

Dell EMC NVDIMM-N Persistent Memory

User 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.

© 2018 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.

Contents

1 Introduction.....	5
2 Change list.....	6
3 NVDIMM-N Overview.....	7
Normal Operation.....	7
Backup to Flash.....	8
Restore from Flash.....	9
4 Hardware.....	10
Server Hardware Configuration.....	10
Modular Chassis Hardware Configuration.....	14
NVDIMM-N Module Details.....	14
Battery.....	15
Minimum Platform Firmware Versions.....	17
5 BIOS.....	18
BIOS Configuration Settings for NVDIMM-N.....	18
BIOS Error Messages.....	22
6 iDRAC NVDIMM-N Management.....	25
iDRAC Graphical User Interface.....	25
NVDIMM-N Status.....	25
BBU Status.....	26
Log Messaging Errata.....	27
Remote Management.....	27
NVDIMM-N Error Reporting.....	27
7 Server Behavior with NVDIMM-Ns.....	30
Shutdown.....	30
Boot.....	31
Automatic Shutdown and Save.....	31
8 DIMM Configuration Changes.....	32
9 Windows.....	33
BIOS Requirements.....	33
Set Up.....	33
Windows Drivers.....	33
Storage Class Memory in Windows Server 2016.....	34
Device manager.....	34
Identifying the right NVDIMM-N disks.....	34

NVDIMM-N health status and properties.....	37
Block Mode.....	37
DAX Mode.....	37
Storage Spaces Support	38
Operational and Diagnostics Logging information.....	38
Storage Class Memory in Windows Server 2019.....	39
NVDIMM-N FW Requirement.....	39
Driver Architecture Overview.....	39
New features in Windows Server 2019 Label support and Namespace management.....	40
PowerShell Cmdlets.....	41
NVDIMM-N Interleaving.....	42
Configuring NVDIMM-N for Hyper-V Virtual Machines.....	43
NVDIMM-N RO Behavior.....	43
Windows Errata.....	43
10 Linux.....	45
Identify and Configure PMEM —Persistent Memory Device.....	45
Installation.....	46
Verify Existing Filesystem.....	46
Read-Only Mode NVDIMM-N.....	46
Interleave.....	46
Interleave Setup.....	46
Interleave Verification.....	47
Read Only Mode NVDIMM-N.....	47
Management Utility	48
ndctl.....	48
mdadm.....	49
RHEL 7.6 features.....	49
Namespace creation.....	49
Linux Errata.....	50
11 ESXi.....	51
Set up.....	51
Storage.....	52
Modules.....	52
Namespaces.....	52
Interleave sets.....	53
Datastore.....	54
Supported Guest OSES with NVDIMM support.....	55
Overall Health Status	55
Operational and Diagnostics Logging information.....	55
Outdated firmware.....	55
NVDIMM-N Errors.....	56
ESXi Errata.....	57

Introduction

DellEMC's NVDIMM-N Persistent Memory is a disruptive Storage Class Memory technology that enables unprecedented performance improvement over legacy storage technologies. Each NVDIMM-N provides 16GB of nonvolatile memory and has the same form factor as a standard 288-Pin DDR4 DIMM. The NVDIMM-N resides in a standard CPU memory slot, placing data close the processor. With its ability to operate at 2666MT/s DDR4 data rates, the NVDIMM-N takes full advantage of the high bandwidth and extremely low latency that is characteristic of the memory bus. For comparison, the table below provides approximate data access times for DDR4 relative to other server storage media.

Table 1. Storage Technology

Storage Technology	Data Access Time
15K SAS Disk	~ 6,000,000 ns
SATA SSD	~ 120,000 ns
NVMe SSD	~ 60,000 ns
DDR4 NVDIMM-N	~ 150 ns

This document provides an overview of the DellEMC NVDIMM-N Persistent Memory solution. It is intended to help the user with initial set-up and configuration, providing information on system behavior when NVDIMM-Ns are installed. This document will also help the user to become familiar with NVDIMM-N manageability and error handling. And lastly, it introduces the user to NVDIMM-N basic setup and configuration in a variety of supported operating systems.

Change list

Table 2. Change list

Version	Changes
1.0	Original Version
1.1	Added ESXi 6.7 support information. Removed Linux errata that is no longer applicable. Edits to remainder of document for clarity.
1.2	Added Modular Server specific information, support for R840, R940xa, changes to the BBU LED behavior and edits to the remainder of document for clarity. NVDIMM-N supported on RHEL 7.5
1.3	Added minimum supported revisions Support for Windows 2019, RHEL 7.6 and ESXi 6.7 U1. Windows 2019 new features and NVDIMM-N RO behavior RHEL 7.6 features and bug fix

NVDIMM-N Overview

The Figure below is an overview of the NVDIMM-N showing its main components and system interfaces. Core to the NVDIMM-N are the DDR4 DRAM devices that allow the NVDIMM-N to operate as an RDIMM. The components that allow the NVDIMM-N to persist data are the Controller, Flash, and Power Voltage Regulators that are also integrated on the DIMM.




Figure 1. NVDIMM-N Overview

Topics:

- [Normal Operation](#)
- [Backup to Flash](#)
- [Restore from Flash](#)

Normal Operation

NVDIMM-Ns are installed in server memory slots. From a hardware perspective, the NVDIMM-Ns appear to the processor as standard DDR4 RDIMMs. They are initialized during BIOS POST, and the CPU can access DRAM data on the NVDIMM-N using standard DDR4 memory load/store transactions.

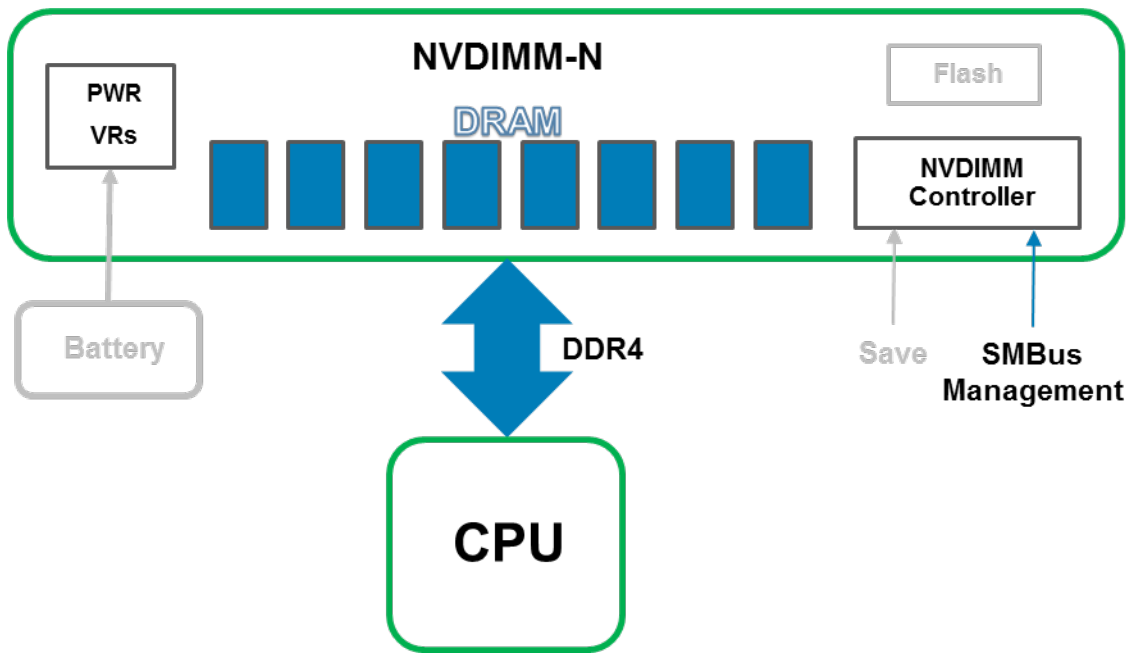


Figure 2. NVDIMM-N Normal Operation

Backup to Flash

In the event of a server shutdown, cold reboot, or power loss, a Save signal is sent to the NVDIMM-N Controller which then triggers the NVDIMM-N Controller to backup all its DRAM contents to its onboard flash storage. The NVDIMM-N Save event is triggered anytime the server is about to power down and power loss to NVDIMM-Ns is imminent. The backup process takes approximately one minute to complete. This duration is independent of the number of NVDIMM-Ns installed in the server because Saves occur in parallel across all NVDIMM-Ns.

A Battery provides backup power to the NVDIMM-N so that it can complete the backup operation even after an power loss. When the Save operation is complete, the Battery is deactivated to allow the server and NVDIMM-Ns to power down completely.

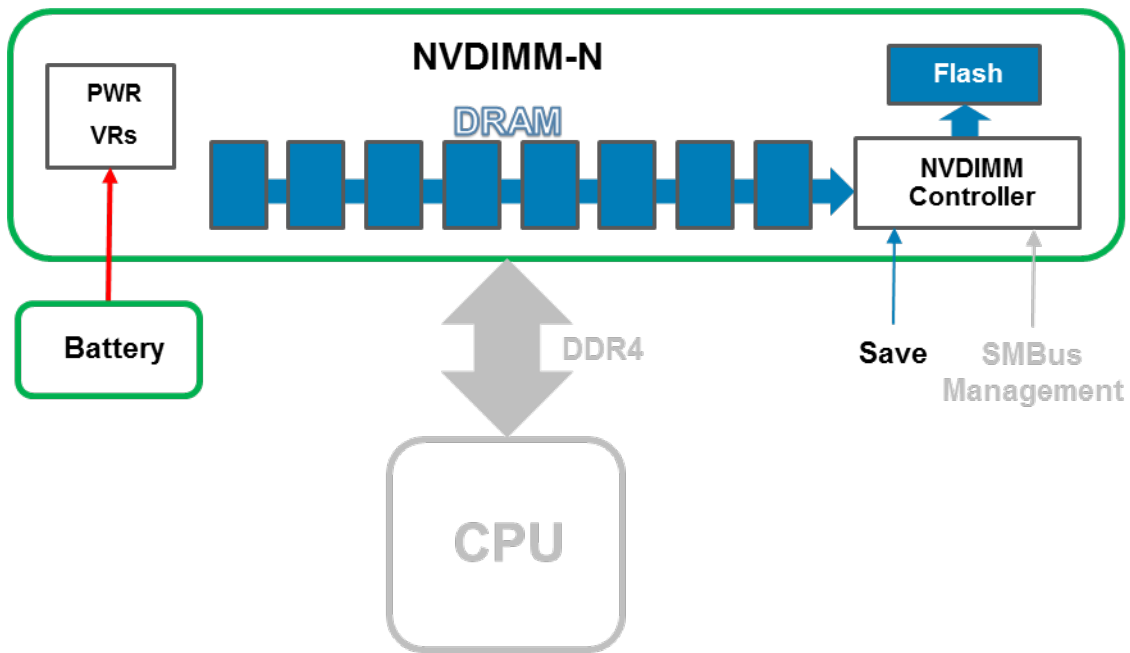


Figure 3. Save Operation

Restore from Flash

On server power-up, BIOS re-initializes the DRAM on the NVDIMM-N. BIOS commands the NVDIMM-N Controller using the SMBus Management Interface to restore its DRAM contents from Flash. The restore process takes approximately one minute to complete. This duration is independent of the number of NVDIMM-Ns installed in the server because Restores occur in parallel across all NVDIMM-Ns. BIOS then exposes the NVDIMM-N to the Server OS as Persistent Memory. Note that Persistent Memory is distinct from System Memory which is required for normal OS operation. Please refer to the respective OS Sections for more details on OS support of Persistent Memory.

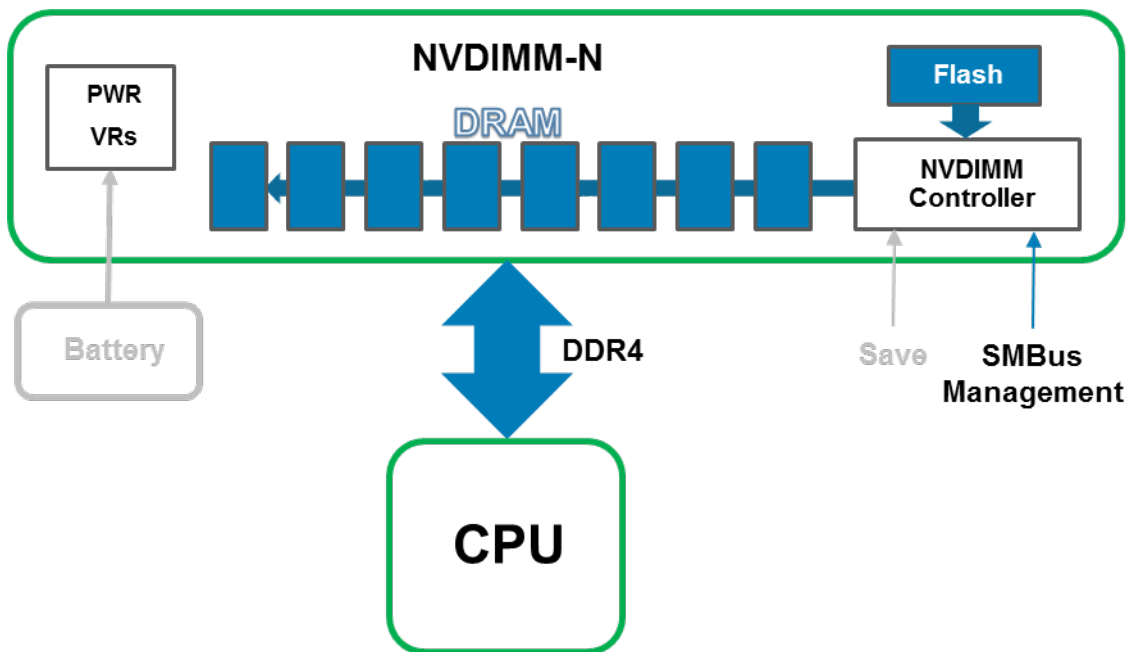


Figure 4. Restore Operation

Topics:

- [Server Hardware Configuration](#)
- [Modular Chassis Hardware Configuration](#)
- [NVDIMM-N Module Details](#)
- [Battery](#)
- [Minimum Platform Firmware Versions](#)

Server Hardware Configuration

NVDIMM-Ns are currently supported in the T640, R640, R740/R740XD, R840, R940, R940xa, MX740c and MX840c PowerEdge Servers. Each server supports from 1x to a maximum of 12x 16GB NVDIMM-Ns for a total max persistent memory capacity of 192GB. Table 1 below provides the NVDIMM-N configurations that have been validated and are fully supported for 2-Socket server configurations. For 4-Socket configurations in the R840, R940, R940xa and MX840c PowerEdge Server, NVDIMM-Ns are only supported in CPU 1 and CPU 2 memory slots. Table 2 below provides the NVDIMM-N configurations that have been validated and are fully supported on the R840, R940, R940xa and MX840c 4-Socket configurations.

Refer to Figure 5 for the R740/R740XD CPU and DIMM slot locations as an example. Please see the respective server's Installation and Service Manual for memory installation guidelines.




Figure 5. R740/R740XD Memory Layout

Table 3. Supported NVDIMM-N Configurations, 2-Socket

NVDIMM-N	NVDIMM-N Capacity	RDIMMs	RDIMM Capacity	DIMM Population Locations
1x	16GB	12x 16GB	192GB	RDIMMs: A1-A6, B1-B6 NVDIMM-N: A7
1x	16GB	12x 32GB	384GB	RDIMMs: A1-A6, B1-B6 NVDIMM-N: A7
1x	16GB	23x 32GB	736GB	RDIMMs: A1-A12, B1-B11 NVDIMM-N: B12
2x	32GB	12x 16GB	192GB	RDIMMs: A1-A6, B1-B6 NVDIMM-Ns: A7, B7
2x	32GB	12x 32GB	384GB	RDIMMs: A1-A6, B1-B6 NVDIMM-Ns: A7, B7
2x	32GB	22x 32GB	704GB	RDIMMs: A1-A11, B1-B11 NVDIMM-Ns: A12, B12
4x	64GB	12x 16GB	192GB	RDIMMs: A1-A6, B1-B6 NVDIMM-Ns: A7-A8, B7-B8
4x	64GB	12x 32GB	384GB	RDIMMs: A1-A6, B1-B6 NVDIMM-Ns: A7-A8, B7-B8
4x	64GB	20x 32GB	640GB	RDIMMs: A1-A10, B1-B10 NVDIMM-Ns: A11-A12, B11-B12
6x	96GB	12x 16GB	192GB	RDIMMs: A1-A6, B1-B6 NVDIMM-Ns: A7-A9, B7-B9
6x	96GB	12x 32GB	384GB	RDIMMs: A1-A6, B1-B6 NVDIMM-Ns: A7-A9, B7-B9
6x	96GB	18x 32GB	576GB	RDIMMs: A1-A9, B1-B9 NVDIMM-Ns: A10-A12, B10-B12

NVDIMM-N	NVDIMM-N Capacity	RDIMMs	RDIMM Capacity	DIMM Population Locations
12x	192GB	12x 16GB	192GB	RDIMMs: A1-A6, B1-B6 NVDIMM-Ns: A7-A12, B7-B12
12x	192GB	12x 32GB	384GB	RDIMMs: A1-A6, B1-B6 NVDIMM-Ns: A7-A12, B7-B12

Note: While other configurations may work, they have not been fully validated and are not currently supported by DellEMC.




Figure 6. MX740c Memory Layout

Table 4. Supported NVDIMM-N Configurations, R940, MX840c 4-Socket 4-Socket

NVDIMM-N	NVDIMM-N Capacity	RDIMMs	RDIMM Capacity	DIMM Population Locations
1x	16GB	24x 16GB	384GB	RDIMMs: A1-A6, B1-B6, C1-C6, D1-D6

NVDIMM-N	NVDIMM-N Capacity	RDIMMs	RDIMM Capacity	DIMM Population Locations
				NVDIMM-N: A7
1x	16GB	24x 32GB	768GB	RDIMMs: A1-A6, B1-B6, C1-C6, D1-D6 NVDIMM-N: A7
1x	16GB	47x 32GB	1504GB	RDIMMs: A1-A12, B1-B11, C1-C12, D1-D12 NVDIMM-N: B12
2x	32GB	24x 16GB	384GB	RDIMMs: A1-A6, B1-B6, C1-C6, D1-D6 NVDIMM-Ns: A7, B7
2x	32GB	24x 32GB	768GB	RDIMMs: A1-A6, B1-B6, C1-C6, D1-D6 NVDIMM-Ns: A7, B7
2x	32GB	46x 32GB	1472GB	RDIMMs: A1-A11, B1-B11, C1-C12, D1-D12 NVDIMM-Ns: A12, B12
4x	64GB	24x 16GB	384GB	RDIMMs: A1-A6, B1-B6, C1-C6, D1-D6 NVDIMM-Ns: A7-A8, B7-B8
4x	64GB	24x 32GB	768GB	RDIMMs: A1-A6, B1-B6, C1-C6, D1-D6 NVDIMM-Ns: A7-A8, B7-B8
4x	64GB	44x 32GB	1408GB	RDIMMs: A1-A10, B1-B10, C1-C12, D1-D12 NVDIMM-Ns: A11-A12, B11-B12
6x	96GB	24x 16GB	384GB	RDIMMs: A1-A6, B1-B6, C1-C6, D1-D6 NVDIMM-Ns: A7-A9, B7-B9
6x	96GB	24x 32GB	768GB	RDIMMs: A1-A6, B1-B6, C1-C6, D1-D6 NVDIMM-Ns: A7-A9, B7-B9

NVDIMM-N	NVDIMM-N Capacity	RDIMMs	RDIMM Capacity	DIMM Population Locations
6x	96GB	42x 32GB	1344GB	RDIMMs: A1-A9, B1-B9, C1-C12, D1-D12 NVDIMM-Ns: A10-A12, B10-B12
12x	192GB	24x 16GB	384GB	RDIMMs: A1-A6, B1-B6, C1-C6, D1-D6 NVDIMM-Ns: A7-A12, B7-B12
12x	192GB	24x 32GB	768GB	RDIMMs: A1-A6, B1-B6, C1-C6, D1-D6 NVDIMM-Ns: A7-A12, B7-B12
12x	192GB	36x 32GB	1152GB	RDIMMs: A1-A6, B1-B6, C1-C12, D1-D12 NVDIMM-Ns: A7-A12, B7-B12

NOTE:

- 1 NVDIMM-Ns are not supported in CPU Sockets 3 and 4 memory slots
- 2 While other configurations may work, they have not been fully validated and are not currently supported by DellEMC.

Modular Chassis Hardware Configuration

The MX7000 Modular chassis currently offers two different servers that support NVDIMM-N: MX740c (2-socket) and MX840c (4-socket). In order for a power loss condition to be detected, the chassis must have at least one Management Module installed. While not a requirement for NVDIMM-N support, dual Management Modules provide additional redundancy and robustness to the NVDIMM-N solution.

If a user is operating a chassis with NVDIMM-N equipped servers and a Management Module firmware update is performed, it is recommended that the chassis be equipped with two Management Modules. If the chassis has a single Management Module, the NVDIMM-Ns in the server may be placed in read-only mode during the Management Module upgrade.

NVDIMM-N Module Details

Figure 6 is a picture of a typical NVDIMM-N, and it shows the location of status LEDs on the NVDIMM-N. Table 3 describes the NVDIMM-Ns LED behavior during different modes of operation.

The DIMM has the following key features:

- JEDEC compliant DDR4 288-pin NVDIMM-N
- 16GB RDIMM, PC4-2666
- SLC Flash for backup
- In-system health monitoring and alert
- On-board I²C temperature sensor with integrated SPD EEPROM

LEDs



Figure 7. NVDIMM-N

Table 5. : NVDIMM-N LED Behavior

NVDIMM-N Operation	Green LED	Blue LED	Amber LED
Powered Off	Off	Off	Off
NVDIMM-N going through power-on and initialization	On	Off	On
NVDIMM-N is operating normally	On	Slow Blink (every 15 seconds)	Off
Save of Restore Operation in progress	On	Fast Blink	Off

Battery

A battery is required to provide backup power to copy contents from DRAM to flash. Although JEDEC based NVDIMM-Ns can utilize Super Caps as backup power, DellEMC's battery is a centralized power solution that provides a more compact, reliable, and integrated power source. Power delivery is integrated into the system board, and does not require individual cables to each NVDIMM-N that is typical of Super Cap based solutions.

On first detection of an power loss, the server activates the battery and triggers the NVDIMM-N to save its DRAM contents to flash. The Save process takes approximately a minute to complete in the T640, R640, R740/R740XD, R840, R940 and R940xa servers and 140 seconds to complete in the MX740c and MX840c servers. During this time the battery is required to provide power for backup. The battery supplies power only to server sub-systems that must be powered for the save to complete. All other subsystems are powered down or placed into a low power state to conserve battery life. After the save completes, the battery is deactivated, and the server is allowed to power down completely. In modular chassis, the batteries also provide power to the system fans to provide cooling during the NVDIMM save after an power loss event.

Table 6. Battery LED Behavior

Battery Operation	Green LED
Powered Off	Off
Idle (not charging and not enabled while the server is on)	Off

Battery Operation	Green LED
Maintenance Charging	On
Critical Charging (battery capacity is below minimum required for NVDIMM-N Save, server is powered on)	Blinking
Battery Enabled and Discharging (Power loss detected, Battery supplying power for NVDIMM-N Save operation)	Blinking

If a battery is not installed or not fully charged, BIOS will restore the NVDIMM-N DRAM contents from its onboard flash, but BIOS will keep the NVDIMM-N in Read-Only mode. This allows OS and applications to have access to NVDIMM-N data, and at the same time protect the NVDIMM-N from potential data loss due to power loss.

The NVDIMM-N battery is not hot swappable. To prevent data loss and potential damage to your system, ensure that the system, LEDs on the system, LEDs on NVDIMM-N, and LEDs on NVDIMM-N battery are turned off before installing or removing the NVDIMM-N Battery.

Figure 7 below shows the system board connections for the NVDIMM-N Battery cables in the R740/R740XD (Items 36 and 38). Figure 8 shows how to install and mount the battery in the R740/R740XD. Please refer to each server's Installation and Service Manual for detailed instructions on how to install or remove the battery.

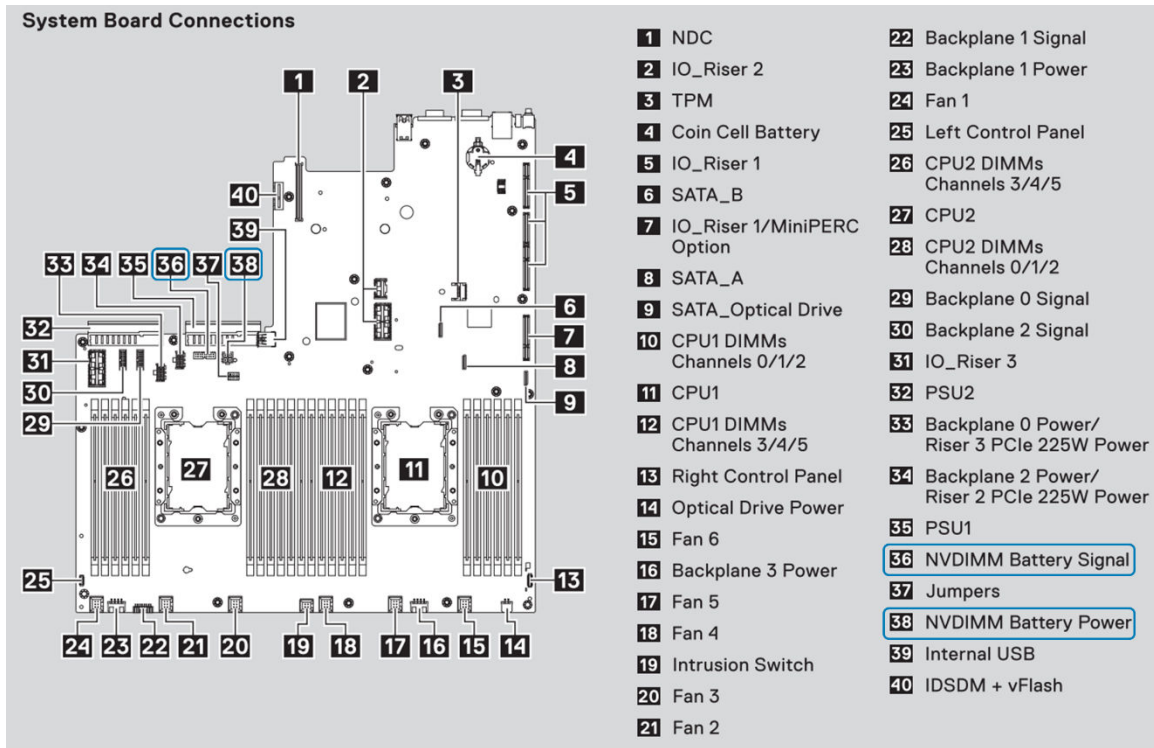


Figure 8. R740/R740XD System Board Connections

NOTE: Connector locations will be different for each server. Please refer to your particular server's Installation and Service Manual for more information.




Figure 9. R740 Battery Installation Instructions

① **NOTE:** Battery installation locations will be different for each server. Please refer to your particular Server's Installation and Service Manual for instructions.

Minimum Platform Firmware Versions

For NVDIMM-N modules to be functional on PowerEdge servers, the minimum platform firmware versions are required to be as follows:

- BIOS: 1.1.7
- iDRAC: 3.00.00.00

① **NOTE:** Certain operating systems require specific minimum versions of BIOS, NVDIMM-N and/or iDRAC firmware. Please refer to the individual sections of the OS for more details.

Topics:

- BIOS Configuration Settings for NVDIMM-N
- BIOS Error Messages

BIOS Configuration Settings for NVDIMM-N

This section focuses only on the BIOS setup options that affect NVDIMM-N operation. For a description of all setup options, please refer to each server's Installation and Service Manual. Persistent Memory BIOS settings are configurable by going into BIOS System Setup. Press F2 at the BIOS screen below to enter BIOS System Setup.



Figure 10. BIOS screen

Select `System BIOS > Memory Settings` to see the following options:




Figure 11. Memory Settings

Node Interleaving Specifies if Non-Uniform Memory Architecture (NUMA) is supported. If this field is set to Enabled, memory interleaving is supported if a symmetric memory configuration is installed. If the field is set to Disabled, the system supports NUMA (asymmetric) memory configurations. This option is set to Disabled by default. Dell EMC recommends that this option be Disabled for NVDIMM-N configurations.

⚠ WARNING: When Persistent Memory Interleaving is enabled in the Persistent Memory System Setup page, Node Interleaving setting also applies to Persistent Memory Interleaving.

Select [Persistent Memory](#) to configure NVDIMM-N setup options:

Persistent Memory Scrubbing Set Persistent Memory Scrubbing mode.

Auto: System automatically scrubs persistent memory during POST when multibit errors have been detected.

One Shot: System will scrub persistent memory during POST on the entire persistent memory range once. Upon next boot, system will go back to "Auto" persistent memory scrub mode.

Enable: System will scrub persistent memory during POST on the entire persistent memory range on everyboot.

ⓘ NOTE: Scrubbing persistent memory on the entire persistent memory range could take more than 60 minutes during system POST depending on system memory population.

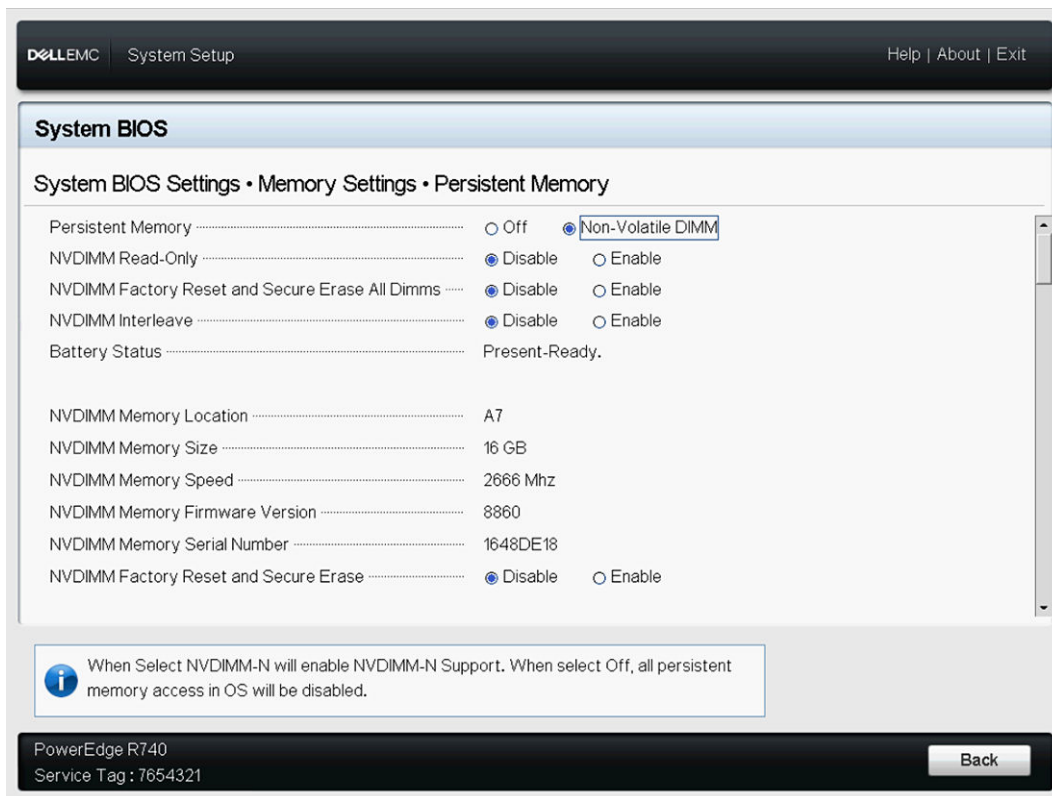



Figure 12. Persistent Memory screen

The following table describes each option that is available in the BIOS setup screen.

Table 7. BIOS setup screen

Option	Description
Node Interleaving	<p>Specifies if Non-Uniform Memory Architecture (NUMA) is supported. If this field is set to Enabled, memory interleaving is supported if a symmetric memory configuration is installed. If the field is set to Disabled, the system supports NUMA (asymmetric) memory configurations. This option is set to Disabled by default. Dell EMC recommends that this option be Disabled for NVDIMM-N configurations.</p> <p> WARNING: When Persistent Memory Interleaving is enabled in the Persistent Memory System Setup page, Node Interleaving setting also applies to Persistent Memory Interleaving.</p> <p>This option is set to Disabled by default.</p>
Persistent Memory	<p>Enables or disables NVDIMM-N operation. If this option is set to Off, NVDIMM-Ns are not mapped to System Physical Address (SPA) and are not presented to the OS. Data that has already persisted to NVDIMM flash are unaffected and are available for OS utilization when this option is returned to Non-Volatile DIMM.</p> <p>This option is set to Non-Volatile DIMM by default.</p>
NVDIMM-N Read-Only	<p>When set to Enable, this forces NVDIMMs to be read-only. NVDIMM-N data is accessible to the OS but locked from any modifications. Read-only is intended for debug or maintenance operations.</p> <p>This option is set to Disable by default.</p>

Option	Description
NVDIMM-N Factory Reset and Secure Erase All Dimms	Enables resetting the NVDIMM-N to factory defaults and clearing data on the NVDIMM-N. If set to Enable , all data on all NVDIMM-N will be erased. This option does not clear the NVDIMM's health information, and should be used to remove data and/or repurpose the NVDIMM-N. The reset and data erasure will occur when BIOS confirms that Setting Changes should be saved. This option is set to Disable by default.
NVDIMM-N Interleave	Enables or disables interleaving on NVDIMM-N. When Enabled, NVDIMM-N interleaving will follow the same interleaving policy that applies to RDIMMs. Volatile RDIMM interleaving policy is not affected by this option. RDIMM system memory and NVDIMM-N persistent memory will remain as two distinct memory regions. This option is set to Disable by default.
Battery Status	Indicates if the NVDIMM-N battery is ready. Battery Status can display one of the following states: <ul style="list-style-type: none"> • Present-Ready • Present-Offline • Not-Present

The following settings are applicable to each NVDIMM-N present in the server

Table 8. NVDIMM-N

Option	Description
NVDIMM-N Memory Location	Specifies the memory slot location of the NVDIMM-N.
NVDIMM-N Memory Size	Specifies information on the capacity of the NVDIMM-N.
NVDIMM-N Memory Speed	Specifies information on the operating speed of the NVDIMM-N.
NVDIMM-N Memory Firmware version	Specifies information on the current firmware version on the NVDIMM-N.
NVDIMM-N Memory Serial Number	Specifies information on the serial number of the NVDIMM-N.
Remaining Rated Write Endurance [%]	Specifies information on the remaining NVDIMM-N flash life-time in terms of percentage.
NVDIMM-N Factory Reset and Secure Erase	Enables resetting and clearing data on the specific NVDIMM-N, and results in loss of data on that specific NVDIMM-N.

Please note that if any option was changed by mistake, press ESC to exit each setup screen, and select **No** when prompted at the screen below to discard all changes. Alternatively, options can be restored to the correct settings before responding **Yes** to the prompt below.

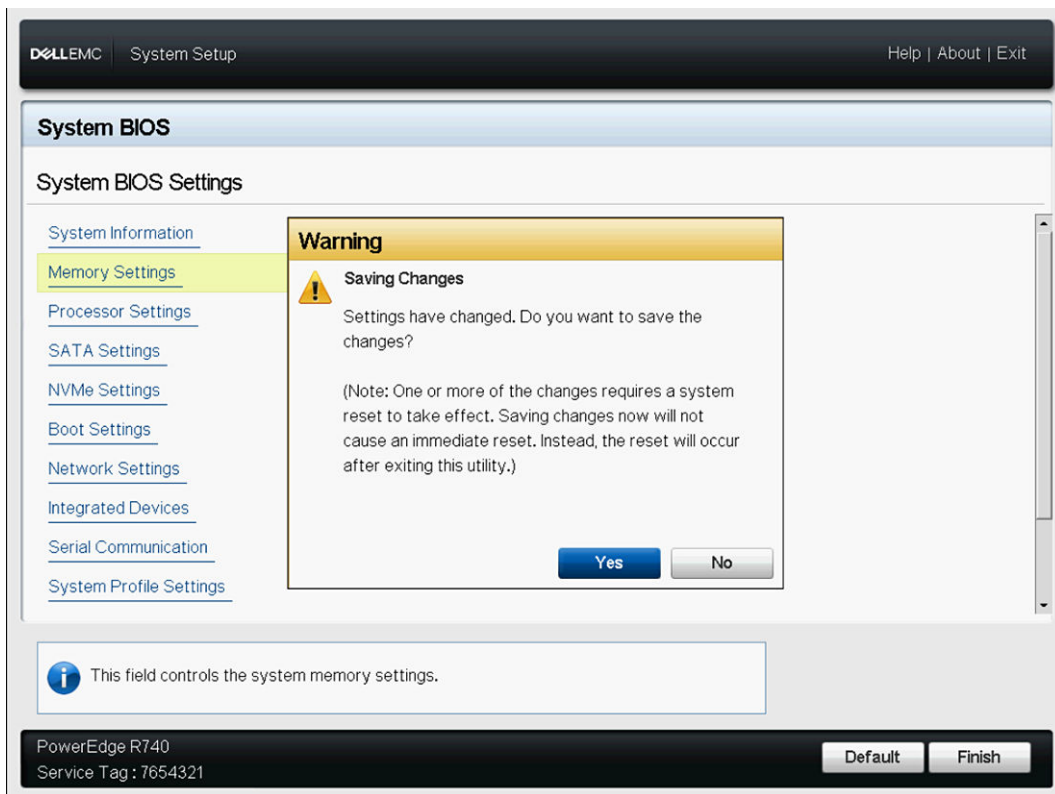


Figure 13. System BIOS Settings screen

BIOS Error Messages

When BIOS detects an NVDIMM-N related error during POST, BIOS displays an F1/F2 prompt and a corresponding error message. Multiple messages will appear when multiple errors were detected. BIOS will also log an event for each error in the Server System Event Log (SEL) and Life Cycle Log (LCL). Please refer to the JEDEC JESD245B spec for more information on each NVDIMM-N related failure. NVDIMM-N [Location] is the memory slot location as shown in Figure 5.

UEFI0302* NVDIMM-N Not Ready or Not Responding on NVDIMM-N located at [Location].

This NVDIMM-N module will be set to read only mode. Remove input power to the system, reseal the NVDIMM-N module, and restart the server. If the issue persists, replace the faulty memory module identified in the message.

UEFI0302* Save Error on NVDIMM-N located at [Location].

This NVDIMM-N module will be set to read only mode. Remove input power to the system, reseal the NVDIMM-N module, and restart the server. If the issue persists, replace the faulty memory module identified in the message.

UEFI0302* Restore Error on NVDIMM-N located at [Location].

This NVDIMM-N module will be set to read only mode. Remove input power to the system, reseal the NVDIMM-N module, and restart the server. If the issue persists, replace the faulty memory module identified in the message.

UEFI0302 Arm Error on NVDIMM-N located at [Location].

This NVDIMM-N module will be set to read only mode. Remove input power to the system, reseal the NVDIMM-N module, and restart the server. If the issue persists, replace the faulty memory module identified in the message.

UEFI0302* Erase Error on NVDIMM-N located at [Location].

This NVDIMM-N module will be set to read only mode. Remove input power to the system, reseat the NVDIMM-N module, and restart the server. If the issue persists, replace the faulty memory module identified in the message.

UEFI0302 Factory Default Error on NVDIMM-N located at [Location].

This NVDIMM-N module will be set to read only mode. Remove input power to the system, reseat the NVDIMM-N module, and restart the server. If the issue persists, replace the faulty memory module identified in the message.

UEFI0302 Set Event Notification Error on NVDIMM-N located at [Location].

This NVDIMM-N module will be set to read only mode. Remove input power to the system, reseat the NVDIMM-N module, and restart the server. If the issue persists, replace the faulty memory module identified in the message.

UEFI0302 Set Energy Source Policy Error on NVDIMM-N located at [Location].

This NVDIMM-N module will be set to read only mode. Remove input power to the system, reseat the NVDIMM-N module, and restart the server. If the issue persists, replace the faulty memory module identified in the message.

UEFI0302 Firmware Operation Error on NVDIMM-N located at [Location].

This NVDIMM-N module will be set to read only mode. Remove input power to the system, reseat the NVDIMM-N module, and restart the server. If the issue persists, replace the faulty memory module identified in the message.

UEFI0302* Persistency Lost Error on NVDIMM-N located at [Location].

This NVDIMM-N module will be set to read only mode. Remove input power to the system, reseat the NVDIMM-N module, and restart the server. If the issue persists, replace the faulty memory module identified in the message.

UEFI0303 Lifetime Percentage <= 5% on NVDIMM-N located at [Location]. The identified NVDIMM-N module is approaching the end of its functional life.

Please plan to replace the NVDIMM-N module during the next scheduled maintenance. For more information, please refer to the server Installation and Service Manual.

UEFI0304 Unable to enable the NVDIMM-N Save operation because Server is not Armed. All NVDIMM-N modules will be set to write-protect mode. Please verify that the battery is operational, and that the server has sufficient power to support this configuration.

UEFI0304 Unable to enable the NVDIMM-N Save operation because Battery is not present. All NVDIMM-N modules will be set to write-protect mode. Please verify that the battery is operational, and that the server has sufficient power to support this configuration.

UEFI0304 Unable to enable the NVDIMM-N Save operation because Battery is not ready. All NVDIMM-N modules will be set to write-protect mode. Please verify that the battery is operational, and that the server has sufficient power to support this configuration.

UEFI0304 Unable to enable the NVDIMM-N Save operation because Server is not Armed. All NVDIMM-N modules will be set to write-protect mode. Please verify that the battery is operational, and that the server has sufficient power to support this configuration, and that NVDIMM Write-Protect mode has not be enabled in BIOS Setup.

UEFI0340 NVDIMM-N Running from Backup Firmware Image at [Location].

Update the NVDIMM-N to the latest firmware. If the issue persists, replace the faulty memory module identified in the message.

UEFI0341* NVDIMM-N Lifetime Percentage <= 1% on NVDIMM-N located at [Location].

This NVDIMM-N module will be set to read only mode. Replace the NVDIMM-N module.

The following two BIOS messages are modular specific and will only be seen on an NVIDMM equipped server in the event that there are no management modules installed in the modular chassis or the installed management modules are not capable of detecting an power loss event.

Table 9. BIOS Error messages with event ID's and event messages

ID	Event Message	
UEFI0364*	Unable to detect chassis power loss, memory may not be persistent.	Ensure there are two functional chassis management modules. If there is only a single chassis management module, only perform a firmware update on it when server modules with persistent memory are powered off.
UEFI0365	Chassis is able to detect power loss, memory is persistent.	No response action required.

* Errors marked with an asterisk may result in data loss.

iDRAC NVDIMM-N Management

iDRAC Graphical User Interface

The image below shows the iDRAC Web GUI Dashboard when remotely managing the server.

The screenshot displays the iDRAC Web GUI Dashboard for an Integrated Dell Remote Access Controller 9 Enterprise. The dashboard is organized into several sections:

- System Health:** A grid of status indicators for various components, all showing green checkmarks: Batteries, Voltages, CPUs, Miscellaneous, Cooling, Intrusion, Memory, and Power Supplies.
- System Information:** A table of system details:

Power State	ON
Model	PowerEdge R740
Host Name	WIN-LS1VT2FMJ4S
Operating System	
Operating System Version	10.0
Service Tag	7654321
BIOS Version	1.0.8
iDRAC Firmware Version	3.00.00.00
iDRAC MAC Address	84:7b:eb:f4:f5:4c
- Virtual Console:** A window showing a remote session with a "Launch Virtual Console" button.
- Recent Logs:** A table of system events:

Severity	Description	Date and Time
✓	An OS graceful shut-down occurred.	Wed 02 Aug 2017 09:31:55
✓	OEM software event.	Wed 02 Aug 2017 09:31:55
✓	C: boot completed.	Wed 02 Aug 2017 09:28:00
- Notes:** A section indicating "There are no work notes to be displayed."

Figure 14. iDRAC Graphical User Interface

NVDIMM-N Status

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

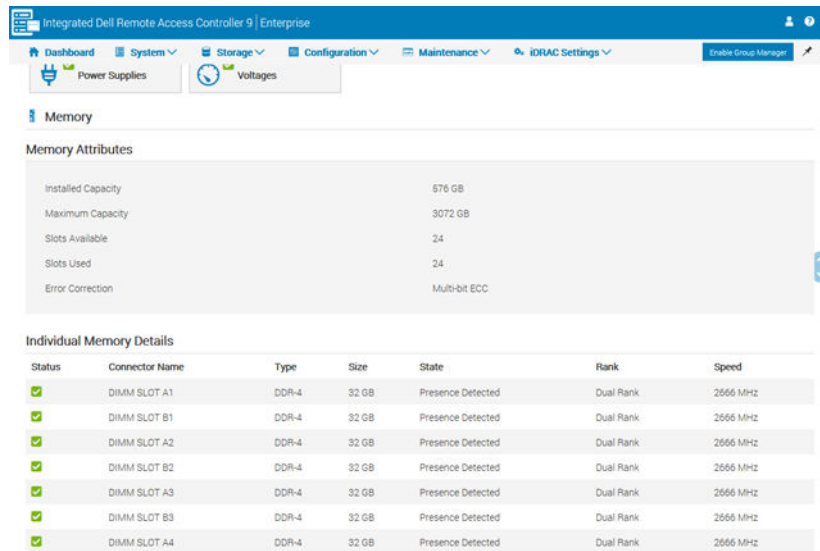


Figure 15. NVDIMM-N Status

NOTE:

- 1 All NVDIMM-N errors will be reported to the OS and logged in the server System Event Log. NVDIMM-N Health Status currently only reflects Correctable Error Threshold Exceeded and Uncorrectable Error status on the NVDIMM-N. Other errors are reported to OS and logged, but are not be reflected in the iDRAC/OM NVDIMM-N Health Status.
- 2 NVDIMM-N DIMMs are currently reported as DDR4 16GB Single-Rank 2666 DIMMs in the Memory Details page. DellEMC standard DDR4 16GB RDIMMs are reported as DDR4 16GB Dual-Rank 2666 DIMMs.

BBU Status

Select the Batteries link on the Dashboard to get more information for the NVDIMM-N Battery health. iDRAC will display the NVDIMM-N Battery status only after BIOS has completed the system inventory (at end of POST), and it has detected that NVDIMM-Ns are installed in the server. The Battery status will not be reported if NVDIMM-Ns are not installed.




Figure 16. BBU Status

Log Messaging Errata

The following errata affects the messaging in the System Event Log:

- When UEFI0340 is logged in the Lifecycle controller log, the System Event Log and Lifecycle controller logs can have entries with the following message “An unsupported event occurred.” This message can be ignored and shall be fixed in a future iDRAC release.
- When UEFI0302 is logged for an Erase Error in the Lifecycle Controller Log, the System Event Log and Lifecycle controller logs can have entries with the following message “An unsupported event occurred.” This message can be ignored and shall be fixed in a future iDRAC release.

Remote Management

When a server with NVDIMM-Ns installed is powered down, the server will initiate a Save to backup NVDIMM-N DRAM data to its onboard flash. The Save process takes approximately 1 minute to complete with 16GB NVDIMM-Ns. During this time the server will be in a reset state, and any remote power on requests from iDRAC will be queued to proceed after the Save is complete. Remote power on requests may appear to be unresponsive while an NVDIMM-N Save is in progress.

NVDIMM-N Error Reporting

The table below provides a list of possible error events that can be logged to the server System Event Log. The table also provides the recommended action for each error.

Table 10. NVDIMM-N Error Reporting

ID	Event Message	Recommended Action
MEM9020	The Non-Volatile Dual In-line Memory Module (NVDIMM) device in the slot [location] is about to reach the end of supported life duration.	During the next scheduled maintenance, replace the Non-Volatile Dual In-line Memory Module (NVDIMM) device. For more information about the usage of NVDIMMs, see the system Owner's Manual on the support site.
MEM9030	The Non-Volatile Dual In-line Memory Module (NVDIMM) device in the slot [location] is not responding and is disabled.	Remove and reinstall the Non-Volatile Dual In-line Memory Module (NVDIMM) device. If the issue persists, contact your service provider. For information about removing and reinstalling the NVDIMM, see the system Owner's Manual on the support site.
MEM9031	The Non-Volatile Dual In-line Memory Module (NVDIMM) device in the slot [location] is unable to save the data during the previous system shutdown operation or power loss.	Remove and reinstall the Non-Volatile Dual In-line Memory Module (NVDIMM) device. If the issue persists, contact your service provider. For information about removing and reinstalling the NVDIMM, see the system Owner's Manual on the support site.
MEM9032	The Non-Volatile Dual In-line Memory Module (NVDIMM) device in the slot [location] is unable to restore the data that was saved in the previous 'save' operation.	Remove and reinstall the Non-Volatile Dual In-line Memory Module (NVDIMM) device. If the issue persists, contact your service provider. For information about removing and reinstalling the NVDIMM, see the system Owner's Manual on the support site.
MEM9033	An unsupported Non-Volatile Dual In-line Memory Module (NVDIMM) device is of unsupported configuration and unable to operate as currently configured.	Review the memory configuration and make sure the configuration is as per memory rules defined in the system Owner's Manual on the support site.
MEM9034	The Non-Volatile Dual In-line Memory Module (NVDIMM) device in the slot [location] is not responding.	Remove and reinstall the Non-Volatile Dual In-line Memory Module (NVDIMM) device. If the issue persists, contact your service provider. For information about removing and reinstalling the NVDIMM, see the system Owner's Manual on the support site.
MEM9035	The Non-Volatile Dual In-line Memory Module (NVDIMM) device in the slot [location] cannot be configured to save data during a power loss because of an issue in the NVDIMM.	Remove and reinstall the Non-Volatile Dual In-line Memory Module (NVDIMM) device. If the issue persists, contact your service provider. For information about removing and reinstalling the NVDIMM, see the system Owner's Manual on the support site.
MEM9036	The Non-Volatile Dual In-line Memory Module (NVDIMM) devices are placed in write-protect mode because the system may not provide sufficient power to save data in case of power loss.	Review any preceding power or battery entries in the log and resolve the issues. Refer to your documentation for details on how to view the log.
MEM9037	The Non-Volatile Dual In-line Memory Module (NVDIMM) device in the slot [location] has reached the end of supported	Replace the Non-Volatile Dual In-line Memory Module (NVDIMM) device. For information about removing and reinstalling

ID	Event Message	Recommended Action
	life duration and is placed in write-protect mode.	the NVDIMM, see the system Owner's Manual on the support site.
MEM9038	The Non-Volatile Dual In-line Memory Module (NVDIMM) device in the slot [location] has lost persistency and is placed in write-protect mode.	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.
MEM9050	The Non-Volatile Dual In-line Memory Module (NVDIMM) device in the slot [location] has regained persistency and is available for use.	No response action required.
BAT0015	The NVDIMM battery is low.	Leave the server powered on to allow the NVDIMM-N Battery to charge. If the issue persists, contact your service provider. For information about removing and reinstalling the NVDIMM, see the system Owner's Manual on the support site.
BAT0017	The NVDIMM battery has failed.	Remove and reinstall the NVDIMM-N Battery. If the issue persists, contact your service provider. For information about removing and reinstalling the NVDIMM, see the system Owner's Manual on the support site.
BAT0019	The NVDIMM battery is absent.	Remove and reinstall the NVDIMM-N Battery. If the issue persists, contact your service provider. For information about removing and reinstalling the NVDIMM, see the system Owner's Manual on the support site.

Server Behavior with NVDIMM-Ns

The server behavior changes slightly when NVDIMM-Ns are installed. This section covers differences that can be observed as the server shuts down and boots up. This section will also describe scenarios where the server will automatically shutdown to ensure that NVDIMM-N DRAM data is securely stored to flash.

Topics:

- [Shutdown](#)
- [Boot](#)
- [Automatic Shutdown and Save](#)

Shutdown

In order for NVDIMM-N data to persist across server power cycles, the NVDIMM-N must save its DRAM data to on boardflash. The Save is triggered by events that would result in power loss to the NVDIMM-N. In MX7000, the entire chassis is shutdown when there is a loss of power and at least one NVDIMM-N equipped server is powered on and armed when power loss occurs. During the Save, the server maintains power to the NVDIMM-Ns, and a battery is used to provide back-up power in case of an power loss. The server remains in a low-power reset state, and the power button on the server's front control panel blinks green once per second to indicate that a save is in progress. The Save takes approximately 1 minute to complete with the 16GB NVDIMM-N DIMMs (140 seconds for MX740c and MX840c servers).



Figure 17. R740 Control Panel

Once the Save completes, the server continues its normal power down sequencing. If the Power Button is pressed (to request server power on) while a SAVE in progress, the server will queue the power-on to occur after the Save.

- ① **NOTE:** When NVDIMMs are installed, verify that the Server is completely powered down before attempting to perform any service actions. This includes removing a server from a modular chassis after initiating a shutdown procedure. If a server equipped with NVDIMMs is removed from a modular chassis prior to the save completing, the save will be interrupted and data loss will occur. The server is completely powered down when all LEDs including NVDIMM-N and NVDIMM-N Battery LEDs are off.
- ① **NOTE:** If a Modular server supporting NVDIMM-N is powered on, the entire chassis will be shutdown when there is a full loss of power or partial loss below minimum chassis power budget requirements.

Users cannot see these LEDs without removing the cover which they should not do during operation. User action is to wait until the LED on front of the system stops blinking.

Boot

Server BIOS restores NVDIMM-N DRAM data from its onboard Flash during boot time. BIOS verifies that the NVDIMM-N Battery is installed and has sufficient charge for a Save event in case of a power loss. BIOS also verifies that the installed server Power Supplies are sufficiently sized for the server configuration. This is to ensure that after a power loss, the PSUs are able to provide enough power to hold up the server until Battery power takes over. Once BIOS verifies that both the Battery and PSUs have enough power to support NVDIMM-Ns, BIOS arms the NVDIMM-Ns and allows the server to boot to OS.

If BIOS detects any issues, BIOS will display an F1/F2 error message on screen and lock the NVDIMM-N in Read-Only mode. OSES will still be able to read NVDIMM-N data at this time, but any changes will be discarded. To restore the NVDIMM-N to normal operation, the user must correct the issue and reboot the server. Please refer to Section BIOS Error Messages 4.2 BIOS Error Messages for more details.

Automatic Shutdown and Save

A primary goal for DellEMC's NVDIMM-N Persistent Memory Solution is to protect customer data. A key ingredient in the data protection is the server's Power Supply. The PSU(s) must have sufficient hold-up power for the Save process to start and for the server to transition to the NVDIMM-N Backup Battery.

Any changes to the server hardware configuration due to hot add/remove of devices or PSU failures will result in a server power budget re-evaluation. If in the case where the power budget exceeds the PSU's ability to hold-up power, the server will immediately shutdown to trigger a Save on the NVDIMM-N and protect customer data to flash. The server will reboot. If the problem persists, BIOS will place the NVDIMM-N into Read-Only mode before booting to the OS.

In a modular chassis, since power is managed at the chassis level, a power loss event or removal of PSUs resulting in power demand exceeding available capacity will trigger an entire chassis shutdown if there are NVDIMM equipped sleds installed and powered on. If the chassis is power constrained, additional components will not be allowed to turn on if the power requirements of those components will push the power consumption of the chassis above the available power. Additionally, if a user elects to perform a "Virtual Reseat" of an NVDIMM equipped sled, the targeted sled will trigger a save on the NVDIMM-N, but not affect the whole chassis.

DIMM Configuration Changes

DellEMC recommends that NVDIMM-N data contents be backed up to external storage before making any changes to the server memory configuration. This applies to both NVDIMM-Ns and RDIMMs. Due to memory Error Correction (ECC) algorithms that are unique to each memory slot and memory configuration, NVDIMM-Ns may generate errors after a memory configuration change. If the server DIMM configuration changes, DellEMC recommends that the NVDIMM-N be erased and reset to factory defaults in BIOS Setup prior to utilization.

In the scenario where the DIMMs (both RDIMMs and NVDIMM-Ns) must be moved or the system board must be replaced, DellEMC recommends that all DIMMs be restored to their exact same locations as in the original server. DellEMC also recommends that BIOS NVDIMM settings are identical.

If an RDIMM must be replaced due to failure, DellEMC recommends that the RDIMM be replaced with an RDIMM of the same speed, organization, and capacity.

Topics:

- [BIOS Requirements](#)
- [Set Up](#)
- [Windows Drivers](#)
- [Storage Class Memory in Windows Server 2016](#)
- [Storage Class Memory in Windows Server 2019](#)
- [Windows Errata](#)

BIOS Requirements

Both Windows 2016 and 2019 require the minimum BIOS version to be at 1.6.13 so that NVDIMM-N modules can be used without any issues.

Set Up

Windows Server 2016 and 2019 natively supports NVDIMM-N devices allowing application and service access to extremely high performance and low latency NVDIMM-N devices on memory bus.

From BIOS Setup, enable the Persistence Memory setting, which is also the default BIOS settings when NVDIMM-N devices are plugged in. Windows Server 2016 and Windows 10 Anniversary Edition currently do not support interleaved sets of NVDIMM-Ns and the NVDIMM-N Interleave setting should be set to Disable, which is the default setting.

Windows Server 2019 supports NVDIMM-N interleaving and this option can be set to Enable.

[2017-07 Cumulative Update for Windows Server 2016 for x64-based Systems \(KB4025334\)](#) is required for NVDIMM-N to function correctly in Windows Server 2016.

Windows Drivers

Windows refers “Storage Class Memory (SCM) “as persistent storage that is directly installed on the memory bus. NVDIMM-N is a form of Storage class memory.

DellEMC PowerEdge platforms implement NVDIMM-N root device using PNP ID “ACPI0012”, this root device contains individual NVDIMM-N devices. During boot process windows loads “Storage Class memory bus driver (Scmbus.sys) for this root device. Scmbus.sys driver enumerates each NVDIMM-N device in the system and loads scmdisk0101.sys as functional driver for each NVDIMM-N device.

DellEMC NVDIMM-N in Windows Server 2016 and 2019 environment can support both block mode and direct access mode (DAX) mode. The following two videos describe the difference between these two modes and how use each mode.

- [Using Non-volatile Memory \(NVDIMM-N\) as Block Storage](#)
<https://channel9.msdn.com/Events/Build/2016/P466>
- [Using Non-volatile Memory \(NVDIMM-N\) as Byte-Addressable Storage](#)

Storage Class Memory in Windows Server 2016

Device manager

Below picture describes the windows device manager view of NVDIMM-N root device and NVDIMM-N disk instances in Windows Server 2016.

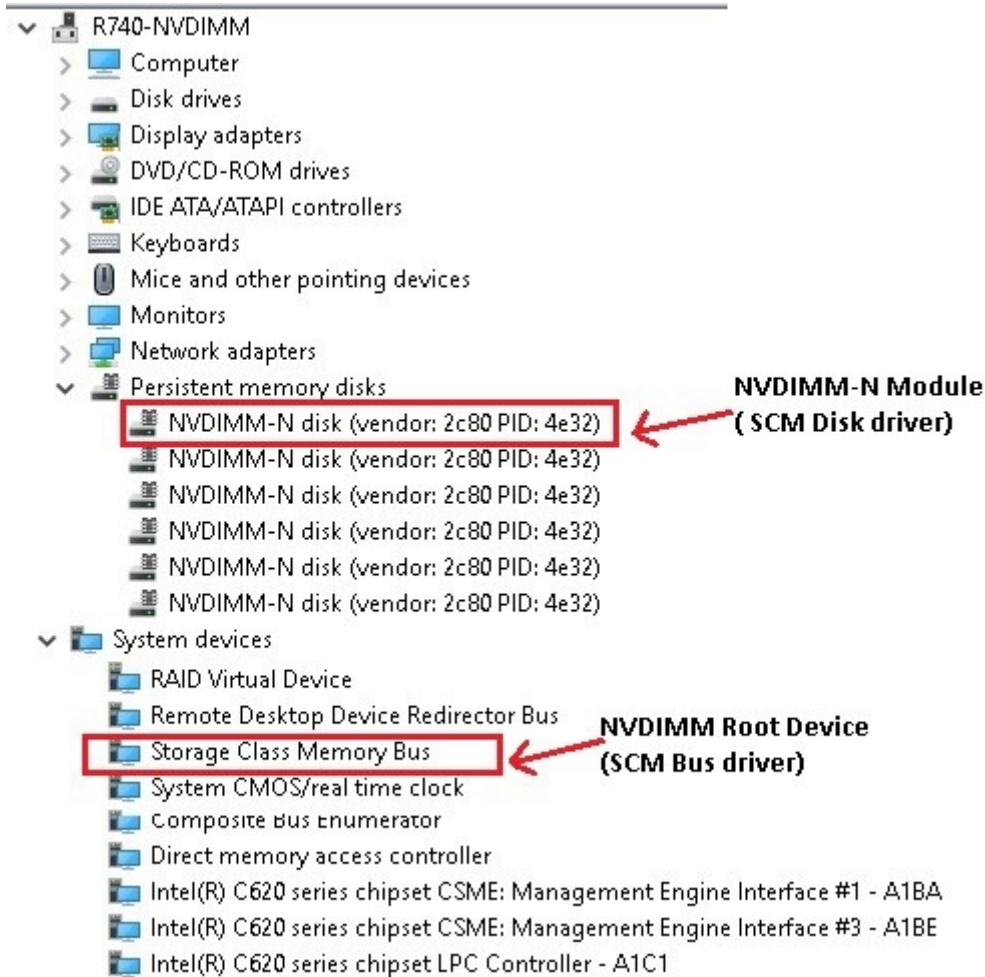


Figure 18. Windows device manager view of NVDIMM-N root device and NVDIMM-N disk instances

Identifying the right NVDIMM-N disks

Windows PowerShell and NVDIMM-N disk properties GUI in device manager provides information that will can be used to uniquely identify the physical NVDIMM-N module.

Using PowerShell:

```
PS C:\Users\Administrator> Get-PhysicalDisk | ? bustype -eq SCM | ft friendlyname, UniqueId, PhysicalLocation, SerialNumber, FirmwareVersion, BusType
friendlyname      UniqueId          PhysicalLocation  SerialNumber      FirmwareVersion  BusType
-----
Vendor 2c80 PID 4e32 {ab2b8927-cca1-e4c4-ae65-ecd3cb60d255} DIMM Socket 0 : Slot 1 802c-0f-1711-1648f4c5 8860 SCM
Vendor 2c80 PID 4e32 {859b53cd-e759-ff8f-03a3-abc6883c00a8} DIMM Socket 0 : Slot 17 802c-0f-1642-146f4660 8860 SCM
Vendor 2c80 PID 4e32 {d3fdd8ad-b4b4-bff2-b4e6-5db618bed75b} DIMM Socket 0 : Slot 33 802c-0f-1642-146f470f 8860 SCM
Vendor 2c80 PID 4e32 {96dff4a9-e2b5-f014-5e53-ffeda701db72} DIMM Socket 0 : Slot 49 802c-0f-1711-164af113 8860 SCM
Vendor 2c80 PID 4e32 {de8dc05f-71a1-607f-c956-ab9ce1c585e3} DIMM Socket 0 : Slot 65 802c-0f-1711-16493233 8860 SCM
Vendor 2c80 PID 4e32 {e9127230-72d8-750c-339a-632dbc05af80} DIMM Socket 0 : Slot 81 802c-0f-1642-146ea236 8860 SCM
Vendor 2c80 PID 4e32 {d7c2ed38-7c0f-f4f1-3f28-06d3f8d48237} DIMM Socket 1 : Slot 1 802c-0f-1711-1648f48f 8860 SCM
Vendor 2c80 PID 4e32 {aee9e2cc-1988-0770-57e1-ff3f879a9da0} DIMM Socket 1 : Slot 17 802c-0f-1711-1649373f 8860 SCM
Vendor 2c80 PID 4e32 {1e1dcdad-18b1-834c-f8c5-02dd6b325672} DIMM Socket 1 : Slot 33 802c-0f-1642-146ef6e7 8860 SCM
Vendor 2c80 PID 4e32 {09dca1ab-17ce-103b-264b-cd7403cddf69} DIMM Socket 1 : Slot 49 802c-0f-1711-1648d885 8860 SCM
Vendor 2c80 PID 4e32 {67e63f1f-cc0b-6257-1ce8-f96d60bb40eb} DIMM Socket 1 : Slot 65 802c-0f-1642-146eea1f 8860 SCM
Vendor 2c80 PID 4e32 {96ed7051-af45-bd96-2e88-8d7a98724af7} DIMM Socket 1 : Slot 81 802c-0f-1642-146ea274 8860 SCM
```

Figure 19. Identifying the right NVDIMM-N disks

The Serial Number for every NVDIMM-N is unique, and physical location values in PowerShell can be mapped to silk screen using the following table

Table 11. DIMM slot location

PowerShell PhysicalLocation	Dimm Slot Location
DIMM Socket 0:Slot 1	A7
DIMM Socket 0:Slot 17	A8
DIMM Socket 0:Slot 33	A9
DIMM Socket 0:Slot 49	A10
DIMM Socket 0:Slot 65	A11
DIMM Socket 0:Slot 81	A12
DIMM Socket 1:Slot 1	B7
DIMM Socket 1:Slot 17	B8
DIMM Socket 1:Slot 33	B9
DIMM Socket 1:Slot 49	B10
DIMM Socket 1:Slot 65	B11
DIMM Socket 1:Slot 81	B12

Using device manager GUI

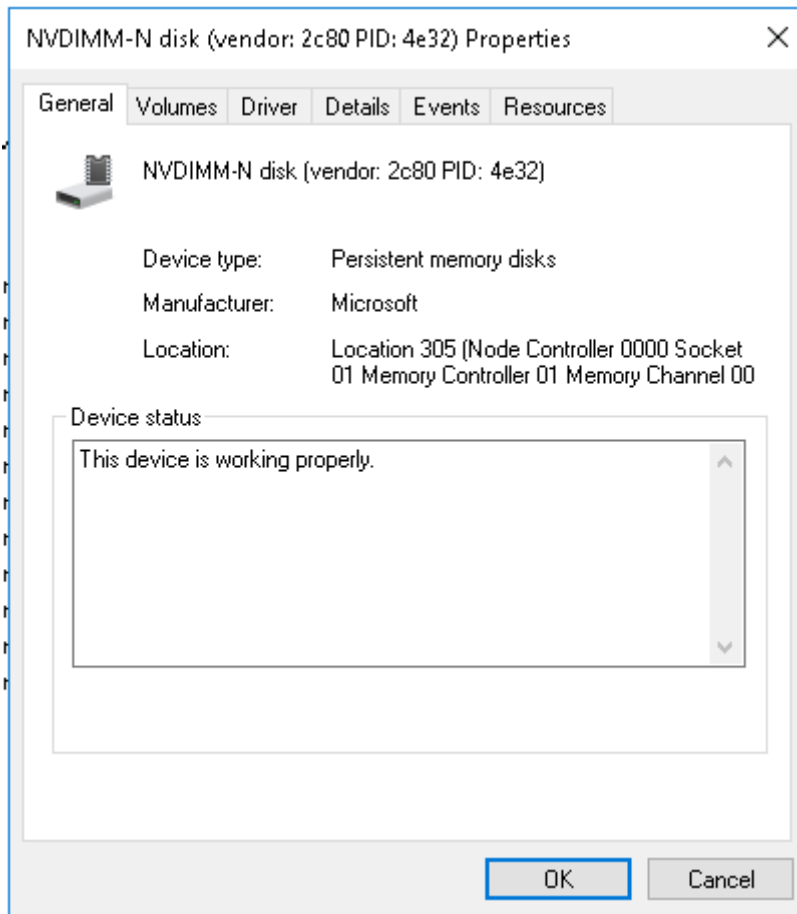


Figure 20. Using device manager GUI

Location information in the above can translated to Physical silk screen using following table

Table 12. DIMM Slot Location

Location	DIMM Slot Location
1	A7
17	A8
33	A9
49	A10
65	A11
81	A12
257	B7
273	B8

Location	DIMM Slot Location
289	B9
305	B10
321	B11
337	B12

NVDIMM-N health status and properties

NVDIMM-N health status can be queried using following PowerShell command.

```
PS C:\Users\Administrator> Get-PhysicalDisk | ? bustype -eq SCM | sort PhysicalLocation | ft friendlyname, PhysicalLocation, SerialNumber, HealthStatus, OperationalStatus, OperationalDetails
```

friendlyname	PhysicalLocation	SerialNumber	HealthStatus	OperationalStatus	OperationalDetails
Vendor 2c80 PID 4e32 DIMM Socket 0 : Slot 1	802c-Of-1711-1648f4e5	Unhealthy	Device Hardware Error (Lost Data Persistence, NVDIMM_N_Error)		
Vendor 2c80 PID 4e32 DIMM Socket 0 : Slot 17	802c-Of-1642-146f4660	Healthy	OK		
Vendor 2c80 PID 4e32 DIMM Socket 0 : Slot 33	802c-Of-1642-146f470f	Healthy	OK		
Vendor 2c80 PID 4e32 DIMM Socket 0 : Slot 49	802c-Of-1711-1648f113	Healthy	OK		
Vendor 2c80 PID 4e32 DIMM Socket 0 : Slot 65	802c-Of-1711-16493233	Healthy	OK		
Vendor 2c80 PID 4e32 DIMM Socket 0 : Slot 81	802c-Of-1642-146e3235	Healthy	OK		
Vendor 2c80 PID 4e32 DIMM Socket 1 : Slot 1	802c-Of-1711-1648f48f	Healthy	OK		
Vendor 2c80 PID 4e32 DIMM Socket 1 : Slot 17	802c-Of-1711-1649377f	Healthy	OK		
Vendor 2c80 PID 4e32 DIMM Socket 1 : Slot 33	802c-Of-1642-146ef6e7	Healthy	OK		
Vendor 2c80 PID 4e32 DIMM Socket 1 : Slot 49	802c-Of-1711-1648de85	Healthy	OK		
Vendor 2c80 PID 4e32 DIMM Socket 1 : Slot 65	802c-Of-1642-146eaa1f	Healthy	OK		
Vendor 2c80 PID 4e32 DIMM Socket 1 : Slot 81	802c-Of-1642-146ea274	Healthy	OK		

Figure 21. NVDIMM-N health status and properties

Windows native driver can handle different health events. For more details on the various health conditions, see the [windows documentation](https://docs.microsoft.com/en-us/windows-server/storage/storage-spaces/storage-class-memory-health) (<https://docs.microsoft.com/en-us/windows-server/storage/storage-spaces/storage-class-memory-health>)

Block Mode

Applications may access NVDIMM-Ns using standard ReadFile/WriteFile Block Interface APIs provided by the Windows SCM Disk Driver. This allows existing applications to immediately take advantage of the NVDIMM-N's high performance and low latency without requiring any modifications. Please note that since the SCM Disk Driver uses Load/Store commands to directly access DRAM on the NVDIMM-N, Disk Filter Drivers may not work. Please refer to the Microsoft video below for more information.

- Using Non-volatile Memory (NVDIMM-N) as Block Storage in Windows Server 2016

<https://channel9.msdn.com/Events/Build/2016/P466>

DAX Mode

Windows server 2016 allows NVDIMM-N Disk to be formatted in direct access mode. Following PowerShell command initialize NVDIMM-N disk, create volume and format volume in DAX mode. Fsutil.exe can be used to verify if volume in in DAX mode

```

PS C:\Users\Administrator> diskpart
PS C:\Users\Administrator> get-disk -Number 1 |ft -autosize
-----
Number  Friendly Name  Serial Number  HealthStatus  OperationalStatus  Total Size  Partition Style
-----
1         Vendor 2c80 PID 4e32  802c-0f-1711-1648f4c5  Healthy         Online         16 GB  RAW

PS C:\Users\Administrator> Initialize-Disk -PartitionStyle GPT
PS C:\Users\Administrator> Format-volume J: Format-NTFS -isDAX true
PS C:\Users\Administrator> New-Volume -FriendlyName DAX-VOL -DriveLetter J | Format-Volume -fileSystem NTFS -isDAX true
-----
DriveLetter  FileSystemLabel  FileSystem  DriveType  HealthStatus  OperationalStatus  SizeRemaining  Size
-----
J             NTFS             NTFS       Fixed     Healthy       OK                 15.92 GB      15.95 GB

PS C:\Users\Administrator> fsutil fsinfo volumeinfo j:
Volume Name :
Volume Serial Number : 0x7ac53c40
Max Component Length : 255
File System Name : NTFS
fs Readwrite
Supports Case-sensitive filenames
Preserves Case of filenames
Supports Unicode in filenames
Preserves & Enforces ACL's
Supports Disk Quotas
Supports Reparse Points
Supports Object Identifiers
Supports Named Streams
Supports Hard Links
Supports Extended Attributes
Supports Open By FileID
Supports non-normal
Is DAX volume

```

Figure 22. DAX Mode

Storage Spaces Support

Windows Server 2016 supports NVDIMM-N devices that allow for extremely fast input/output (I/O) operations. One attractive way of using such devices is as a write-back cache to achieve low write latencies. [Microsoft blog](#) discusses how to set up a mirrored storage space with a mirrored NVDIMM-N write-back cache as a virtual drive. To setup storage spaces configuration on NVDIMM-N, see [Configuring Storage Spaces with a NVDIMM-N write-back cache](#).

Operational and Diagnostics Logging information

The NVDIMM-N bus driver, scmbus.sys, writes its logs to the Microsoft-Windows-Scmbus provider. To see them, open Event Viewer and navigate to Applications and Services Logs -> Microsoft -> Windows -> Scmbus

The NVDIMM-N disk driver, scmdisk0101.sys, writes its logs to the Microsoft-Windows-ScmDisk0101 provider. To see them, open Event Viewer and navigate to Applications and Services Logs -> Microsoft -> Windows -> ScmDisk0101

Following is the example NVDIMM-N Disk

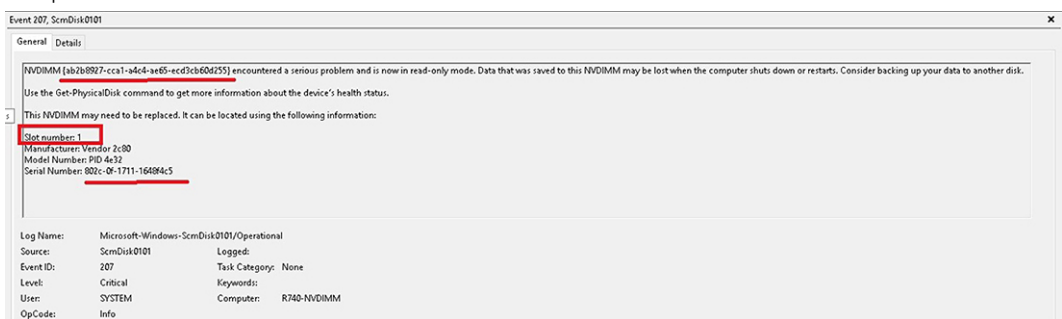


Figure 23. Operational and Diagnostics Logging information: NVDIMM-N Disk

By using the Uniqueid or serial number one can identify the NVDIMM-N physical location

```

PS C:\Users\Administrator> Get-PhysicalDisk -FriendlyName * | Format-Table -Property FriendlyName, uniqueID, PhysicalLocation, SerialNumber, healthStatus, OperationalStatus, OperationalDetails
FriendlyName    uniqueID          PhysicalLocation    SerialNumber    healthStatus OperationalStatus    OperationalDetails
-----
Vendor 2c80 PID 4e32 (4b298927-ccaf-a6c4-a6d5-ec4fcb04255) DIMM Socket 0, Slot 1 802c-0f-3721-1648f4c5 Unhealthy Device Hardware Error (Lost Data Persistence, Threshold Exceeded, W039M_N Error)

PS C:\Users\Administrator> Get-PhysicalDisk -SerialNumber 802c-0f-3721-1648f4c5 | Format-Table -Property FriendlyName, uniqueID, PhysicalLocation, SerialNumber, healthStatus, OperationalStatus, OperationalDetails
FriendlyName    uniqueID          PhysicalLocation    SerialNumber    healthStatus OperationalStatus    OperationalDetails
-----
Vendor 2c80 PID 4e32 (4b298927-ccaf-a6c4-a6d5-ec4fcb04255) DIMM Socket 0, Slot 1 802c-0f-3721-1648f4c5 Unhealthy Device Hardware Error (Lost Data Persistence, Threshold Exceeded, W039M_N Error)

```

Figure 24. NVDIMM-N physical location

Storage Class Memory in Windows Server 2019

NVDIMM-N FW Requirement

Windows Server 2019 requires NVDIMM-N devices to support labels in order for the OS to configure namespaces. NVDIMM-N Firmware 8860 doesn't have label space functionality and devices with this firmware cannot be configured by Windows Server 2019.

The NVDIMM-N firmware must be updated to FW 9324 or above so that the devices can be successfully configured and used in Windows Server 2019.

It is also recommended to update BIOS, iDRAC and the Operating System as part of the scheduled update.

Driver Architecture Overview

There are two types of device objects related to persistent memory in Windows Server – Physical NVDIMMs and Logical Persistent Memory Disks on top of the physical driver.

In Device Manager, physical NVDIMMs appear under "Memory devices", while logical disks are under "Persistent memory disks".

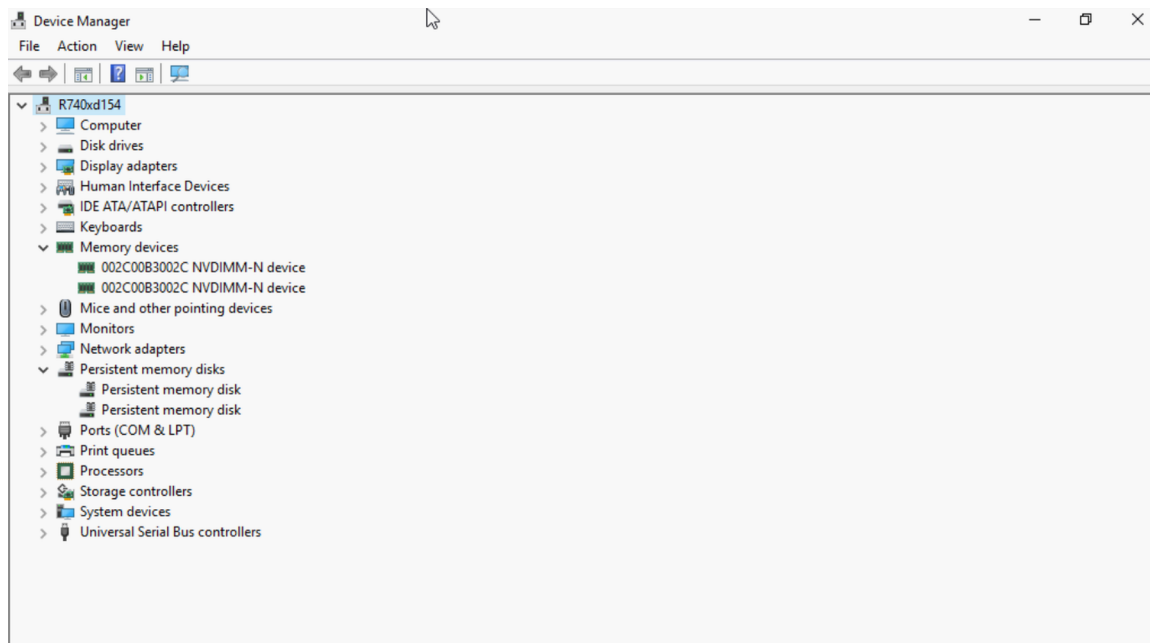


Figure 25. Device Manager

All NVDIMM-N devices are controlled by the nvdimms.sys driver, while the logical disks are controlled by the pmem.sys driver. Both types of device objects are created by scmbus.sys, the bus driver for persistent memory. This bus driver object can be found in Device Manager under "System Devices"

New features in Windows Server 2019 Label support and Namespace management

With Windows Server 2019, OS provides support for Label and Namespace management.

When NVDIMM-N modules are exposed to the OS first time or after factory-reset/sanitize operation, namespaces are not enumerated automatically.

Users need to create namespaces before using them as storage devices. User can run the commands below to manage namespaces on NVDIMM-N.

```
PS C:\Users\Administrator.WELABS> Get-Command -Module PersistentMemory
```

CommandType	Name	Version	Source
Cmdlet	Get-PmemDisk	1.0.0.0	PersistentMemory
Cmdlet	Get-PmemPhysicalDevice	1.0.0.0	PersistentMemory
Cmdlet	Get-PmemUnusedRegion	1.0.0.0	PersistentMemory
Cmdlet	Initialize-PmemPhysicalDevice	1.0.0.0	PersistentMemory
Cmdlet	New-PmemDisk	1.0.0.0	PersistentMemory
Cmdlet	Remove-PmemDisk	1.0.0.0	PersistentMemory

Figure 26. Available PMEM PowerShell Cmdlets

Examples:

```
PS C:\Users\Administrator.WELABS> Get-PmemUnusedRegion
```

RegionId	TotalSizeInBytes	DeviceId
1	17179869184	{1}
2	17179869184	{11}

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

DeviceId	DeviceType	HealthStatus	OperationalStatus	PhysicalLocation	FirmwareRevision
1	002C00B3002C NVDIMM-N device	Healthy	{Ok}	A7	9324
11	002C00B3002C NVDIMM-N device	Healthy	{Ok}	A8	9324

```
PS C:\Users\Administrator.WELABS> Get-PmemDisk
```

```
PS C:\Users\Administrator.WELABS> Get-PmemUnusedRegion 1 | New-PmemDisk
```

Creating new persistent memory disk. This may take a few moments.

```
PS C:\Users\Administrator.WELABS> Get-PmemUnusedRegion 2 | New-PmemDisk
```

Creating new persistent memory disk. This may take a few moments.

```
PS C:\Users\Administrator.WELABS> Get-PmemDisk
```

DiskNumber	Size	HealthStatus	AtomicityType	CanBeRemoved	PhysicalDeviceIds	UnsafeShutdownCount
9	16 GB	Healthy	None	True	{1}	0
10	16 GB	Healthy	None	True	{11}	0

Figure 27. List PMEM Unused regions, PMEM Physical Devices and PMEM Disks


```
PS C:\Users\Administrator.WELABS> Get-PmemDisk | fl *

DiskNumber      : 9
UniqueId        : {ba266808-3759-40d8-ad0e-bb71f4a23fcd}
HealthStatus    : Healthy
SizeInBytes     : 17179869184
AtomicityType   : None
CanBeRemoved    : True
UnsafeShutdownCount : 0
PhysicalDeviceIds : {1}

DiskNumber      : 10
UniqueId        : {2214f5df-2a6d-4854-83f4-4d7a988e762d}
HealthStatus    : Healthy
SizeInBytes     : 17179869184
AtomicityType   : None
CanBeRemoved    : True
UnsafeShutdownCount : 0
PhysicalDeviceIds : {11}

PS C:\Users\Administrator.WELABS> Remove-PmemDisk
cmdlet Remove-PmemDisk at command pipeline position 1
Supply values for the following parameters:
DiskNumber: 9
Removing the persistent memory disk. This may take a few moments.

PS C:\Users\Administrator.WELABS> Get-PmemDisk | fl *

DiskNumber      : 9
UniqueId        : {2214f5df-2a6d-4854-83f4-4d7a988e762d}
HealthStatus    : Healthy
SizeInBytes     : 17179869184
AtomicityType   : None
CanBeRemoved    : True
UnsafeShutdownCount : 0
PhysicalDeviceIds : {11}
```

Figure 28. List/Remove PMEM Disks

PowerShell Cmdlets

#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), RFC, 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 [Storage class memory] 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.

NVDIMM-N Interleaving

Windows Server 2019 supports NVDIMM-N interleaving which can further enhance system performance across various applications. This option can be configured through the BIOS Setup Menu.

In the OS, this is visible as a single logical device across several physical devices on that socket. Two interleaved persistent memory disks should be visible, one for each socket in the Device Manager, if there are NVDIMM-N modules plugged in on both the sockets.

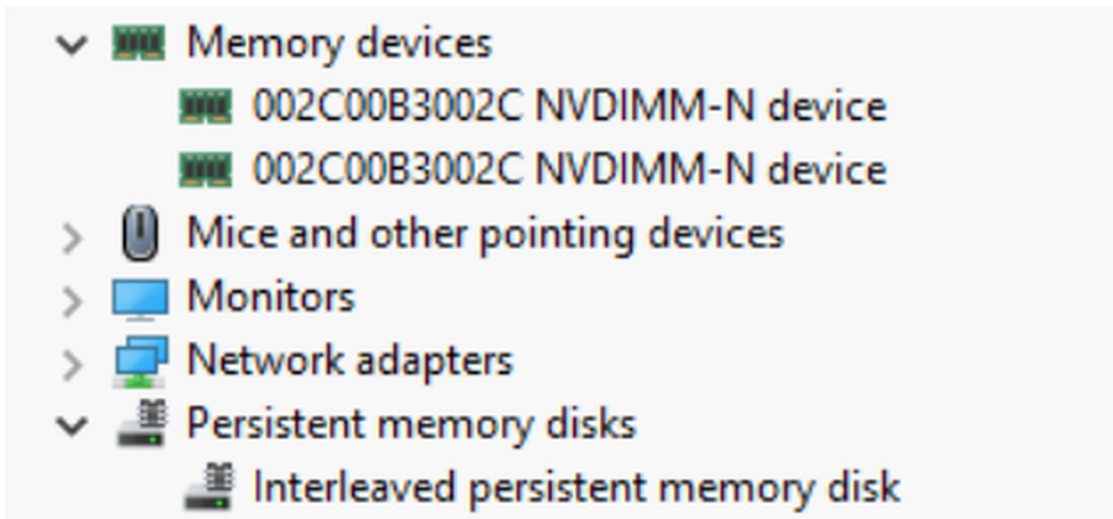


Figure 29. Interleaved Persistent Memory Disk in Device Manager

```

Administrator: Windows PowerShell
Windows PowerShell
Copyright (C) Microsoft Corporation. All rights reserved.

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

RegionId TotalSizeInBytes DeviceId
-----
1        34359738368 {1, 11}

PS C:\Users\Administrator.WELABS> Get-PmemPhysicalDevice

DeviceId DeviceType HealthStatus OperationalStatus PhysicalLocation FirmwareRevision Persistent
memory size
-----
1        002C00B3002C NVDIMM-N device Healthy {Ok} A7 9324 16 GB
11       002C00B3002C NVDIMM-N device Healthy {Ok} A8 9324 16 GB

PS C:\Users\Administrator.WELABS> Get-PmemUnusedRegion 1 | New-PmemDisk
Creating new persistent memory disk. This may take a few moments.
PS C:\Users\Administrator.WELABS> Get-PmemDisk

DiskNumber Size HealthStatus AtomicityType CanBeRemoved PhysicalDeviceIds UnsafeShutdownCount
-----
9          32 GB Healthy None True {1, 11} 0

```

Figure 30. Visibility in PowerShell

Configuring NVDIMM-N for Hyper-V Virtual Machines

The article referenced here, [Cmdlets for configuring persistent memory devices for Hyper-V VMs](#), provides details about configuring Hyper-V VMs with JEDEC compliant NVDIMM-N.

NVDIMM-N RO Behavior

Windows Server 2019

By design, Windows Server 2019 manages NVDIMM-N in a manner that differs from that of Windows Server 2016. In Windows Server 2016, if an NVDIMM-N that is in read/write mode loses power, it goes into read-only mode.

But in Windows Server 2019, if an NVDIMM-N that is in read/write mode loses power, it continues to function in read/write mode.

Similarly, if the NVDIMM-N modules are forced to Read Only through the BIOS Setup Menu or through any module related failures, these devices would still be read/write accessible under Windows Server 2019.

To make Windows Server 2019 manage NVDIMM-N in the same way as Windows Server 2016 RTM, a registry entry needs to be modified.

For details refer to <https://support.microsoft.com/en-us/help/4492760/nvdimm-in-read-write-mode-after-backup-power-removed-in-windows-server>

Windows Errata

The following errata affects Windows OSes and are expected to be fixed in a future OS patch.

- In Windows Server 2016, the Storage Spaces GUI shows NVDIMM-N bus type as UNKNOWN. It doesn't have any functionality loss or inability in creation of Pool, Storage VD
- Windows Server 2016 and Windows 10 Anniversary Edition do not support interleaved sets of NVDIMM-Ns

- Hyper-V role in Windows Server 2016 and Windows 10 Anniversary Edition do not support NVDIMM-Ns, i.e., NVDIMM-N cannot be directly exposed to VMs.

NVDIMM-N hardware is supported on versions 7.3, 7.4 , 7.5 and 7.6 versions of RHEL.

Topics:

- [Identify and Configure PMEM —Persistent Memory Device](#)
- [Installation](#)
- [Verify Existing Filesystem](#)
- [Read-Only Mode NVDIMM-N](#)
- [Interleave](#)
- [Management Utility](#)
- [RHEL 7.6 features](#)
- [Linux Errata](#)

Identify and Configure PMEM —Persistent Memory Device

When the OS is up running, verify that NVDIMM-Ns are populated correctly.

Go into root user

```
$ su
```

Identify whether NVDIMM-Ns appear as `/dev/pmem0`, `/dev/pmem1`, .. `/dev/pmemN`, where N is a natural number.

```
# ls /dev/pmem*
```

Check the size of NVDIMM-N devices `/dev/pmem0` to `/dev/pmemN`

```
# lsblk
```

Create xfs file system for all pmemN devices

```
# mkfs.xfs /dev/pmem0
```

Create directories on all nvdimmN

```
# mkdir -p /mnt/nvdimm0
```

```
# mount -t xfs -o dax /dev/pmem0 /mnt/nvdimm0
```

Save the mount point and option so that devices are mounted on next reboot.

```
# echo "/dev/pmem0 /mnt/nvdimm0 xfs dax 0 0" >> /etc/fstab
```

Write or copy files onto `/mnt/nvdimm0`

```
# echo "writing into nvdimm" >> /mnt/nvdimm0/write.txt
```

```
# shutdown
```

Installation

Dump the RHEL ISO onto the USB stick via *dd* command.

```
# dd if=/home/dell/RHEL7.3.iso of=/dev/sdb bs=4M conv=noerror, sync
```

BIOS boots the Linux kernel from USB.

Follow the on-screen steps to finish installing RHEL.

After the installation is completed, reboot the server.

For detailed installation instructions, please refer to https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux/7/html/Installation_Guide/index.html

Verify Existing Filesystem

Press PowerOn button on the server

When the OS is up running,

```
$ su
```

Verify contents stored on NVDIMM-Ns

```
# cat /mnt/nvdim0  
# shutdown
```

Read-Only Mode NVDIMM-N

Press PowerON button on the server

Follow the guidance in Section 4 to setup BIOS.

Enable the Persistence Memory, disable “NVDIMM Interleave”, enable “NVDIMM Read-Only”.

When OS is up and running,

```
$ su
```

To see whether it is writable

```
# touch /mnt/nvdim0/write.txt  
# shutdown
```

Interleave

For R740/R640, insert NVDIMM-Ns in memory slots of both CPU Sockets. For this example 6 NVDIMM-Ns are installed on CPU0, and 6 are installed on CPU1.

Interleave Setup

Press PowerOn button on the server

Follow the guidance in Section 4 to setup BIOS.

Enable the Persistence Memory, enable “NVDIMM Interleave”, disable “NVDIMM Read-Only”.

Install RHEL or start OS if it is already installed

When OS is up running,

```
$ su
```

CPU0's 6 NVDIMM-Ns show up as /dev/pmem0, CPU1's 6 NVDIMM-Ns appear as /dev/pmem1.

```
# ls /dev/pmem*
```

View the size of /dev/pmem0 and /dev/pmem1, each should be around 6*16 GB = 96GB because each NVDIMM-N is 16 GB.

```
# lsblk
```

Create xfs file system for /dev/pmem0 and /dev/pmem1

```
# mkfs.xfs /dev/pmem0
```

Create directory /mnt/nvdimm0 and /mnt/nvdimm1

```
# mkdir -p /mnt/nvdimm0
```

Mount /dev/pmem0 and /dev/pmem1

```
# mount -t xfs -o dax /dev/pmem0 /mnt/nvdimm0
```

Save the mount point and option so that devices will be mounted on next reboot.

```
# echo "/dev/pmem0 /mnt/nvdimm0 xfs dax 0 0" >> /etc/fstab
```

Copy files to /mnt/nvdimm0 and /mnt/nvdimm1, then shutdown.

```
# echo "writing into nvdimm" >> /mnt/nvdimm0/write.txt  
# shutdown
```

Interleave Verification

Press Power On button on the server

When the OS is up running,

```
$ su
```

Verify contents of /mnt/nvdimm0, and do the same thing for /mnt/nvdimm1

```
# cat /mnt/nvdimm0  
# shutdown
```

Two mounted devices, /mnt/nvdimm0 and /mnt/nvdimm1 will appear, each size is approximately 96GB, or 6x NVDIMM-Ns.

The files stored on /mnt/nvdimm0 and /mnt/nvdimm1 will persist across reboots.

Read Only Mode NVDIMM-N

Press PowerON button on the server

Follow the guidance in Section 4 to setup BIOS.

Enable the Persistence Memory, enable “NVDIMM Interleave”, enable “NVDIMM Read-Only”.

When OS is up and running,

```
$ su
```

Mount /dev/pmem0 and /dev/pmem1

```
# mount -t xfs -o dax /dev/pmem0 /mnt/nvdimm0
```

To see whether it is writable on /dev/pmem0 and /dev/pmem1

```
# touch /mnt/nvdimm0/write.txt
# shutdown
```

Management Utility

Management Utility ‘ndctl’ && mdadm

- 1 Press PowerOn button on the server.
- 2 Follow the guidance in Section 4 to setup BIOS.
- 3 Enable the Persistence Memory, disable “NVDIMM Interleave”, disable “NVDIMM Read-Only”.
- 4 Install RHEL or start OS if it is already installed.

ndctl

Install ndctl -- developer option

When OS is up running, enable the Ethernet and get an IP address

```
$ ifconfig -a
$ git clone https://github.com/pmem/ndctl.git
$ cd ndctl
$ ./autogen.sh
$ ./configure CFLAGS='-g -O0' --prefix=/usr --sysconfdir=/etc --libdir=/usr/lib64
$ make
# make install
```

Install ndctl – user option

ndctl rpm packages can be installed from the RHEL 7.4 releases via

```
# yum install ndctl
```

Check health status of NVDIMM-N

```
# ndctl list --dimms --health --idle'
Sample output of 'ndctl list --dimms --health --idle'
{
  "dev": "nmem0",
  "id": "802c-0f-1711-1648dd20",
  "state": "disabled",
  "health": {
```



```
    "health_state": "ok",
    "temperature_celsius": 27.000000,
    "life_used_percentage": 3
  }
}
```

mdadm

Create software RAID on NVDIMM-Ns.

Say there are 6 devices, /dev/pmem0.. /dev/pmem5.

Create directory /mnt/md0 /mnt/md1 /mnt/md2 /mnt/md5 /mnt/md6.

```
$ mkdir -p /mnt/md0
```

Create RAID 0

```
$ mdadm --create --verbose /dev/md0 --level=0 --raid-devices=6 /dev/pmem0 /dev/pmem1 /dev/
pmem2 /dev/pmem3 /dev/pmem4 /dev/pmem5
$ cat /proc/mdstat
$ mkfs.xfs /dev/md0
$ mount /dev/md0 /mnt/md0
$ lsblk
```

Before you unmount the device, Read/Write data into /mnt/md0

```
$ umount /dev/md0
$ mdadm --stop /dev/md0
$ mdadm --zero-superblock /dev/pmem0 /dev/pmem1 /dev/pmem2 /dev/pmem3 /dev/pmem4 /dev/pmem5
$ lsblk
```

Creating RAID 1 is similar as above

```
$ mdadm --create --verbose /dev/md1 --level=1 --raid-devices=6 /dev/pmem0 /dev/pmem1 /dev/
pmem2 /dev/pmem3 /dev/pmem4 /dev/pmem5
$ cat /proc/mdstat
$ mkfs.xfs /dev/md1
$ mount /dev/md1 /mnt/md1
$ lsblk
```

Before unmounting the device, Read/Write data into /mnt/md1

```
$ umount /dev/md1
$ mdadm --stop /dev/md1
$ mdadm --zero-superblock /dev/pmem0 /dev/pmem1 /dev/pmem2 /dev/pmem3 /dev/pmem4 /dev/pmem5
$ lsblk
```

RAID 5/5/6/10 Volumes can be created in a similar way.

RHEL 7.6 features

Namespace creation

Starting from RHEL 7.6, ndctl utility is available natively.

NVDIMM-N modules that have FW 9324 are not visible in the OS as this firmware has label support. In other words, namespaces do not get enumerated automatically and using ndctl, namespaces have to be created first before mounting a file system and accessing them.

NOTE: It is also recommended that an OS patch is installed first to ensure proper functionality of the ndctl utility for creating namespaces. This patch can be downloaded here <https://access.redhat.com/solutions/3920221>.

Run through the commands below as a super user to create /dev/pmem devices.

To display DIMM list run:

- `$ ndctl list -D`

Run command below to create namespaces. This command if used as is by default creates /dev/pmem devices. In order to create namespaces in other modes, refer to the link <https://www.mankier.com/1/ndctl-create-namespace>

This command should be run as many times as the number of NVDIMM-N modules plugged into the system.

- `$ ndctl create-namespace`

For more information on how to use ndctl utility refer to the users guide here <https://docs.pmem.io/ndctl-users-guide>

Linux Errata

Following errata effects RHEL 7.6:

- Patch to make ndctl utility work can be downloaded from here <https://access.redhat.com/solutions/3920221>

Topics:

- [Set up](#)
- [Storage](#)
- [Supported Guest OSES with NVDIMM support](#)
- [Overall Health Status](#)
- [Operational and Diagnostics Logging information](#)
- [NVDIMM-N Errors](#)
- [ESXi Errata](#)

Set up

Beginning with ESXi version 6.7 NVDIMM-N hardware is supported. The NVDIMM-N devices will be detected on startup with auto-generated namespaces. All NVDIMM-N devices will have capacity pooled into a single logical memory array for access by ESXi virtual machines.

NVDIMM-N Hardware is now supported on ESXi 6.7U1. Features and functionality should be identical to ESXi 6.7

Firmware requirements:

Table 13. Firmware requirements

Firmware	Minimum Required Versions		
	T640, R640, R740, R740XD, R940	R840, R940xa	MX740c, MX840c
IDRAC	3.15.15.15	3.19.19.19	3.20.20.20
BIOS	1.3.7	1.1.1	1.0.0
NVDIMM-N	9324	9324	9324

Please locate your product on support.dell.com to locate update packages for any of these firmware bundles.

Hardware requirements:

ESXi currently only supports NVDIMM-N hardware when the host system has an Intel Xeon™ Gold (51xx, 61xx) or Platinum (81xx) processor SKUs. Silver and Bronze SKUs are not sufficient for ESXi solution support at this time, but may be added in the future.

For BIOS Setup, default Persistent Memory settings can be used. However, both Non-Interleaved and Interleaved settings are supported by ESXi.

When adding a new VM to the ESXi host, follow these steps:

- 1 Click on “Add other device” and select “NVDIMM” from the menu.
- 2 Select the total capacity in MB of persistent memory you wish to expose to the VM.
 - Size should not exceed total NVDIMM-N capacity in the host system.

Storage

Modules

Below picture shows how the NVDIMM-N devices will appear in the Persistent Memory tab under the Storage heading in the left menu bar. The first column is an ID to uniquely identify the NVDIMM-N in the event the information is needed for troubleshooting. The next column indicates free space and it is expected to be "0 B" for all NVDIMM-N devices that are fully mapped and operated correctly. Health should be Normal. Detailed explanation of Health section is provided in "Overall health Status" below.

Translation of ID to physical NVDIMM-N in host system can be seen below.

Table 14. Translation of ID to physical NVDIMM-N in host system

ID (Hex)	Decimal	Physical Location (L->R)	Silk screen slot name
0x1	16	Socket 0 , IMC0, Channel 0 , Slot 1	A7
0x11	17	Socket 0 , IMC0, Channel 1 , Slot 1	A8
0x21	33	Socket 0 , IMC0, Channel 2 , Slot 1	A9
0x101	257	Socket 0 , IMC1, Channel 0 , Slot 1	A10
0x111	273	Socket 0 , IMC1, Channel 1 , Slot 1	A11
0x121	289	Socket 0 , IMC1, Channel 2 , Slot 1	A12
0x1001	4097	Socket 1 , IMC0, Channel 0 , Slot 1	B7
0x1011	4113	Socket 1 , IMC0, Channel 1 , Slot 1	B8
0x1021	4129	Socket 1 , IMC0, Channel 2 , Slot 1	B9
0x1101	4353	Socket 1 , IMC1, Channel 0 , Slot 1	B10
0x1111	4369	Socket 1 , IMC1, Channel 1 , Slot 1	B11
0x1121	4385	Socket 1 , IMC1, Channel 2 , Slot 1	B12

Namespaces

All Namespace names are automatically generated by ESXi when the NVDIMM-N modules are detected on startup. These Namespaces will persist over warm reset and cold boot of the system after first detection/creation.




Figure 31. Namespaces

Interleave sets

If Interleaving is Disabled in BIOS F2 Setup, then each NVDIMM-N will be in an individual Interleave Set as seen in the GUI below.




Figure 32. Interleave sets while Interleaving is Disabled

If Interleaving is Enabled in BIOS F2 setup, then the total NVDIMM-N capacity will be split into pools based on CPU socket. A total of two interleave sets will display with the aggregate capacity for the CPU socket displaying as one Interleave Set.




Figure 33. Interleave sets while Interleaving is Enabled

Datastore

This is the Datastore area which includes information about the traditional storage accessible to ESXi and the persistent memory Datastore.

NOTE: The PMEM Datastore is created on the first system boot after NVDIMM-Ns are detected in a system by ESXi. This datastore is fixed at that time and will not expand if additional NVDIMM-N modules are added after the fact.

To expand the datastore:

- 1 Migrate your data.
- 2 Boot into BIOS F2 Setup.
- 3 Select Persistent Memory as indicated in section 5 of this guide.
- 4 Select Factory Default and Secure Erase for all NVDIMM-N modules in the system.
- 5 Follow prompts to reboot and continue on into ESXi. The additional NVDIMM-N modules will now be detected.

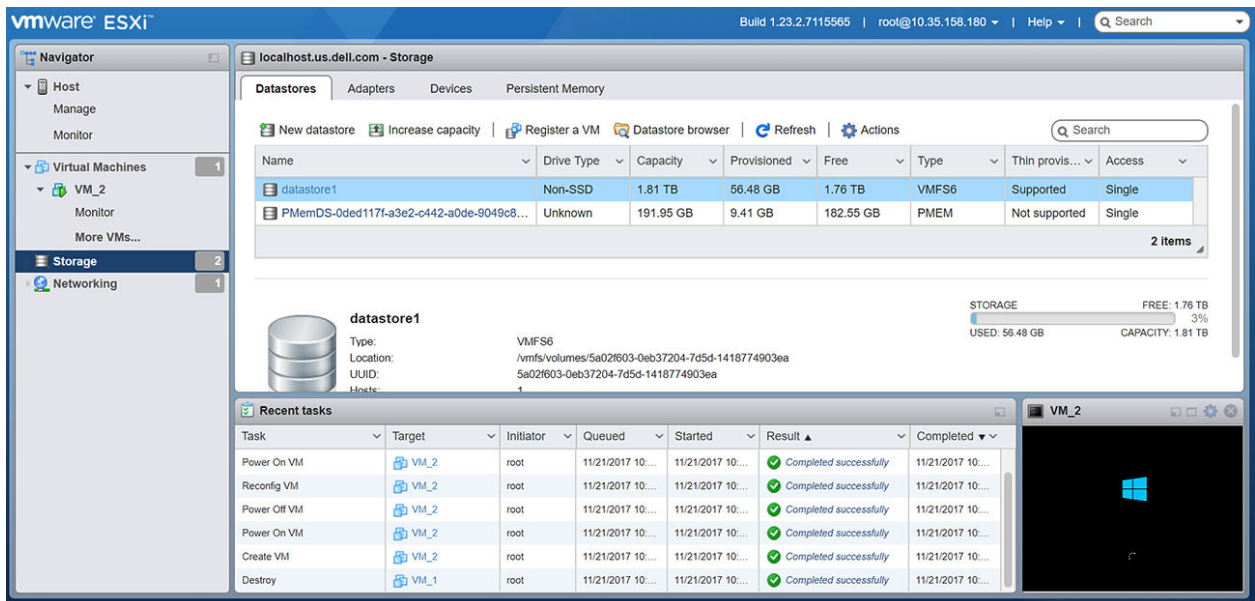


Figure 34. Datastores

Supported Guest OSES with NVDIMM support

- Windows Server 2016 Build 14393 and above
- Windows 10 Anniversary Update Version 1607 and above
- RedHat Enterprise Linux 7.4 and above
- SUSE Linux Enterprise 12 SP2 and above
- Photon OS 1.0 Revision 2 and above
- CentOS 7.4 and above
- Ubuntu 17.04 and above

Overall Health Status

Health status of the NVDIMM-N modules is represented in a tabular column of the Modules and Namespace section of the ESXi interface. Please take corrective action as indicated in the table below.

Redundant Table:

Table 15. Health status of the NVDIMM-N modules

Status	Action
Normal	None
Maintenance Needed	Please consult ESXi, SEL, and Lifecycle logs for instructions to troubleshoot the NVDIMM-N status.
All data loss in the event of power loss	<ol style="list-style-type: none">1 Do not shut off the system. Migrate data from the NVDIMM-N to another location as soon as possible.2 After the NVDIMM-N data is successfully migrated, consult ESXi, SEL and Lifecycle logs for instructions to troubleshoot the NVDIMM-N status.3 Confirm the battery is healthy.

Operational and Diagnostics Logging information

ESXi has many available logs to diagnose system behavior located in the “Monitor” section of the user interface. You can also create a complete Support Bundle here which will aggregate all logs into a simple link for download.

Outdated firmware

ESXi requires NVDIMM-N modules to have a minimum 9324 firmware image. If they have older firmware they will not behave correctly.

In the event of out dated firmware on the memory (lower than 9324), the system will boot into the ESX hypervisor and the DIMMs will be visible in the UI. **No namespaces will be populated** and the DIMMs cannot be mounted to a VM guest OS. Since this lower firmware is unsupported, other error behaviors such as Overall Health Status, Address Range Scrub, and UI error reporting are not tested and may not function properly until DIMM firmware is up to a current version.

Though unsupported, Storage->Modules and Storage->Interleave Sets appear to display the correct information.

When NVDIMMs with 9324 and 8860 FW revisions are mixed, the ones with 9324 rev can be used and will be indicated as “Normal” in the Web GUI whereas the others will be indicated as “Maintenance Needed”.

When an NVDIMM-N firmware update fails, the DIMM will operate off of backup firmware. For NVDIMM-N modules shipped in 2017 and 2018 this backup firmware is expected to be at revision 8860. Unless the NVDIMMs are re-updated to revision 9324, they cannot be used in ESXi.

The screenshot shows the 'Modules' section in the ESXi interface. A table lists 12 NVDIMM-N modules, each with a 16 GB capacity and a 'Maintenance needed' health status. The table includes columns for ID, Capacity, Free, and Health. A search bar and a 'Refresh' button are visible at the top of the table area.

ID	Capacity	Free	Health
0x1	16 GB	16 GB	Maintenance needed
0x11	16 GB	16 GB	Maintenance needed
0x21	16 GB	16 GB	Maintenance needed
0x101	16 GB	16 GB	Maintenance needed
0x111	16 GB	16 GB	Maintenance needed
0x121	16 GB	16 GB	Maintenance needed
0x1001	16 GB	16 GB	Maintenance needed
0x1011	16 GB	16 GB	Maintenance needed
0x1021	16 GB	16 GB	Maintenance needed
0x1101	16 GB	16 GB	Maintenance needed
0x1111	16 GB	16 GB	Maintenance needed
0x1121	16 GB	16 GB	Maintenance needed

Figure 35. NVDIMM-N modules status

NVDIMM Namespaces do not get created when the NVDIMM-N devices are running of firmware revision 8860.

The screenshot shows the 'Storage' view in the ESXi interface, specifically the 'Persistent Memory' tab. The 'Namespaces' section is selected in the left-hand menu. The main area shows a table with columns for Name, Capacity, Health, State, and Location id. The table is empty, displaying the message 'No namespaces to display'.

Name	Capacity	Health	State	Location id
No namespaces to display				

Figure 36. NVDIMM Namespaces

NVDIMM-N Errors

In the majority of cases where NVDIMM-N modules have an error, the system level behavior will be for the bad NVDIMM-N module to be put into a read-only state. This will cause any VMs mapped to **ANY** NVDIMM-N modules in the system to fail to power-on including VMs which are not mapped to the bad NVDIMM-N module. If an NVDIMM-N is flagged as having an error, resolve the issue with that NVDIMM-N then attempt to access VMs connected to the PMEM datastore again. Please see ESXi documentation for more information about migrating data and restoring access to impacted VMs.




Figure 37. NVDIMM-N Errors

Refer to table below for the Overall Health Status message shown on ESXi Web Client in the event of the following errors:

Table 16. Overall Health Status message shown on ESXi Web Client

Status	Errors
Save Error – pre OS boot	Maintenance Needed
Restore Error – pre OS boot	Maintenance Needed
Erase Error – pre OS boot	Maintenance Needed
Arm Error – pre OS boot	Maintenance Needed
Controller Error – pre OS boot	Maintenance Needed
Controller Error at OS run time	No OHS error until next boot, logged in SEL real time
Invalid Firmware or Firmware running from Backup Slot	Maintenance Needed
NVDIMM-N Battery Failure at OS run-time	All Data Loss
NVDIMM-N Battery Failure – pre OS boot	Maintenance Needed

All the above error messages are also accompanied by equivalent messages in the System Event Log and Lifecycle Controller Log as described in the Sections 5 and 6 of this User Guide.

ESXi Errata

None at this time