

# 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.

**Copyright © 2017 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

<b>1 Introduction.....</b>	<b>5</b>
<b>2 NVDIMM-N Overview.....</b>	<b>6</b>
Normal Operation.....	6
Backup to Flash.....	7
Restore from Flash.....	8
<b>3 Hardware.....</b>	<b>9</b>
Server Hardware Configuration.....	9
NVDIMM-N.....	12
Battery.....	13
<b>4 BIOS.....</b>	<b>16</b>
BIOS Configuration Settings for NVDIMM-N.....	16
BIOS Error Messages.....	20
<b>5 iDRAC NVDIMM-N Management.....</b>	<b>22</b>
iDRAC Graphical User Interface.....	22
NVDIMM-N Status.....	22
BBU Status.....	23
Remote Management.....	24
NVDIMM-N Error Reporting.....	24
<b>6 Server Behavior with NVDIMM-Ns.....</b>	<b>27</b>
Shutdown.....	27
Boot.....	27
Automatic Shutdown and Save.....	28
<b>7 DIMM Configuration Changes.....</b>	<b>29</b>
<b>8 Windows.....</b>	<b>30</b>
Set Up.....	30
Windows Drivers.....	30
Storage Class Memory.....	30
Device manager.....	30
Identifying the right NVDIMM-N disks.....	31
NVDIMM-N health status and properties.....	34
Block Mode.....	34
DAX Mode.....	34
Storage Spaces Support .....	35
Operational and Diagnostics Logging information.....	35
Windows Errata.....	36

<b>9 Linux.....</b>	<b>37</b>
Identify and Configure PMEM —Persistent Memory Device.....	37
Installation.....	38
Verify Existing Filesystem.....	38
Read-Only Mode NVDIMM-N.....	38
Interleave.....	38
Interleave Setup.....	38
Interleave Verification.....	39
Read Only Mode NVDIMM-N.....	39
Management Utility .....	40
ndctl.....	40
mdadm.....	41
Linux Errata.....	41
mdadm.....	41

# 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 support in Microsoft Windows Server 2016 and RHEL 7.3. VMWare support is planned for a future release

## 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 AC 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](#)
- [NVDIMM-N](#)
- [Battery](#)

## Server Hardware Configuration

NVDIMM-Ns are currently supported in the T640, R640, R740/R740XD, and R940 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 R940 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 R940 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.

**Table 2. : 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

NVDIMM-N	NVDIMM-N Capacity	RDIMMs	RDIMM Capacity	DIMM Population Locations
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
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.

**Table 3. : Supported NVDIMM-N Configurations, R940 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: 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
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.



**Figure 5. R740/R740XD Memory Layout**

## NVDIMM-N

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 I<sup>2</sup>C temperature sensor with integrated SPD EEPROM




Figure 6. NVDIMM-N

Table 4. : 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, Dell EMC'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 AC 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, and 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.

Table 5. Battery LED Behavior

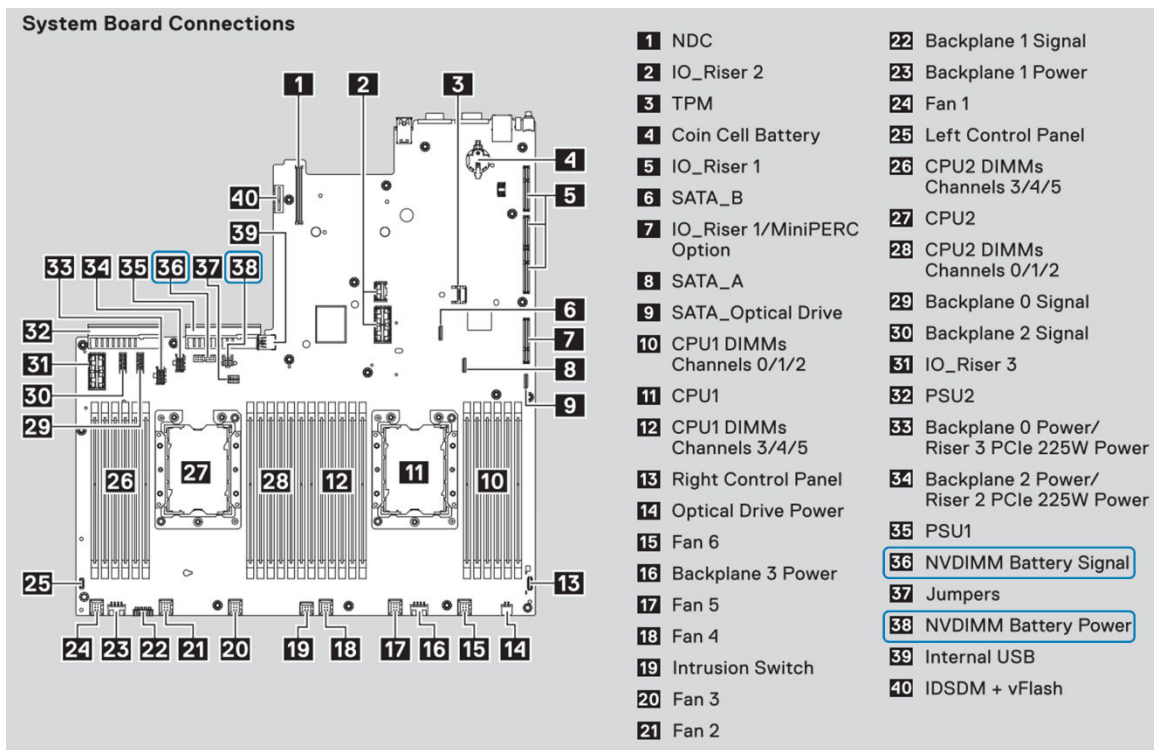
Battery Operation	Green LED	Amber LED
Powered Off	Off	Off
Idle (not charging and not enabled while the server is on)	Off	Off

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

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 AC 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 7. 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 8. R740 Battery Installation Instructions**

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

Topics:

- [BIOS Configuration Settings for NVDIMM-N](#)
- [BIOS Error Messages](#)

## BIOS Configuration Settings for NVDIMM-N

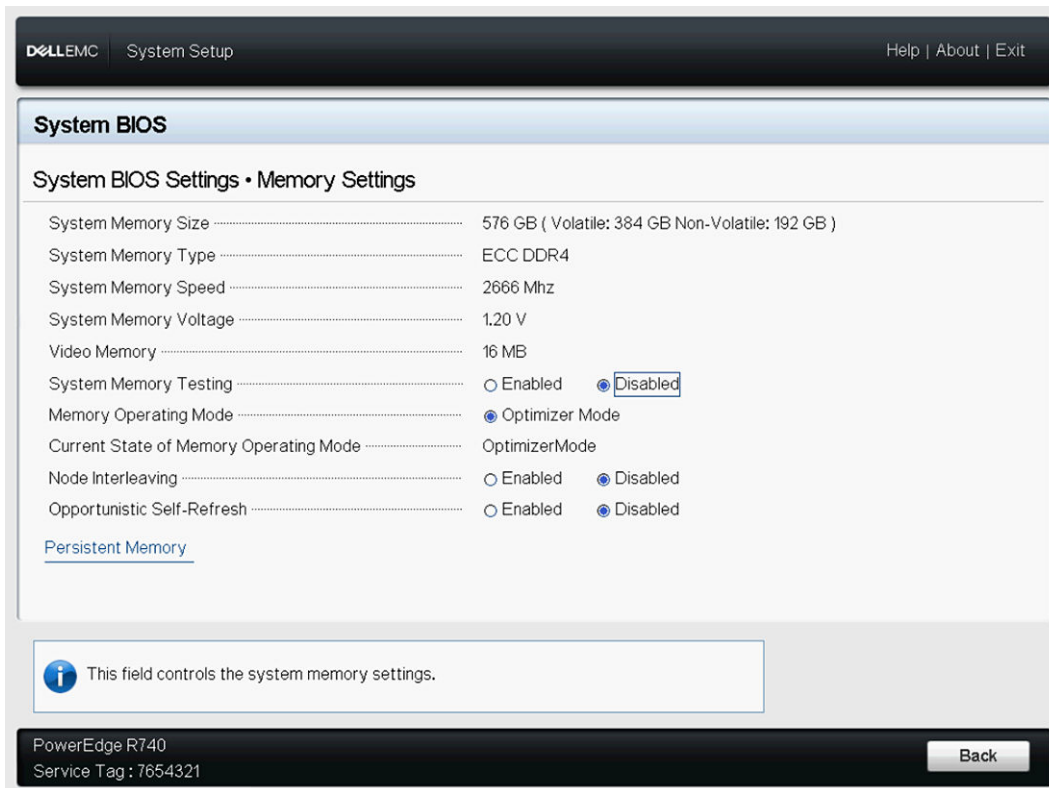
This section focuses only on the BIOS setup options that affect NVDIMM-N operation. For a description of all setup options, please refer to each server's Installation and Service Manual. Persistent Memory BIOS settings are configurable by going into BIOS System Setup. Press F2 at the BIOS screen below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



**Figure 9. BIOS screen**

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






**Figure 10. 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.

**When Persistent Memory Interleaving is enabled in the Persistent Memory System Setup page, this setting also applies to Persistent Memory Interleaving. DellEMC recommends that that this option be Disabled for NVDIMM-N configurations.**

Select Persistent Memory to configure NVDIMM-N setup options:



**Figure 11. Persistent Memory screen**

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

**Table 6. BIOS setup screen**

Option	Description
Node Interleaving	<p>Specifies if Non-Uniform Memory Architecture (NUMA) is supported. If this field is set to <b>Enabled</b>, memory interleaving is supported if a symmetric memory configuration is installed. If the field is set to <b>Disabled</b>, the system supports NUMA (asymmetric) memory configurations.</p> <p>When Persistent Memory Interleaving is enabled in the <u>Persistent Memory System Setup</u> page, this setting also applies to Persistent Memory Interleaving. Dell EMC recommends that that this option be <b>Disabled</b> for NVDIMM-N configurations.</p> <p>This option is set to <b>Disabled</b> by default.</p>
Persistent Memory	<p>Enables or disables NVDIMM-N operation. If this option is set to <b>Off</b>, 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 <b>Non-Volatile DIMM</b>.</p> <p>This option is set to <b>Non-Volatile DIMM</b> by default.</p>
NVDIMM-N Read-Only	<p>When set to <b>Enable</b>, 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 <b>Disable</b> 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 <b>Enable</b> , 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 <b>Disable</b> 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 <b>Disable</b> by default.
Battery Status	Indicates if the NVDIMM-N battery is ready. <b>Battery Status</b> can display one of the following states: <ul style="list-style-type: none"> <li>• Present-Ready</li> <li>• Present-Offline</li> <li>• Not-Present</li> </ul>

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

**Table 7. 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.
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 12. 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** 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, reseal 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, reseal 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, reseal 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, reseal 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, reseal the NVDIMM-N module, and restart the server. If the issue persists, replace the faulty memory module identified in the message.

**UEFI0302** NVDIMM-N Lifetime Percentage <= 1% 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.

**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 writeprotect 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 been enabled in BIOS Setup.

\* 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 a green checkmark for 'OK' status: 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
OK	An OS graceful shut-down occurred.	Wed 02 Aug 2017 09:31:55
OK	OEM software event.	Wed 02 Aug 2017 09:31:55
OK	C: boot completed.	Wed 02 Aug 2017 09:28:00
- Notes:** A section indicating 'There are no work notes to be displayed.'

Figure 13. iDRAC Graphical User Interface

## NVDIMM-N Status

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



**Figure 14. 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 has detected that NVDIMM-Ns are installed in the server. The Battery status will not be reported if NVDIMM-Ns are not installed.

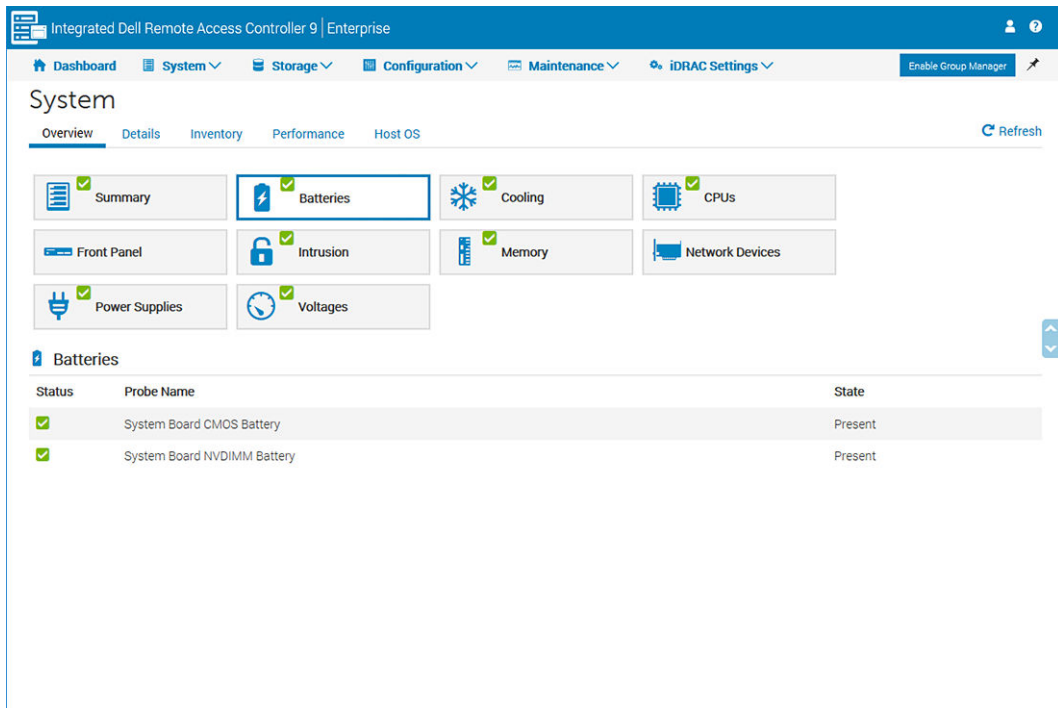


Figure 15. BBU Status

## 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 8. 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.



ID	Event Message	Recommended Action
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 life duration and is placed in write-protect mode.	Replace the Non-Volatile Dual In-line Memory Module (NVDIMM) device. For information about removing and reinstalling 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

ID	Event Message	Recommended Action
		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 onboard flash. The Save is triggered by events that would result in power loss to the NVDIMM-N. During the Save, the server maintains power to the NVDIMM-Ns, and a battery is used to provide back-up power in case of AC 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.



**Figure 16. 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.

**Note:** When NVDIMMs are installed, verify that the Server is completely powered down after removing AC Power before attempting to perform any service actions. The server is completely powered down when all LEDs including NVDIMM-N and NVDIMM-N Battery LEDs are off.

## 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 AC power loss. BIOS also verifies that the installed server Power Supplies are sufficiently sized for the server configuration. This is to ensure that after AC 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.

# DIMM Configuration Changes

Dell EMC 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, Dell EMC 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, Dell EMC recommends that all DIMMs be restored to their exact same locations as in the original server. Dell EMC also recommends that BIOS NVDIMM settings are identical.

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

# Windows

Topics:

- [Set Up](#)
- [Windows Drivers](#)
- [Storage Class Memory](#)
- [Windows Errata](#)

## Set Up

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

For BIOS Setup, enable the Persistence Memory setting and Disable the NVDIMM-N Interleave setting. These are also the default BIOS settings. Windows Server 2016 and Windows 10 Anniversary Edition currently do not support interleaved sets of NVDIMM-Ns.

[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 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 in Windows Server 2016](#)

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

- [Using Non-volatile Memory \(NVDIMM-N\) as Byte-Addressable Storage in Windows Server 2016](#)

<https://channel9.msdn.com/Events/Build/2016/P470>

## Storage Class Memory

## Device manager

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




Figure 17. 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-a4c4-ae65-ecd3cb60d255}	DIMM Socket 0 : Slot 17	802c-0f-1642-146f4600	8860	SCM
Vendor 2c80 PID 4e32	{859b53cd-e733-f8f0-03a3-abce893c00a8}	DIMM Socket 0 : Slot 17	802c-0f-1642-146f4600	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-ferda701db72}	DIMM Socket 0 : Slot 49	802c-0f-1711-164af113	8860	SCM
Vendor 2c80 PID 4e32	{de8dc85f-71a1-607f-c956-ab9ce1c585e3}	DIMM Socket 0 : Slot 65	802c-0f-1711-16493233	8860	SCM
Vendor 2c80 PID 4e32	{e9127230-73db-750c-339a-632db15af10}	DIMM Socket 0 : Slot 81	802c-0f-1642-146ea236	8860	SCM
Vendor 2c80 PID 4e32	{d7c2ed36-7c0f-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	{1e14cda8-18b1-834c-f8c5-02dd69b325672}	DIMM Socket 1 : Slot 33	802c-0f-1642-146ef6e7	8860	SCM
Vendor 2c80 PID 4e32	{09dc11e1-17ce-103b-268b-cd7403cddf09}	DIMM Socket 1 : Slot 49	802c-0f-1711-1648a682	8860	SCM
Vendor 2c80 PID 4e32	{67e6371f-cc0b-6257-1ce8-f96d60bb90e5}	DIMM Socket 1 : Slot 65	802c-0f-1642-146ea1f8	8860	SCM
Vendor 2c80 PID 4e32	{9e6d7051-af45-bd96-2e88-8d7a98724af7}	DIMM Socket 1 : Slot 81	802c-0f-1642-146ea274	8860	SCM

Figure 18. Identifying the right NVDIMM-N disks

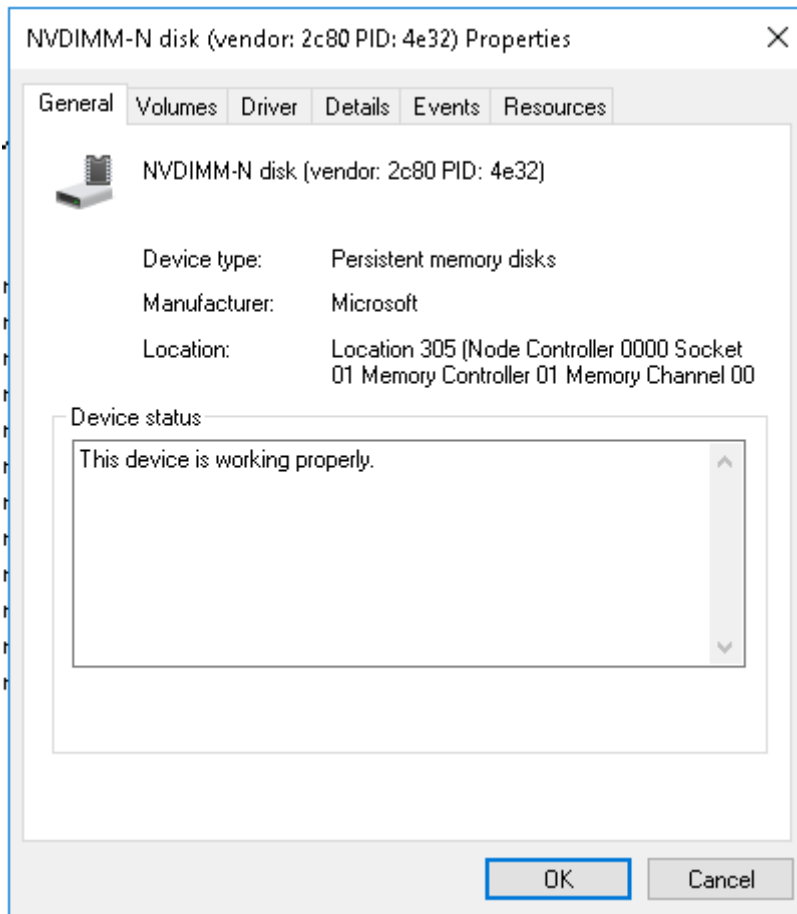
The Serial Number for every NVDIMM-N is unique, and physical location values in PowerShell can be mapped to silk screen using the following table

**Table 9. DIMM slot location**

<b>PowerShell PhysicalLocation</b>	<b>Dimm Slot Location</b>
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 19. Using device manager GUI**

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

**Table 10. 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 20. 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> activate-disk
PS C:\Users\Administrator> get-disk -Number 1 |Initialize-Disk -PartitionStyle GPT
PS C:\Users\Administrator> detach-volume and format to DAX mode
PS C:\Users\Administrator> get-disk -Number 1 | New-Volume -FriendlyName DAX-VOL -DriveLetter J | Format-Volume -fileSystem NTFS -isDAX $true
-----
DriveLetter FileSystemLabel FileSystem DriveType HealthStatus OperationalStatus SizeRemaining Size
-----
J 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 21. 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

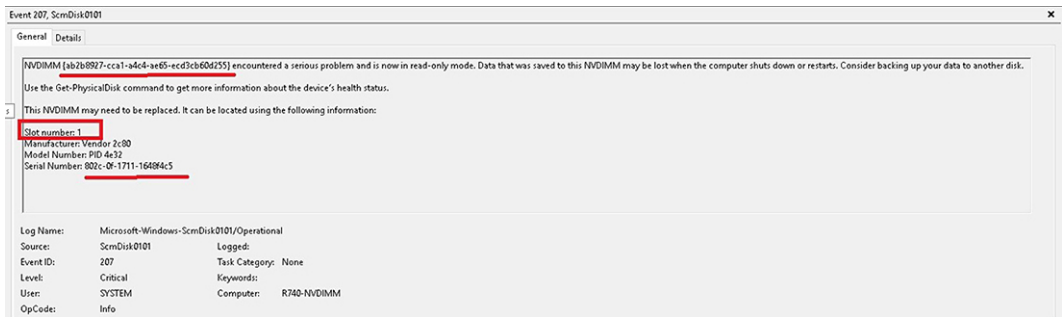


Figure 22. Operational and Diagnostics Logging information: NVDIMM-N Disk

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

```

PS C:\Users\Administrator> Get-PhysicalDisk -FriendlyName "802c-0f-3721-1648f4c5" | fl FriendlyName, uniqueID, PhysicalLocation, SerialNumber, healthStatus, OperationalStatus, OperationalDetails
FriendlyName      uniqueID          PhysicalLocation      SerialNumber          HealthStatus OperationalStatus      OperationalDetails
-----
Vendor 2c80 PID 4e32 (4b2b8927-ccaf-ae04-ae05-ec4fcb04255) DIMM Socket 0, Slot 1 802c-0f-3721-1648f4c5 Unhealthy Device Hardware Error (Lost Data Persistence, Threshold Exceeded, W039PLN Error)

PS C:\Users\Administrator> Get-PhysicalDisk -SerialNumber 802c-0f-3721-1648f4c5 | fl FriendlyName, uniqueID, PhysicalLocation, SerialNumber, healthStatus, OperationalStatus, OperationalDetails
FriendlyName      uniqueID          PhysicalLocation      SerialNumber          HealthStatus OperationalStatus      OperationalDetails
-----
Vendor 2c80 PID 4e32 (4b2b8927-ccaf-ae04-ae05-ec4fcb04255) DIMM Socket 0, Slot 1 802c-0f-3721-1648f4c5 Unhealthy Device Hardware Error (Lost Data Persistence, Threshold Exceeded, W039PLN Error)

```

Figure 23. NVDIMM-N physical location

# 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.

Topics:

- [Identify and Configure PMEM — Persistent Memory Device](#)
- [Installation](#)
- [Verify Existing Filesystem](#)
- [Read-Only Mode NVDIMM-N](#)
- [Interleave](#)
- [Management Utility](#)
- [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 7.3/7.4 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 RHEL7.3.

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](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 7.3 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 7.3/7.4 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.

## Linux Errata

### mdadm

MD software RAID Volumes can be created while NVDIMM-Ns are in Read/Write mode. If NVDIMM-Ns are set to Read-Only Mode in BIOS after reboot, you can still mount the md array on a read/write file systems. However, your written data will not be saved after shutdown.

Workaround:

- 1 Avoid setting NVDIMM-N as Read Only Mode in BIOS.
- 2 Set up /dev/mdX as read-only by

```
$ mdadm --readonly /dev/mdX
```