initialize the memory because one or more errors have occurred during the NVDIMM initialization in the slot <slot Number>.

Recommended Action : Manually remove and reinstall the NVDIMM. If the issue persists, contact your service provider. For more information about removing and reinstalling an NVDIMM, see the product Installation and Service Manual available on the support site.
 - **UEFI0348** : The Remaining Rated Endurance value of NVDIMM installed in slot <slot Number> is less than or equal to 1%.

Recommended Action : Power off the server and replace the NVDIMM immediately. For more information about the Remaining Rated Endurance, see the product Installation and Service Manual available on the support site.
 - **UEFI0349** : Unable to initialize the NVDIMMs because different types of NVDIMM SKUs or controller revisions are installed.

Recommended Action : Power off the server and replace the NVDIMMs to ensure all NVDIMM SKUs or controller revisions are same. For more information about the NVDIMMs SKUs or controller revision, see the product Installation and Service Manual available on the support site.
 - **UEFI0350** : The NVDIMM in the memory slot <slot Number> is replaced or removed from a previously configured Persistent Memory (PM) region.

Recommended Action : Create a new Persistent Memory (PM) region. For more information about creating a PM region, see the platform Installation and Service Manual available on the support site.
 - **UEFI0351**: The NVDIMMs on socket <socket Number> are installed in an unsupported (configuration) manner. Major Error code <major MRC error code> Minor Error code <minor MRC error code>. Data in the Persistent Memory (PM) region may not be accessible.

Recommended Action : For more information about creating a PM region, see the platform Installation and Service Manual available on the support site.
 - **UEFI0352** : Unable to use the Persistent Memory (PM) region configuration of the NVDIMM in the memory slot <slot Number>.

Recommended Action : Verify the NVDIMM population configuration and retry the operation. For more information about NVDIMMs, see the platform Installation and Service Manual available on the support site.
 - **UEFI0353** : The Persistent Memory (PM) configuration information of the NVDIMM in the memory slot <slot Number> is corrupted because the checksum or the header type is not valid.

Recommended Action : Remove and reinstall the NVDIMMs or create a new Persistent Memory (PM) configuration. If the issue persists, contact your service provider. For more information about creating a PM region, see the platform Installation and Service Manual available on the support site.
 - **UEFI0354** : The DDR4 memory to NVDIMM ratio on the socket <socket Number> is not optimal to provide minimal performance requirement.


Recommended Action : Reconfigure DDR4 memory to a ratio between 1:4 and 1:16. For more information about reconfiguring DDR4 memory, see the platform Installation and Service Manual available on the support site.
 - **UEFI0355** : All the NVDIMMs are disabled because the value of memory capacity of all the installed NVDIMMs has exceeded the maximum value supported by the processor.

Recommended Action : Reconfigure the processor to support the value of memory capacity of all the installed NVDIMMs. For more information about reconfiguring the processor, see the platform Installation and Service Manual available on the support site.
 - **UEFI0356** : The data in the Persistent Memory DIMM located in memory slot <slot label> is not accessible because the DIMM is locked and the passphrase is incorrect.

Recommended Action : Update the Persistent Memory Passphrase to the correct passphrase or perform a Secure Erase operation on the Dual inline memory module (DIMM). Secure erase, erases all persistent data.
 - **UEFI0357** : The Cryptographic Erase operation on the Intel Persistent Memory DIMM with serial number <serial number> in slot <slot number> is successfully completed.

Recommended Action : N/A
 - **UEFI0358** : Unable to complete the Cryptographic Erase operation on the Intel Persistent Memory DIMM with serial number <serial number> in slot <slot number>.

Recommended Action : Retry the operation. If the issue persists, contact your service provider.
 - **UEFI0359** : The Overwrite DIMM operation on the Intel Persistent Memory DIMM with serial number <serial number> in slot <slot number> is successfully completed.

Recommended Action : N/A
-  **NOTE: This is part of the DCPMM Sanitize function.**
- **UEFI0360** : Unable to complete the Overwrite DIMM operation on the Intel Persistent Memory DIMM with serial number <serial number> in slot <slot number>.

Recommended Action : Retry the operation. If the issue persists, contact your service provider.

 **NOTE: This is part of the DCPMM Sanitize function.**

- **UEFI0361** : The Factory Default operation for the Intel Persistent Memory DIMMs in the system has completed successfully.
Recommended Action : N/A
- **UEFI0362** : Unable to complete the Factory Default operation on the Intel Persistent Memory DIMMs.
Recommended Action : Retry the operation. If the issue persists, contact your service provider.
- **UEFI0367** : The Create Goal operation on the Intel Persistent Memory DIMMs is successfully completed.
Recommended Action : N/A
- **UEFI0368** : Unable to complete the Create Goal operation on the Intel Persistent Memory DIMMs.
Recommended Action : Retry the operation. If the issue persists, contact your service provider.
- **UEFI0369** : The Complex Programmable Logic Device (CPLD) is successfully armed for the Asynchronous DRAM Refresh (ADR) signal.
Recommended Action : N/A
- **UEFI0370** : The Complex Programmable Logic Device (CPLD) cannot arm for the Asynchronous DRAM Refresh (ADR) signal. It might be due to Intel Persistent Memory DIMMs failed to initialize.
Recommended Action : Reboot the system. If the issue persists, contact your service provider.
- **UEFI0372** : The Remaining Rated Endurance value of the NVDIMM installed in the slot <slot number> is equal to 0%.
Recommended Action : Turn off the server and replace the NVDIMM immediately. For more information about the Remaining Rated Endurance, see the product Installation and Service Manual available on the support site.
- **UEFI0373** : The NVDIMM installed in slot <slot number> requires a maintenance.
Recommended Action : Consider replacing the DIMM during the next maintenance cycle. For more information about the NVDIMM Health Status, see the product Installation and Service Manual available on the support site.
- **UEFI0374** : The NVDIMM installed in slot <slot Number> is in critical condition.
Recommended Action : Power off the server and replace the NVDIMM immediately. For more information about the NVDIMM Health Status, see the product Installation and Service Manual available on the support site.
- **UEFI0375** : Unable to apply the Persistent Memory (PM) region configuration of the NVDIMM in the memory slot <slot Number>.
Recommended Action : Verify the NVDIMM population configuration and retry the operation. For more information about NVDIMMs, see the platform Installation and Service Manual available on the support site.
- **UEFI0376** : All the NVDIMMs are disabled because the installed processor does not support the Intel Persistent Memory DIMMs.
Recommended Action : Reconfigure the processor to support the Intel Persistent Memory DIMMs. For more information about reconfiguring the processor, see the Installation and Service Manual of the system available on the support site.
- **UEFI0377** : The Non-Volatile Dual In-Line memory module (NVDIMM) in the memory slot <slotNum> has encountered a transient fatal failure during the previous boot.
Recommended Action : Remove and reinstall the Non-Volatile Dual In-line memory module (NVDIMM) device. If the issue persists, contact your service provider. For more information about installation and removal, see the Installation and Service Manual of the server available on the support site.
- **UEFI0378** : The Non-Volatile Dual In-Line memory module (NVDIMM) in the memory slot <slotNum> has encountered a thermal shutdown event during the previous boot.
Recommended Action : Do one of the following: 1) Power off the server. 2) Disconnect the input power, wait for 30 seconds, and then reconnect to the power source. 3) Power on the server. 4) If the issue persists, contact your service provider.
- **PWR2281** : Unable to perform the memory arming operation because the PSU configuration of the server is insufficient to guarantee data flush time in the event of power loss.
Recommended Action : Do the following and retry the operation:
 - Turn off the server.
 - Ensure the PSUs are installed as recommended in the user guide.
 - Turn on the server.

For more information on supported PSUs, see the Installation and Service Manual of the system available on the support site.

iDRAC Intel DCPMM management

iDRAC GUI

DCPMM firmware version

DCPMM FW version is displayed under **System > Inventory > Firmware inventory tab**.

DDR4 DIMM A7	01.02.00.5365
DDR4 DIMM A8	01.02.00.5365
DDR4 DIMM B12	01.02.00.5365

Figure 7. DCPMM firmware version

DCPMM hardware status

Select the Memory link on the Dashboard to get more information about memory health.

DCPMM hardware status is displayed under **System > Inventory > Hardware inventory tab**.

— DIMM A7

BankLabel: A
CacheSize: 0 MB
CurrentOperatingSpeed: 2666 MHz
DeviceDescription: DIMM A7
DeviceType: Memory
FQDD: DIMM.Socket.A7
InstanceID: DIMM.Socket.A7
LastSystemInventoryTime: 2019-04-03T19:51:17
LastUpdateTime: 2019-03-28T20:45:29
ManufactureDate: Mon Aug 20 07:00:00 2018 UTC
Manufacturer: Intel
MemoryTechnology: Intel persistent
MemoryType: DDR-4
Model: DDR4 DIMM
NonVolatileSize: 129024 MB
PartNumber: NMA1XBD128GQS
PrimaryStatus: Ok
Rank: Single Rank
RemainingRatedWriteEndurance: 100 %
SerialNumber: 0000029B
Size: 131072 MB
Speed: 2666 MHz
SystemEraseCapability: Supported
VolatileSize: 0 MB

Figure 8. DCPMM hardware status

DCPMM goal configuration using iDRAC GUI

1. Log on to iDRAC GUI interface.
2. Navigate to **Configuration > BIOS settings > Memory settings > Persistent Memory Settings tab.**




Figure 9. DCPMM goal configuration using iDRAC GUI

3. Navigate to **Intel Persistent Memory > Region Configuration > Create Goal Config.**

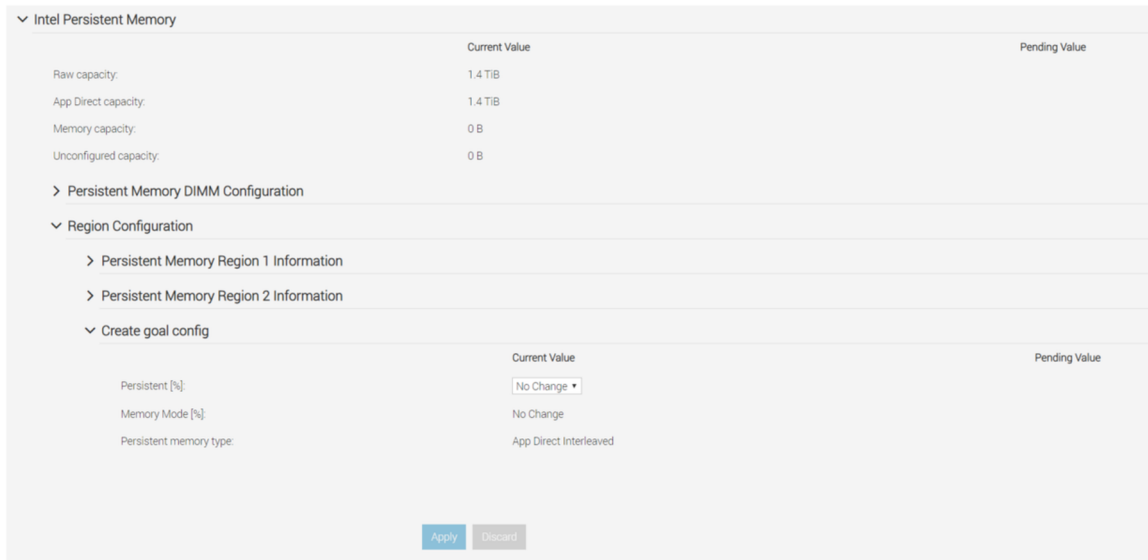


Figure 10. DCPMM goal configuration

- Change Persistent percentage as 100% to configure Intel DCPMM into 100% App-direct Mode and 0% to configure the DIMMs into 100% memory mode.

NOTE: The Persistent memory type field is for configuring a new goal. It is not for reading the status of the current goal. Use instructions under DCPMM Hardware Status to confirm current configuration.

- Click Apply and reset the system.
- Goals will be applied into operating modes (either App-direct or Memory) in the next power cycle.

DCPMM remaining rated write endurance

DCPMM remaining lifetime is reflected as **Remaining Rated Write Endurance** in the iDRAC GUI. It is displayed under **System > Overview > Memory**.

Key limitations and feature information:

- The feature does not work if the system is halted in BIOS such as the BIOS setup. If the system stays in this state for long periods, the endurance for all DCPMM reports as 0% as status cannot be retrieved. To resolve, the user needs to boot into the operating system and wait up to 24 hours for the next automatic poll of status.
- The values do not change frequently and are polled once per day. If the user performs a System Erase or Repurpose and Retire, it destroys the stored values for this feature. DCPMM will be displayed as 0% after such an operation until the next automatic poll within 24 hours.

Individual Memory Details									
Status	Connector Name	Memory Technology	Type	Size	Remaining Rated Write Endurance	State	Rank	Size	Speed
✓	DIMM SLOT A1	DRAM	DDR-4	32 GB	-	Presence Detected	Dual Rank		2666 MHz
✓	DIMM SLOT A2	DRAM	DDR-4	32 GB	-	Presence Detected	Dual Rank		2666 MHz
✓	DIMM SLOT A3	DRAM	DDR-4	32 GB	-	Presence Detected	Dual Rank		2666 MHz
✓	DIMM SLOT A4	DRAM	DDR-4	32 GB	-	Presence Detected	Dual Rank		2666 MHz
✓	DIMM SLOT A5	DRAM	DDR-4	32 GB	-	Presence Detected	Dual Rank		2666 MHz
✓	DIMM SLOT A6	DRAM	DDR-4	32 GB	-	Presence Detected	Dual Rank		2666 MHz
✓	DIMM SLOT A7	Intel persistent memory	DDR-4	512 GB	100%	Presence Detected	Single Rank		2666 MHz
✓	DIMM SLOT A8	Intel persistent memory	DDR-4	512 GB	100%	Presence Detected	Single Rank		2666 MHz
⚠	DIMM SLOT A9	Unknown	Information Not Available	0 GB	-	Absent	Information Not Available		0 MHz
✓	DIMM SLOT A10	Intel persistent memory	DDR-4	512 GB	100%	Presence Detected	Single Rank		2666 MHz
✓	DIMM SLOT A11	Intel persistent memory	DDR-4	512 GB	100%	Presence Detected	Single Rank		2666 MHz

Figure 11. Individual memory details

DCPMM security

Memory mode

In Memory mode DCPMMs operate as volatile system memory. User passphrase is not supported and this BIOS setting will be greyed out.

App-direct

Users have the option to enable Passphrase protection of DCPMM regions. The intent of the passphrase is to protect against unauthorized access to data stored on the DCPMM region. If the DCPMMs are moved from one server to another server, the user must re-enter the security passphrase in BIOS setup before the data can be accessed.

If the customer chooses to enable passphrase protection or not, BIOS locks the DCPMM before booting to the operating system or UEFI Shell. This means that all security changes are controlled by the Dell BIOS and operating system level security changes including Passphrase management and DCPMM erasing functions will not be supported. All these functions must be driven through the BIOS setup.

NOTE: As mentioned in section DIMM Configuration Changes, the only migration scenario that is supported is a slot for slot replacement between motherboards. Adding or removing individual DCPMM for any reason will likely result in data loss and trigger the need for goal and security reconfiguration.

The passphrase to lock or encrypt the data at rest on the DCPMM in App-direct is configurable in the BIOS setup. If the field is not empty, every boot the supplied passphrase is used to attempt to unlock all DCPMM in the system.

The following use cases are related to unsupported migration scenarios:

- When changing passphrase in the BIOS setup, the existing passphrase only needs to be entered once per session. Entering and existing the field multiple times will not reprompt for the passphrase again (until the next boot session).
- Passphrase can be cleared by entering empty string in BIOS setup passphrase field.

NOTE: To clear the passphrase, keep the passphrase field blank and hit Enter.




Figure 12. Clearing Passphrase

Cryptographic erase and DCPMM sanitize

There are two ways to erase the persistent region content (App-direct) of the DCPMM:

- Crypto Erase
- Sanitize

Both erase methods can be executed using BIOS setup options. User can choose to perform an erase on all or a subset of installed DCPMMs.

Crypto erase

The Crypto Erase function erases the App-direct Region Key (PM-RK) forcing the system to reboot.

Cryptographic erase option can be accessed by going to: **System BIOS Settings > Memory Settings > Persistent Memory > Intel Persistent Memory > Persistent Memory DIMM Configuration**

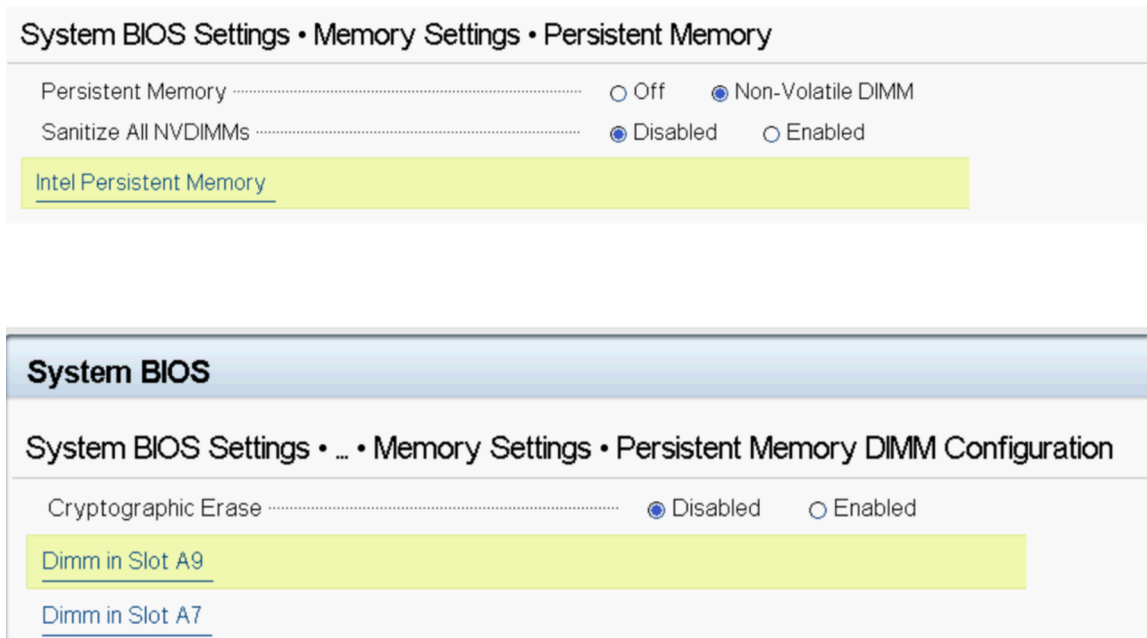


Figure 13. Crypto erase

NOTE: It is not recommended to erase part of DCPMMs installed on the system when App-direct interleaved region is configured. This operation renders all the data on the interleave set invalid.

Sanitize

DCPMM Sanitization is a long operation which runs on all the selected persistent memories in parallel.

This process does a cryptographic erase first and writes zeros to all accessible persistent media regions on the DCPMM. Also, it destroys any existing memory goal contents.

The region of the DCPMM where the goal is configured will be empty and on the next boot the memory will default to 100% Memory Mode. If the system does not have the proper ratio of RDIMM or LRDIMM memory to DCPMM, it is expected on the following boot for an error to occur indicating that the ratio is not optimized.

Sanitize option can be accessed by going to: **System BIOS Settings > Memory Settings > Persistent Memory**




Figure 14. Sanitize

Sanitize can take up to 15 minutes with fully loaded 128 GB DIMM configuration, 30 minutes with 256 GB and 1 hour with 512 GB.

i **NOTE: Sanitize is not supported when DCPMMs are configured in Memory mode. When the Sanitize operation is running, a prompt appears in BIOS indicating an Overwrite. Overwrite is the name for the second firmware command that is conducted. The first command which happens quickly and will not be displayed on-screen is the Crypto Erase (firmware command name is "Secure Erase").**

DIMM configuration changes

The following DCPMM migration scenarios are supported:

- Replacement of System Board due to failure

All DIMMs must be re-populated in the exact same slots. DCPMMs and data content will be available for customer application access after the board has been restored to the same configuration as the original board. System Restore will automatically restore the BIOS configuration on the replacement board, including the DCPMM Passphrase, if it is set.

- Replacing bad DIMM

In the case of a failed DCPMM, any data associated with the DCPMM would be lost. The region and interleave set on the failed DCPMM must be recreated once it is replaced. The user must use BIOS Setup to create a new goal for the affected DCPMM.

- ① **NOTE:** Any remaining persistent memory data on Intel DCPMM must be backed up prior creating a goal. The goal creation process deletes all namespaces, regions, and data stored on the DCPMMs on the selected CPUs. If a Security Passphrase is enabled, the new persistent memory region is protected with the system DCPMM passphrase.
- ① **NOTE:** Adding or removing DCPMMs to an existing DCPMM configuration is not supported, and not validated. It is recommended that customers back up all DCPMM data to another storage device before making any DCPMM configuration changes. Once the server is reconfigured to the new DCPMM configuration, the customer can create a goal configuration and restore data back to DCPMMs.

Dell EMC supports Intel DCPMM with Microsoft Windows 2019 in Memory mode and App-direct mode.

NOTE: Keep Windows updated with the monthly cumulative updates.

Topics:

- [PMEM in App-direct mode](#)
- [PMEM disk with interleave sets](#)
- [PMEM in memory Mode](#)
- [Windows troubleshooting and event monitoring](#)

PMEM in App-direct mode

In App-direct mode, Windows creates two types of device objects for DCPMMs :

- Physical INVDIMM device
- Logical Persistent Memory disk

Logical persistent memory disks get created after creating namespaces on PMEM Physical disks.

When the system first boots into the Operating System, PMEM devices get enumerated as physical INVDIMM devices under Memory Devices in Device Manager.

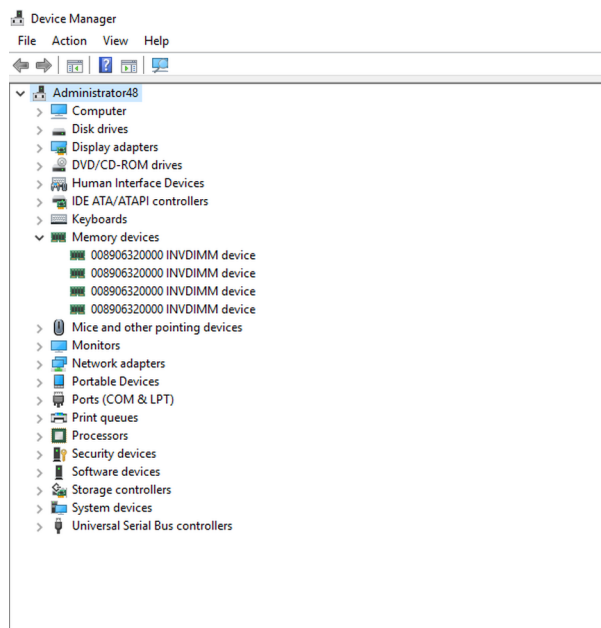


Figure 15. Memory Devices in Device Manager

After PMEM disks are configured over PowerShell, logical persistent memory disks show under Device Manager.




Figure 16. Memory Devices in Device Manager

PMEM Disk management

Windows currently supports only one namespace per interleave set (this is independent of the number of physical devices in the interleave set). The option to interleave DCPMMs can be selected during goal creation as described in App Direct and Memory mode configurations.

PMEM disks have to be created with the help of "New-Pmemdisk" command by providing relevant region IDs. Once PMEM disks are configured, PMEM volumes can be used as normal disks.

Windows supports the following PowerShell cmdlets to manage persistent memory:

- `Get-PmemDisk`
 - Returns one or more logical persistent memory disks.
 - The returned object has information about size, atomicity type, health status, and underlying physical devices.
- `Get-PmemPhysicalDevice`
 - Returns one or more physical persistent memory devices (NVDIMMs).
 - The returned object has information about size(s), RFIC, device location, and health/operational status.
- `New-PmemDisk`
 - Creates a new disk out of a given unused region.
 - Writes out the labels to create the namespace then rebuilds the SCM stacks to expose the new logical device.
 - Optional parameters:
 - `FriendlyName` gives the persistent memory disk a friendly name. Default is "PmemDisk<N>".
 - `AtomicityType` lets you set BTT. Default is "none".
- `Remove-PmemDisk`
 - Removes the given persistent memory disk. It accepts the output of `Get-PmemDisk`.
 - Deletes the namespace's labels and then rebuilds the SCM stacks to remove the logical device.
 - Requires user confirmation, which can be overridden with `Force`.
- `Get-PmemUnusedRegion`
 - Returns aggregate PMEM regions available for provisioning a logical device.
 - Returned object has a unique region ID, total size, and list of physical devices that contribute to the unused region.
- `Initialize-PmemPhysicalDevice`
 - Writes zeroes to the label storage area, writes new label index blocks, and then rebuilds the SCM stacks to reflect the changes.

- Requires user confirmation, which can be overridden with `Force`.
- This cmdlet is intended as a "big hammer" recovery mechanism. It is not recommended for normal use.

List PMEM physical disks and check their health status

The following image displays command usage to list all usage to list all PMEM physical devices and get their health. Physical location gives the location of the DIMM on the motherboard.

```
PS C:\Users\Administrator> Get-PmemPhysicalDevice
DeviceId DeviceType HealthStatus OperationalStatus PhysicalLocation FirmwareRevision Persistent memory size Volatile memory size
-----
1 008906320000 INVDIMM device Healthy {Ok} A7 102005375 126 GB 0 GB
1001 008906320000 INVDIMM device Healthy {Ok} B7 102005375 126 GB 0 GB
1011 008906320000 INVDIMM device Healthy {Ok} B8 102005375 126 GB 0 GB
11 008906320000 INVDIMM device Healthy {Ok} A8 102005375 126 GB 0 GB

PS C:\Users\Administrator> Get-PmemUnusedRegion
RegionId TotalSizeInBytes DeviceId
-----
1 135291469824 {1}
2 135291469824 {11}
4 135291469824 {1001}
5 135291469824 {1011}
```

Figure 17. List PMEM Physical Disks and their health status

NOTE: If Health Status is not Healthy and Operational Status is not OK, the issue needs to be rectified before creating namespaces. Customers can run into this scenario if DCPMMs were previously used with another operating system and were booted with Windows without Sanitizing. If such a scenario occurs, right-click and uninstall all the Memory devices and Persistent Memory disks from Device Manager and then scan for HW changes under Action in Device Manager. This resolves any issues with the driver stack.

Create PMEM Disks

```
PS C:\Users\Administrator> New-PmemDisk -RegionId 1,2,4,5
Creating new persistent memory disk. This may take a few moments.
Creating new persistent memory disk. This may take a few moments.
Creating new persistent memory disk. This may take a few moments.
Creating new persistent memory disk. This may take a few moments.
PS C:\Users\Administrator> Get-PmemDisk
DiskNumber Size HealthStatus AtomicityType CanBeRemoved PhysicalDeviceIds UnsafeShutdownCount
-----
4 126 GB Healthy None True {1} 4
5 126 GB Healthy None True {11} 2
6 126 GB Healthy None True {1001} 2
7 126 GB Healthy None True {1011} 2

PS C:\Users\Administrator> get-disk -FriendlyName "Persistent memory disk"
Number Friendly Name Serial Number HealthStatus OperationalStatus Total Size Partition Style
-----
4 Persistent memory disk 030180891a09c42a6cd24b47bc97e... Healthy Online 126 GB RAW
5 Persistent memory disk 0301808982ebc27b571edf4280762... Healthy Online 126 GB RAW
6 Persistent memory disk 03018089b305ff37f7de704e9a0f1... Healthy Online 126 GB RAW
7 Persistent memory disk 03018089c27158d1b96dc74d9640b... Healthy Online 126 GB RAW
```

Figure 18. Create PMEM Disks

Remove PMEM disks

```
PS C:\Users\Administrator> Get-PmemDisk | Remove-PmemDisk
This will remove the persistent memory disk(s) from the system and will result in data loss.
Remove the persistent memory disk(s)?
[Y] Yes [A] Yes to All [N] No [L] No to All [S] Suspend [?] Help (default is "Y"): Y
Removing the persistent memory disk. This may take a few moments.

This will remove the persistent memory disk(s) from the system and will result in data loss.
Remove the persistent memory disk(s)?
[Y] Yes [A] Yes to All [N] No [L] No to All [S] Suspend [?] Help (default is "Y"): A
Removing the persistent memory disk. This may take a few moments.
Removing the persistent memory disk. This may take a few moments.
Removing the persistent memory disk. This may take a few moments.
PS C:\Users\Administrator> Get-PmemDisk
PS C:\Users\Administrator> Get-PmemUnusedRegion

RegionId TotalSizeInBytes DeviceId
-----
1 135291469824 {1}
2 135291469824 {11}
4 135291469824 {1001}
5 135291469824 {1011}

PS C:\Users\Administrator> Get-PmemPhysicalDevice

DeviceId DeviceType HealthStatus OperationalStatus PhysicalLocation FirmwareRevision Persistent memory size Volatile memory size
-----
1 008906320000 INVDIMM device Healthy {Ok} A7 102005375 126 GB 0 GB
1001 008906320000 INVDIMM device Healthy {Ok} B7 102005375 126 GB 0 GB
1011 008906320000 INVDIMM device Healthy {Ok} B8 102005375 126 GB 0 GB
11 008906320000 INVDIMM device Healthy {Ok} A8 102005375 126 GB 0 GB
```

Figure 19. Remove PMEM disks

Windows 2019 does not support redundant volume creation on PMEM disks by using Windows VDS (Virtual disk service).

In order to create redundant volumes, use the storage spaces method.

For information about storage spaces method refer to: <https://docs.microsoft.com/en-us/windows-server/storage/storage-spaces/deploy-standalone-storage-spaces>.

PMEM disk with interleave sets

Interleaved sets can often be created to make multiple persistent memory devices be displayed as a single logical disk to windows server. For PMEM Disk with the interleave set, "App-direct Interleave" should be selected during goal configuration.

PMEM disk creation with interleave sets

When Interleave set is enabled, BIOS assigns single Region Id for DCPMMs connected to the same processor.

```
PS C:\Users\Administrator> Get-PmemPhysicalDevice

DeviceId DeviceType HealthStatus OperationalStatus PhysicalLocation FirmwareRevision Persistent memory size Volatile memory size
-----
1 008906320000 INVDIMM device Healthy {Ok} A7 102005375 126 GB 0 GB
1001 008906320000 INVDIMM device Healthy {Ok} B7 102005375 126 GB 0 GB
1011 008906320000 INVDIMM device Healthy {Ok} B8 102005375 126 GB 0 GB
11 008906320000 INVDIMM device Healthy {Ok} A8 102005375 126 GB 0 GB

PS C:\Users\Administrator> Get-PmemUnusedRegion

RegionId TotalSizeInBytes DeviceId
-----
1 270582939648 {1, 11}
3 270582939648 {1001, 1011}
```

Figure 20. PMEM Disk creation with interleave sets

In above example two Interleave sets have been created, one Interleave set using slots A7, A8 and another Interleave set using B7, B8.

We can create PMEM Disk also using Region id.

```

PS C:\Users\Administrator> Get-PmemUnusedRegion | New-PmemDisk
Creating new persistent memory disk. This may take a few moments.
Creating new persistent memory disk. This may take a few moments.
PS C:\Users\Administrator> Get-PmemDisk

```

DiskNumber	Size	HealthStatus	AtomicityType	CanBeRemoved	PhysicalDeviceIds	UnsafeShutdownCount
4	252 GB	Healthy	None	True	{1, 11}	6
5	252 GB	Healthy	None	True	{1001, 1011}	4

```

PS C:\Users\Administrator>

```

Figure 21. PMEM Disk creation using Region id

PMEM in memory Mode

When Intel DCPMM is configured in memory mode, operating system sees it as system memory.

Persistent memory size is shown as zero and volatile memory size accounts for the entire size of the DCPMM.

```

PS C:\Users\Administrator> Get-PmemPhysicalDevice

```

DeviceId	DeviceType	HealthStatus	OperationalStatus	PhysicalLocation	FirmwareRevision	Persistent memory size	Volatile memory size
1	008906320000 INVDIMM device	Healthy	{Ok}	A7	102005375	0 GB	126 GB
1001	008906320000 INVDIMM device	Healthy	{Ok}	B7	102005375	0 GB	126 GB
1011	008906320000 INVDIMM device	Healthy	{Ok}	B8	102005375	0 GB	126 GB
11	008906320000 INVDIMM device	Healthy	{Ok}	A8	102005375	0 GB	126 GB

```

PS C:\Users\Administrator> Get-PmemUnusedRegion
PS C:\Users\Administrator>

```

Figure 22. PMEM in memory Mode

Windows troubleshooting and event monitoring

If any of the PMEM physical devices or logical devices are not functioning properly, it is suggested to check Windows event logs.

To see the logs, open Event Viewer and navigate to : **Applications and Services Logs > Microsoft > Windows**

The names of all persistent memory drivers logs start with “PersistentMemory.”

All runtime errors are logged to the “Operational” log. This log captures the complete operation of PMEM physical device(NVDIMM) and PMEM logical device (PMEMDisk).




Figure 23. Windows troubleshooting and event monitoring

Identify and configure persistent memory device

Listing DCPMM devices

To list all physical devices in the system, run command

```
ndctl list -DH
```

Create namespace

The configuration of namespaces will decide how much memory capacity user wants to expose to the OS.

To configure namespace, run the following command:

```
ndctl create-namespace
```

Following command needs to be executed for each region created when the goals were created for the system:

```
ndctl create-namespace -r regionX
```

After creating name spaces, all name spaces are shown using the following command:

```
ndctl list -N
```

To correlate a namespace to a pmem device, use the following command:

```
lsblk
```

NOTE: Ensure you delete all the namespaces before repurposing the DCPMMs. Repurposing includes DCPMM migration, removal and reconfiguration between memory and App-direct mode

Mount file system on namespace device

After you see the PMEM devices in /device folder, mount the filesystems by using the following commands:

```
mkfs.xfs -f /dev/pmemX
```

(X is a natural number ranging from 0 to number of Logical PMEM Devices)

```
mkdir /mnt/pmemX  
mount -o dax /dev/pmemX /mnt/pmemX
```

To write data into the device, run the following command:

```
cd /mnt/pmemX  
echo "Hello World" >>test.txt
```

Reset the system and the data should be persistent over power cycle.

Delete namespaces

Namespace can be deleted using NDCTL command:

```
ndctl destroy-namespace <namespace>
```

<namespace> is namespaceX.Y device name which can be retrieved using the following command:

```
ndctl list -N
```

To delete all the namespaces on the system, you may have to send the command multiple times

 **NOTE: Reconfiguration of the goal or region from BIOS setup will automatically delete all the namespace on the system.**

Management utility

Linux distributions use native tool ndctl to manage nonvolatile devices.

For comprehensive list of commands and syntax, see open-source resource pmem.io.

Check DCPMM health status

Following NDCTL command shows the health status of each DCPMMs installed on the system:

```
ndctl list -DHl
```

Health Information includes:

Table 6. Health information

Health fields	Comment
Overall Health State	Critical, Non-critical, ok
DIMM Temperature	In degree C
Spare Percentage/Lifetime percentage	In %
User configured alarm temperature	Smart controller temperature threshold for alarm
User configured alarm spares	Smart spares threshold for alarm
Shutdown state	Dirty or Clean. Determines if data has been successfully cleaned to persistent domain during previous power cycle.

Linux errata

1. If an uncorrectable error occurs in the early metadata region of the DCPMM, the system can get into an infinite loop of being unresponsive when booting into Linux.

Workaround: Boot with "*modprobe.blacklist=nd_pmem*" on the kernel command line to stop the consumption of the error and enable boot progress. Afterwards, wait sufficient time for the Address Range Scrub feature to scrub the memory and identify the bad address location so it will be mapped out on future boots.

2. Linux boot fails to Emergency Mode when DCPMM is configured to automount during boot in fstab.

Workaround: Add "*x-systemd.device-timeout=0*" to the mount options for the PMEM partition in fstab.

Example: `/dev/pmem5 /mnt/somedir ext4 defaults, x-systemd.device-timeout=0,dax 0 2`

3. Dell BIOS does not support boot from DCPMM. Some Linux operating system (Red Hat Enterprise Linux 7.6, Red Hat Enterprise Linux 8.0) distros are offering early support of this feature, but Dell BIOS does not support this feature.

Workaround:None.

VMware ESXi

Dell EMC supports Intel DCPMM from vSphere 6.7 EP10 (Build #13981272) and above. The support is available for Intel DCPMM Memory mode and App-direct modes.

The persistent memory inventory details are available as part of Host client. See **Managing Persistent Memory** in the **vSphere Single Host Management - VMware Host Client** which detail the persistent memory data that are presented to the user.

Following sections talk about how ESXi displays the persistent memory details in different operating modes.

Topics:

- [PMEM in App-direct mode](#)
- [PMEM in Memory mode](#)
- [PMEM health status](#)
- [ESXi troubleshooting and maintenance](#)

PMEM in App-direct mode

Log in to the Host client using ESXi credentials. Navigate to **Storage > Persistent Memory** to view the persistent memory modules, the interleave sets created and the namespaces that are created on ESXi.




Figure 24. Persistent Memory modules populated in the system




Figure 25. Number of Interleave sets created on the system

The number of interleave sets are created based on the 'interleave' option that is selected during creation of the goal. When **Interleaved** is selected, the number of interleave sets exposed would be equal to the number of CPU sockets in the system. When it is set to **not interleaved**, individual interleave sets are created for each DCPMM module. Unlike Linux, ESXi automatically creates namespaces on the interleave sets exposed.




Figure 26. Namespaces created on ESXi out of the interleave sets exposed

After creating the namespaces, ESXi automatically creates a pMem datastore and mounts it as a datastore for the users to consume it.




Figure 27. PMEM Datastore

PMEM in Memory mode

When Intel DCPMM is configured in memory mode, ESXi sees it as system memory.

NOTE: The inventory details in the vSphere HTML client will not be available if Intel DCPMM is set to Memory mode. There is no datastore created in memory mode as ESXi creates in Intel DCPMM App-direct mode.

PMEM health status

ESXi provides multiple health status for PMEM such as "Maintenance needed", "All data loss" and "Normal".

ESXi reports the health of the specific DCPMMs as 'Maintenance needed', when there is a health error.



Figure 28. PMEM health warning

ESXi reports the health of the specific DCPMMs as "All data loss", when there is a fatal error.




Figure 29. PMEM multibit error

ESXi troubleshooting and maintenance

NOTE: Go through `/var/log/vmkernel.log` to see if there are any errors or warnings reported related to persistent memory. The log entries related to persistent memory would be appended with "NVD" and/or "IntelINVDimm"

- Configuring Intel DCPMM in memory mode shows a reduction of system memory in system BIOS.
 - For example, a system configuration with 4 * 128 GB Intel DCPMMs, configured in memory mode provides a system memory of 504 GB instead of 512 GB.
 - It is for reserving space for metadata.
- PMEM datastore is not mounted automatically in ESXi.
 - Refer to `/var/log/vmkernel.log` and check for 'shut down counter' related errors to DIMM handles.
 - When there is a critical error to any of the Intel DCPMMs, datastore is not mounted by ESXi.

- ESXi boot takes longer when Intel DCPMMs are configured as noninterleaved App-Direct.
 - It can take anywhere 70 to 80 seconds per interleaved set to 2 hours depending on the capacity, configuration, and number of sockets. This is an expected behavior due to ARS performed on each DCPMM interleave set sequentially. The Vmkernel log may contain entries similar to "Attempt to allocate zero bytes, allocating 1 byte." These messages can be safely ignored.
- If DCPMMs contain a critical error, the ESXi Host Client (HTML client) GUI does not respond when the user moves to Persistent Memory section under Storage. The host server is unresponsive when it reads the PMEM region address range with the critical error. For more information see VMware KB 70661: <https://kb.vmware.com/s/article/70661>.
- In ESXi, Windows as a guest operating system does not display Storage Class Memory device in the device manager, for the PMEM devices that are attached to the virtual machine (VM).

Virtual BIOS for the VM hides the device from device manager. User can use the PowerShell cmdlets to view the PMEM devices that are attached to the windows guest operating systems running on ESXi.

- 'esxstop' utility in ESXi exposes some dummy performance counters when DCPMMs are configured in the system. Counter names might be displayed as 'Cache hit', 'Cache miss', 'nBuffers', 'Flush interval'. These are dummy performance counters and no functionalities that are associated with it for now.
- ESXi displays more number of NUMA nodes on systems that are configured with DCPMMs in App-direct mode. This is an expected behavior as the NUMA nodes are created for volatile and nonvolatile address ranges. The nonvolatile address ranges are displayed as 0 MB.
- ESXi logs a warning in the Vmkernel log 'Unable to register PMEM filesystem for APD notifications!'. This log entry can be safely ignored as the PSA (Pluggable Storage Architecture) APD(All Paths Down) plug in is not supported for PMEM type volumes. For more information see VMware KB 2145444: <https://kb.vmware.com/s/article/2145444>.

System diagnostics

System diagnostics in the Lifecycle Controller does not test Intel DCPMM in App-direct mode to avoid destroying customer data.

 **NOTE: System diagnostics are not recommended to investigate Intel DCPMM failures when in Memory mode due to extremely long test time.**

Firmware update

Dell DUP update

Download DCPMM DUP from www.dell.com/support/drivers.

NOTE: Make sure BIOS is up to date with the latest version before applying DCPMM DUP.

1. Boot to the operating system and execute DUP package.
 - a) On a Windows system, double click the .exe DUP file. Reboot the system and DUP will be automatically executed.
 - b) On Linux system:
 1. Navigate to the folder that contains DUP bin file
 2. Give read write permissions to the DUP bin file (Chmod 777)
 3. Execute the DUP bin file (./)
2. If operating systems is not accessible, log in to iDRAC. Navigate to **Maintenance > System Update**, upload and install the **.exe** file downloaded from the Dell support site.

The screenshot shows the iDRAC web interface. At the top, there is a navigation bar with links for Dashboard, System, Storage, Configuration, Maintenance, and iDRAC Settings. Below this, the 'Maintenance' section is active, with sub-tabs for Lifecycle Log, Job Queue, System Update (selected), System Event Log, Troubleshooting, Diagnostics, and SupportAssist. Under 'System Update', there are three sub-sections: Manual Update (selected), Automatic Update, and RollBack. The 'Manual Update' section contains a 'Location Type' dropdown menu set to 'Local'. Below that is a 'Single Update Location' section with an 'Update Path*' field and a 'Choose File' button. The text 'No file chosen' is displayed next to the button. At the bottom of this section is an 'Upload' button.

Figure 30. System update via iDRAC

3. The system will automatically reboot and flash the firmware using Lifecycle Controller. Multiple reboots are expected.
4. Firmware can also be updated using the Lifecycle Controller:
 - a. Launch Lifecycle controller(LC) during boot.
 - b. Launch "Firmware Update" from Lifecycle Controller.
 - c. Select FTP server or local drive that hosts the downloaded DUP and launch the update.
 - d. DUP update gets staged through LC.

For more information on updating firmware's using Lifecycle Controller see: *Lifecycle Controller User's Guide* available at www.dell.com/idracmanuals .