



Hewlett Packard
Enterprise

Intel Optane persistent memory 200 series for HPE User Guide

Part Number: 30-B41F6FB9-002

Published: November 2021

Edition: 2

Intel Optane persistent memory 200 series for HPE User Guide

Abstract

This document includes installation, maintenance, and configuration information for the Intel Optane persistent memory 200 series for HPE and is for the person who installs, administers, and troubleshoots HPE ProLiant Gen10 Plus and HPE Synergy systems. Hewlett Packard Enterprise assumes you are qualified in the servicing of computer equipment and trained in recognizing hazards in products with hazardous energy levels.

Part Number: 30-B41F6FB9-002

Published: November 2021

Edition: 2

© Copyright 2021 Hewlett Packard Enterprise Development LP

Notices

The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Confidential computer software. Valid license from Hewlett Packard Enterprise required for possession, use, or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Links to third-party websites take you outside the Hewlett Packard Enterprise website. Hewlett Packard Enterprise has no control over and is not responsible for information outside the Hewlett Packard Enterprise website.

Acknowledgments

Intel[®], Optane[™], and Xeon[®], are trademarks of Intel Corporation in the U.S. and other countries.

Linux[®] is the registered trademark of Linus Torvalds in the U.S. and other countries.

Red Hat[®] Enterprise Linux[®] is a registered trademark of Red Hat, Inc. in the United States and other countries.

SUSE[®] is a registered trademark of SUSE LLC in the United States and other countries.

Windows Server[®] is either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

VMWare vSphere[®] is a registered trademark of VMware, Inc. in the United States and/or other jurisdictions.

Table of contents

1 Introduction

1.1 Intel Optane persistent memory 200 series for HPE

1.1.1 Persistent memory modes

1.1.2 Memory cache ratios

2 Security features

2.1 Passwords

2.2 Encryption

2.3 Sanitization

2.4 Signed firmware

2.5 Firmware rollback protection

3 Component identification

3.1 Intel Optane persistent memory 200 series for HPE label identification

4 Installation

4.1 System requirements

4.2 Memory population information

4.3 Persistent memory module handling guidelines

4.4 Installing a DIMM or persistent memory module

5 Configuring the system

5.1 Configuration overview

5.1.1 Configuration tools

5.2 Setting the goal configuration

5.2.1 Setting the goal configuration using UEFI System Utilities

5.2.2 Setting the goal configuration with ipmctl

5.2.3 Setting a goal configuration using HPE iLO RESTful API

5.3 Creating namespaces

5.3.1 Creating namespaces using UEFI System Utilities

5.3.2 Creating namespaces using ipmctl

5.3.3 Creating namespaces using ndctl in Linux

5.4 Enabling key management

5.4.1 Encrypting persistent memory modules with local key management

5.4.2 Encrypting persistent memory modules with remote key management

5.4.2.1 Using a key management server

5.4.2.1.1 Supported key managers

5.4.2.1.2 Configuring key manager servers

5.4.2.1.2.1 Key manager server options

5.4.2.1.3 Adding key manager configuration details

5.4.2.1.3.1 Key manager configuration details

5.4.2.1.4 Testing the key manager configuration

5.4.2.1.5 Viewing key manager events

5.4.2.1.6 Clearing the key manager log

5.5 Other BIOS/Platform Configuration (RBSU) options

6 Management tools

6.1 Managing Intel Optane persistent memory for HPE

6.2 UEFI System Utilities

6.2.1 Changing the goal configuration using UEFI System Utilities

6.2.2 Deleting the goal configuration using UEFI System Utilities

6.2.3 Changing persistent memory module passwords

6.2.4 Viewing the status of persistent memory modules

6.2.5 Changing the key management mode

6.2.6 Disabling key management

6.2.7 Disabling encryption for a persistent memory module

6.2.8 Changing Performance Options using UEFI System Utilities

6.3 HPE iLO RESTful API

6.3.1 HPE iLO RESTful API overview

6.3.1.1 Data model overview

6.3.1.1.1 Examples: Retrieving memory resources

6.3.2 Managing Intel Optane persistent memory for HPE with the HPE iLO RESTful API

6.3.3 Persistent memory module provisioning with the HPE iLO RESTful API

6.3.4 Examples: Provisioning persistent memory modules

6.3.4.1 Using RESTful Interface Tool rawpost to configure 100% App Direct interleaved

6.3.4.2 Using python to configure 100% App Direct interleaved

6.3.4.3 Using Postman to configure 100% App Direct interleaved

6.3.5 Example: Managing Intel Optane persistent memory for HPE using curl

6.4 RESTful Interface Tool

6.4.1 Launching the RESTful Interface Tool

6.4.2 Discovery commands

6.4.2.1 Device discovery

6.4.2.2 Device configuration discovery

6.4.2.3 Persistent interleave regions discovery

6.4.2.4 Persistent memory summary

6.4.3 Configuration commands

6.4.3.1 Show pending configuration

6.4.3.2 Apply predefined configuration

6.4.3.3 Apply a user-defined configuration

6.4.3.4 Clear pending configuration

6.4.3.5 Show recommended configuration

6.5 ipmctl tool

6.5.1 Installing ipmctl for Linux

6.5.2 Showing persistent memory module configurations with ipmctl

6.5.3 Deleting a goal configuration with ipmctl

6.5.4 Deleting namespace with ipmctl

6.5.5 Determining the memory mode with ipmctl

7 Maintenance

7.1 Persistent memory module relocation guidelines

7.2 Manually backing up persistent memory module data

7.3 Removing a DIMM or persistent memory module

7.4 Replacing a system board

7.5 Migrating a persistent memory module

7.5.1 Migrating a persistent memory module encrypted with local key management

7.5.2 Migrating a persistent memory module encrypted with remote key management

7.6 Sanitizing a persistent memory module

7.6.1 Sanitization policies

7.6.2 Sanitization guidelines

7.6.3 Sanitization with UEFI System Utilities

7.6.4 Sanitization with the HPE iLO RESTful API

7.6.5 Sanitization with ipmctl

7.7 Recommissioning a persistent memory module with a lost password

7.8 Updating persistent memory module firmware

8 Linux support of Intel Optane persistent memory 200 series for HPE

8.1 nmem devices

8.1.1 nmem device properties

8.1.2 Listing nmem devices

8.2 Regions

8.2.1 Region properties

8.2.2 Listing regions

8.3 Namespaces

8.3.1 Namespace properties

8.3.2 Creating a namespace

8.3.3 List all namespaces

8.3.4 Changing namespace mode

8.3.5 Deleting namespaces

8.4 Initialization of pmem devices

8.5 Showing the amount of memory in the system

8.6 Filesystems

8.7 I/O statistics

9 VMware support of Intel Optane persistent memory 200 series for HPE

10 Windows Server support of Intel Optane persistent memory 200 series for HPE

11 Troubleshooting

11.1 Known issues

11.1.1 System boot fails due to persistent memory file systems

11.2 Troubleshooting resources

12 Websites

13 Support and other resources

13.1 Accessing Hewlett Packard Enterprise Support

13.2 Accessing updates

13.3 Remote support

13.4 Warranty information

13.5 Regulatory information

13.6 Documentation feedback

Intel Optane persistent memory 200 series for HPE

Intel Optane persistent memory 200 series for HPE offers the flexibility to deploy memory as dense memory (Memory mode) or fast storage (App Direct mode) and enables per-socket memory capacity of up to 4.0 TB. Persistent memory modules, together with traditional volatile DRAM DIMMs, provide fast, high-capacity, cost-effective memory and storage to transform big data workloads and analytics by enabling data to be stored, moved, and processed quickly.

Persistent memory modules use the standard DIMM form factor and are installed alongside DIMMs in a server memory slot. Intel Optane persistent memory 200 series for HPE is designed for use only with third-generation Intel Xeon Scalable processors, and is available in the following capacities:

- 128 GB
- 256 GB
- 512 GB

Persistent memory modes

Intel Optane persistent memory 200 series for HPE can be configured to operate in two modes.

App Direct mode

When configured for App Direct mode, persistent memory modules function as persistent memory.

Memory mode

When configured for Memory mode, persistent memory modules function as volatile system memory while DRAM capacity operates as a cache. For more information, see Memory cache ratios.

Memory cache ratios

Persistent memory modules can be carved into volatile and persistent regions.

For the volatile region, the ratio of volatile memory capacity to DRAM capacity affects performance. The recommended volatile memory capacity to DRAM capacity is between 4:1 and 16:1:

- 4:1—largest cache; most likely to get cache hits.
- 8:1
- 16:1—smallest cache; least likely to get cache hits.

Users are strongly recommended not to configure for Memory Mode when the ratio is at 1:1 or lower due to more severe performance impact.

A message will be logged to the Integrated Management Log (IML) if a cache ratio that is not recommended is used.

The following table contains example configurations, with 100% of the persistent memory module capacity allocated to volatile memory. Cache ratios improve as some of that memory is carved into persistent memory.

Persistent memory module capacity ¹	Persistent memory module configuration	DIMM capacity ¹	DIMM configuration	Ratio
1024 GB	8 x 128 GB	128 GiB	8 x 16 GiB	8:1
		256 GiB	8 x 32 GiB	4:1
		512 GiB	8 x 64 GiB	2:1 ²
		1024 GiB	8 x 128 GiB	1:1 ²
2.0 TB	8 x 256 GB	128 GiB	8 x 16 GiB	16:1
		256 GiB	8 x 32 GiB	8:1
		512 GiB	8 x 64 GiB	4:1
		1024 GiB	8 x 128 GiB	2:1 ²
		2048 GiB	8 x 256 GiB	1:1 ²
4.0 TB	8 x 512 GB	128 GiB	8 x 16 GiB	32:1 ³
		256 GiB	8 x 32 GiB	16:1
		512 GiB	8 x 64 GiB	8:1
		1024 GiB	8 x 128 GiB	4:1
		2048 GiB	8 x 256 GiB	2:1 ²

¹ Capacity per processor.

² Not recommended. No benefit from caching.

³ Not recommended.

Security features

Intel Optane persistent memory for HPE provides a number of features to secure and protect your data:

- Passwords
- Encryption
- Sanitization
- Signed firmware
- Firmware rollback protection

Passwords

Persistent memory modules support password-based locking with a 32 byte binary password. If locked, data on the persistent memory module is not accessible until it is unlocked. If a persistent memory module is locked and the password is lost, the persistent memory module can be sanitized to regain access to the hardware, but it does not provide access to the data.

HPE ProLiant and HPE Synergy Gen10 Plus server products provide two methods for managing persistent memory module passwords:

- Local key management
- Remote key management

Only one key management method can be selected at one time to manage passwords.

Local key management

Local key management is available in servers with HPE Trusted Platform Module (TPM) 2.0 installed. When enabled, the server generates a 32 byte random value to use as a password for each persistent memory module.

Persistent memory module passwords are stored in a flash memory shared by HPE iLO and the system firmware. Each password in the password database is encrypted using the HPE TPM 2.0, a tamper resistant part.

During POST, the server extracts the passwords from the database and unlocks all persistent memory modules. Passwords can be exported to a USB key for migration into another server. This migration file is encrypted with a key generated from a transient password (an ASCII string) that the user must provide. To import the file on another server, the user must enter the same transient password.

This file also serves as a backup to restore the passwords if the server system board fails.

Remote key management

Remote key management is available in servers with the HPE iLO enrolled in and connected to a key management server. Persistent memory module passwords are automatically generated, managed, and stored on a key management server. The remote key management feature requires an HPE iLO Advanced License.

Encryption

Persistent memory modules encrypt all data written to the media using 256 bits XTS-AES algorithms.

For volatile memory regions, the persistent memory module generates a new encryption key at power-up and holds the key in a volatile register, which is lost on power loss. Although the media is naturally persistent, it makes the volatile memory region effectively volatile.

For persistent memory regions, the persistent memory module remembers the encryption key across power cycles, ensuring that the data is still accessible.

- If the persistent memory module password is enabled, the encryption key is itself encrypted by another key derived from the password. This "key wrapping" prevents an unauthorized user from reading the media content.

The encryption key is available only if the proper password is presented to the persistent memory module, and the persistent memory module holds it in volatile registers.

- If the persistent memory module password is not enabled, the encryption key is stored on the media. Although the user data is encrypted, an unauthorized user might be able to decrypt it.

Both cases facilitate an "instant secure erase" sanitization feature. By changing the encryption key, all data becomes undecipherable.

With a persistent memory module password, the encryption key is never exposed. Without a password, the encryption key is never exposed to the system, but an unauthorized user with physical access to the medium might have retrieved the encryption key before erase and used it later. That tampering would be physically evident.

Sanitization

Media sanitization is defined by NIST SP800-88 Guidelines for Media Sanitization (Rev 1, Dec 2014) as "a general term referring to the actions taken to render data written on media unrecoverable by both ordinary and extraordinary means."

The specification defines the following levels:

- **Clear:** Overwrite user-addressable storage space using standard write commands; might not sanitize data in areas not currently user-addressable (such as bad blocks and over-provisioned areas).
- **Purge:** Overwrite or erase all storage space that might have been used to store data using dedicated device sanitize commands, such that data retrieval is "infeasible using state-of-the-art laboratory techniques."
- **Destroy:** Ensure that data retrieval is "infeasible using state-of-the-art laboratory techniques" and render the media unable to store data (such as disintegrate, pulverize, melt, incinerate, or shred).

Intel Optane persistent memory for HPE supports the purge level using a cryptographic erase technique and an overwrite technique.

HPE ProLiant and HPE Synergy Gen10 Plus server products support sanitizing persistent memory modules during POST. Use the RESTful Interface Tool or UEFI System Utilities to schedule sanitization on the next boot.

Cryptographic erase technique

This technique allows "instant secure erase," scrambling all persistent contents on the persistent memory module in less than a second, no matter what the capacity. Persistent media reads back random-looking data (data encrypted with a now-lost key).

It also ensures that data in bad or worn-out parts of the media is undecipherable, even if these areas could be accessed. This technique is stronger than overwrite techniques, which might be unable to overwrite such areas.

Sanitization is allowed under this technique, even if the persistent memory module is locked with a password. This ensures that the persistent memory module hardware is usable even if you forget the password.

Overwrite technique

Persistent memory modules also support an overwrite technique. It complies with the clear level by default, but it also complies with the purge level if it is successful at overwriting bad and worn-out parts of the media.

If encryption is enabled (a password has been set on the persistent memory module), this operation overwrites the media with encrypted zeroes. If encryption is not enabled or if the `CryptoEraseOverwrite` command is used, the operation overwrites the media with zeroes.

NIST SP800-88 Guidelines for Media Sanitization (Rev 1, Dec 2014) is available for download from the NIST website (<http://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-88r1.pdf>).

Signed firmware

Persistent memory module firmware images are cryptographically signed. The image includes a cryptographic hash value (for example, SHA-256) that is encrypted using RSA public-private key encryption.

The hash value is encrypted using the private key. The persistent memory module decrypts the hash value using the public key.

The private key is kept in a FIPS 140-2 level 3 or level 4 (tamper-resistant or tamper-proof) code-signing appliance. The appliance enforces access controls so that it only accepts authentic, signed images. The persistent memory module rejects images that do not decrypt correctly.

Firmware rollback protection

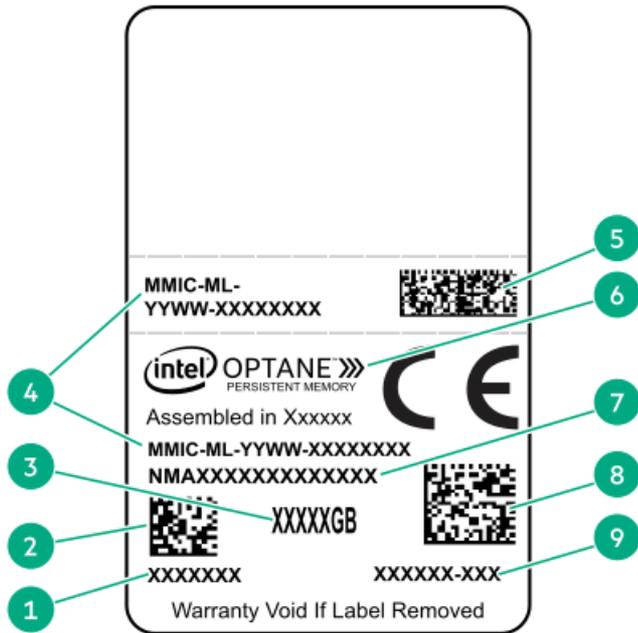
Firmware images are identified by version number, similar to 01.02.03.0405.

The second field (for example, 02) represents the security version number. This number is incremented in releases that contain security improvements. Persistent memory modules cannot accept firmware images with an older security version number than the image that is running. This protection prevents the firmware from rolling back to an earlier image that might contain exploitable features. Firmware roll-back can only be done with individual binary files via BIOS embedded FW update applications.

Component identification



Intel Optane persistent memory 200 series for HPE label identification



Item	Description	Example
1	Work Order Number	XXXXXXX
2	Work Order Number Barcode	XXXXXXX
3	Capacity	128 GB 256 GB 512 GB
4	Unique ID number	8089-A2-1802-1234567A
5	Serial Number and Part Number Barcode	S8089A218040000168APNMAXXXXXXXXXXXXX
6	Product Name	Intel® Optane™ persistent memory
7	Part Number	1234567A
8	Serial Number Barcode	8089-A2-1802-1234567A
9	PBA Number	XXXXXX-XXX

For more information about product features, specifications, options, configurations, and compatibility, see the product QuickSpecs on the Hewlett Packard Enterprise website (<https://www.hpe.com/support/persistentmemoryQS>).



System requirements

i IMPORTANT:

Hewlett Packard Enterprise recommends that you implement best-practice configurations such as clustered configurations for high availability (HA).

The following hardware components are required:

- HPE DDR4 Standard Memory RDIMMs or LRDIMMs
- Intel Optane persistent memory 200 series for HPE
- Third-generation Intel Xeon Scalable processors

Supported firmware versions:

- System ROM version 1.40 or later
- Server Platform Services (SPS) Firmware version 04.04.04.053
- HPE iLO 5 Firmware version 2.44
- HPE Innovation Engine Firmware version 1.0.0.20 or later

Download the required firmware and drivers from the Hewlett Packard Enterprise website (<https://www.hpe.com/info/persistentmemory>).

Supported operating systems:

- Windows Server 2016 with persistent memory drivers from Hewlett Packard Enterprise
- Windows Server 2019
- Red Hat Enterprise Linux 7.9 and later
- Red Hat Enterprise Linux 8.2 and later
- SUSE Linux Enterprise Server 12 SP5 and later
- SUSE Linux Enterprise Server 15 SP2 and later
- VMware vSphere 7.0 U2 and later
- VMware vSphere 6.7 U3 with P03

Hardware and licensing requirements for optional encryption of the persistent memory modules:

- HPE TPM 2.0 (local key encryption)
- HPE iLO Advanced License (remote key encryption)
- Key management server (remote key encryption)

Memory population information

DIMMs and persistent memory modules are installed in specific configurations based on the workload requirements of the server. Supported configurations are optimized for persistent memory capacity, volatile memory capacity, and performance.

- Persistent memory capacity—the available capacity is equal to the persistent memory module capacity.
- Volatile memory capacity:
 - App Direct mode—the volatile capacity is equal to the DRAM capacity (the capacity of all the non-persistent memory modules installed).
- Memory tier capacity—the memory tier capacity is the total capacity of all installed memory (DRAM and persistent memory modules).

ⓘ IMPORTANT:

If the installed memory exceeds the processor capacity, the system will map out all but one DIMM channel and operate in App Direct mode. A message will be logged to the IML for exceeding capacity. To resolve the issue, remove the memory that exceeds the processor capacity.

- Performance:
 - Uses all channels to efficiently utilize processor resources.
 - Memory mode—more regular DIMMs provide a better cache ratio.

Before beginning the installation, review the memory population guidelines on the [Hewlett Packard Enterprise website](#).

Persistent memory module handling guidelines

 **CAUTION:**

Failure to properly handle persistent memory modules can damage the component and the system board connector.

When handling a persistent memory module, observe the following guidelines:

- Avoid electrostatic discharge.
- Always hold persistent memory modules by the side edges only.
- Avoid touching the connectors on the bottom of the persistent memory module.
- Never wrap your fingers around a persistent memory module.
- Avoid touching the components on the sides of the persistent memory module.
- Never bend or flex the persistent memory module.

When installing a persistent memory module, observe the following guidelines:

- Before seating the persistent memory module, open the persistent memory module slot and align the persistent memory module with the slot.
- To align and seat the persistent memory module, use two fingers to hold the persistent memory module along the side edges.
- To seat the persistent memory module, use two fingers to apply gentle pressure along the top of the persistent memory module.

For more information, see the Hewlett Packard Enterprise website (<https://www.hpe.com/support/DIMM-20070214-CN>).

Installing a DIMM or persistent memory module

For server-specific steps used in this procedure, see the server user guide on the Hewlett Packard Enterprise website:

- HPE ProLiant Gen10 Plus servers (<https://www.hpe.com/info/proliantgen10plus-docs>)
- HPE Synergy Gen10 Plus compute modules (<https://www.hpe.com/info/synergy-docs>)

IMPORTANT:

Hewlett Packard Enterprise recommends that you implement best-practice configurations such as clustered configurations for high availability (HA).

Prerequisites

Before beginning the installation, review the memory population guidelines on the [Hewlett Packard Enterprise website](#).

Procedure

1. Observe the following alerts:

CAUTION:

DIMMs and persistent memory modules are keyed for proper alignment. Align notches on the DIMM or persistent memory module with the corresponding notches in the slot before installing the component. Do not force the DIMM or persistent memory module into the slot. When installed properly, not all DIMMs or persistent memory modules will face in the same direction.

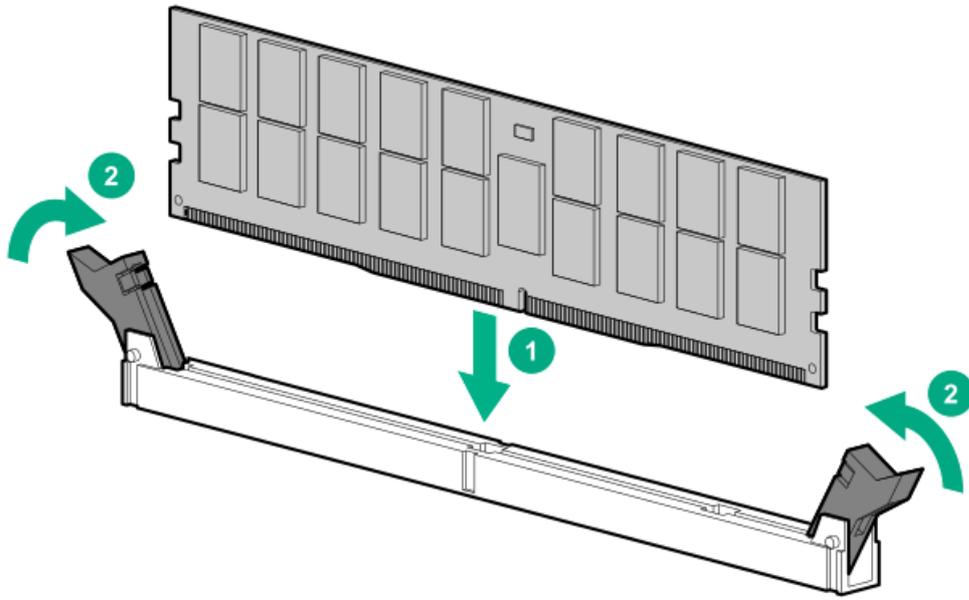
CAUTION:

Electrostatic discharge can damage electronic components. Be sure you are properly grounded before beginning this procedure.

CAUTION:

Failure to properly handle persistent memory modules can damage the component and the system board connector.

2. Power down the server:
 - a. Shut down the OS as directed by the OS documentation.
 - b. To place the server in standby mode, press the Power On/Standby button. When the server enters standby power mode, the system power LED changes to amber.
 - c. Disconnect the power cords (rack and tower servers).
3. Do one of the following:
 - Extend the server from the rack.
 - Remove the server from the rack, if necessary.
 - Remove the server or server blade from the enclosure.
4. Place the server on a flat, level work surface.
5. Remove the access panel.
6. Remove all components necessary to access the DIMM slots.
7. Install the DIMM or persistent memory module.



8. Install any components removed to access the DIMM slots.
9. Install the access panel.
10. Slide or install the server into the rack.
11. If removed, reconnect all power cables.
12. Power up the server.

Configuration overview

Configure Intel Optane persistent memory for HPE by:

1. Setting a "goal configuration," which defines the regions of volatile memory and persistent memory on a persistent memory module.
2. Creating namespaces on top of the resulting persistent regions.
3. (Optional) Enabling local or remote key management.
4. (Optional) Encrypting the persistent memory modules.

i IMPORTANT:

Always follow recommendations from your software application provider for high-availability best practices to ensure maximum uptime and data protection.

Configuration tools

A number of tools are available to configure and maintain Intel Optane persistent memory for HPE:

Embedded tools

- UEFI System Utilities
- ipmctl tool (under the UEFI Shell)

REST/iLO-based tools

- HPE iLO RESTful API
- RESTful Interface Tool

OS-based tools

- Windows PowerShell cmdlets
- ipmctl tool (Linux and Windows)
- ndctl tool (Linux)

Setting the goal configuration

Goal configurations that define regions of volatile memory and persistent memory are stored in the metadata of the persistent memory modules. Because persistent memory modules are on the system memory bus, changing the goal configuration requires a system reboot. During the next boot, system firmware detects the goal configuration request and reconfigures the persistent memory modules.

Goal configurations must adhere to the recommended memory cache ratios. Selecting a non-recommended ratio will generate messages in the IML.

i IMPORTANT:

When data must be preserved, Hewlett Packard Enterprise strongly recommends that you perform a manual backup of all user data on the persistent memory modules before changing the goal configuration or performing relocation procedures.

i IMPORTANT:

Always follow recommendations from your software application provider for high-availability best practices to ensure maximum uptime and data protection.

Setting the goal configuration using UEFI System Utilities

i IMPORTANT:

Be sure to observe all pop-up messages displayed in UEFI System Utilities that pertain to persistent memory. Failure to follow the instructions in these messages might cause persistent memory data loss.

Procedure

1. From the System Utilities screen, select **System Configuration > BIOS/Platform Configuration (RBSU) > Memory Options**.

2. Confirm the following default settings:

- **Maximum Memory Bus Frequency**—Auto
- **Memory Patrol Scrubbing**—Enabled
- **Memory Remap**—No Action

3. Select **Persistent Memory Options**, and confirm the following selection:

Persistent Memory Address Range Scrub—Enabled.

4. Select **PMM Options > Goal Configuration Options**.

Goal Configuration Options displays the most recent configuration settings, but not necessarily the active configuration. The configuration settings defined in this screen are only applied at the next server reboot.

5. Make the following selections:

- **Volatile Memory Capacity**—% of persistent memory module capacity providing volatile memory:
 - **Memory mode**—Select **100%**.
 - **App Direct mode**—Select **0%**.
- **Persistent Memory Interleaving**—Enabled or Disabled.

6. Select **Apply Goal Configuration**.

Goal configuration settings will be applied at the next reboot.

7. Select **PMM Options > Security Options > Security Freeze Lock**—Disabled.

8. Confirm your selections.

9. To save your changes, press the **F12** key.

10. To commit the goal configuration and persistent memory options, reboot the server.

Setting the goal configuration with ipmctl

The ipmctl tool can be run in the UEFI command line, Windows OS, or in Linux:

```
create
[-dimm [(DimmIDs)]]
-goal
[-socket (SocketIDs)]
[MemoryMode=(0|%)]
[PersistentMemoryType=(AppDirect|AppDirectNotInterleaved)]
```

Sample goal configurations

Command	Description
<code>ipmctl create -goal</code>	Defaults to 100% interleaved persistent memory
<code>ipmctl create -goal MemoryMode=0 Reserved=100</code>	100% unconfigured
<code>ipmctl create -goal MemoryMode=100</code>	100% volatile memory
<code>ipmctl create -goal MemoryMode=0 PersistentMemoryType=AppDirect</code>	100% interleaved persistent memory
<code>ipmctl create -goal MemoryMode=0 PersistentMemoryType=AppDirectNotInterleaved</code>	100% non-interleaved persistent memory
<code>ipmctl create -goal MemoryMode=80 PersistentMemoryType=AppDirect</code>	<ul style="list-style-type: none">• 80% volatile memory• 20% interleaved persistent memory

The values must adhere to the recommended memory cache ratio. Selecting a nonrecommended ratio might impact system performance and will generate messages in the IML.

Setting a goal configuration using HPE iLO RESTful API

The HPE iLO RESTful API can be accessed using a number of tools. Hewlett Packard Enterprise recommends using the RESTful Interface Tool.

The `rawpost` command takes in a JSON file. The following example displays the JSON file and a batch script to use the RESTful Interface Tool to configure a server for 100% AppDirect with interleaving enabled.

Memorychunk-rawpost.txt

```
{
  "path": "/redfish/v1/Systems/1/MemoryDomains/PROC1MemoryDomain/MemoryChunks",
  "body": {
    "AddressRangeType": "PMEM",
    "Oem": {
      "Hpe": {
        "MemoryChunkSizePercentage": 100
      }
    },
    "InterleaveSets": [{
      "Memory": {
        "@odata.id": "/redfish/v1/Systems/1/Memory/proclimm6/"
      }
    }, {
      "Memory": {
        "@odata.id": "/redfish/v1/Systems/1/Memory/proclimm7/"
      }
    }
  ]
}
```

Windows batch script

```
@echo off

set argC=0
for %%x in (*) do Set /A argC+=1
if %argC% LSS 3 goto :failCondition
goto :main

:failCondition
@echo Usage:
@echo ilorest-script-memory-remote.bat [URL] [USERNAME] [PASSWORD]
goto :EOF

:main
@echo Logging in...
ilorest.exe --nologo login %1 -u %2 -p %3
@echo rawpost to Memory Chunk collection...
ilorest.exe --nologo rawpost memorychunk-rawpost.txt
@echo Note: Status of 202 is success.
```

Creating namespaces



Creating namespaces using UEFI System Utilities

IMPORTANT:

Be sure to observe all pop-up messages displayed in UEFI System Utilities that pertain to persistent memory. Failure to follow the instructions in these messages might cause persistent memory data loss.

NOTE:

If you are using Intel Optane persistent memory for HPE with VMware vSphere, creating namespaces is not required. VMware vSphere automatically creates namespaces upon rebooting.

Namespaces define persistent memory regions on the persistent memory modules.

Procedure

1. From the System Utilities screen, select **System Configuration > BIOS/Platform Configuration (RBSU) > Memory Options > Persistent Memory Options**.
2. Select **PMM Options > Advanced Options**, and make the following selections:
 - **Apply Default Namespaces**—Enabled or Disabled.

This selection creates namespace metadata on the next boot for interleave sets that do not already have one. For Linux systems, Hewlett Packard Enterprise recommends using OS tools such as ndctl for this purpose.
 - **Delete Namespaces**—Immediately deletes any active namespaces.
3. To save your changes, press the **F12** key.
4. To commit the goal configuration and persistent memory options, reboot the server.

Creating namespaces using ipmctl

Default namespaces can be created by using the `ipmctl` tool in the UEFI command line.

```
Shell> ipmctl create -namespace -region id
```

Creating namespaces using ndctl in Linux

Linux supports multiple namespace modes, which can be created by using the `ndctl` command in Linux.

```
ndctl create-namespace [<options>]
```

For more information on using the `ndctl` command to create or change namespaces, see the following:

- [Namespaces](#)
- The ndctl documentation at <https://docs.pmem.io/ndctl-user-guide/>

Enabling key management



NOTE:

The ipmctl OS tool does not support key management features. To enable key management and enable or disable encryption of persistent memory modules, use the following procedure in the UEFI System Utilities.

Prerequisites

Before enabling local or remote key management, confirm the following:

- The goal configuration is set and Intel Optane persistent memory for HPE is configured based on your server workload requirements.
- For local key management:
 - The server has an HPE TPM 2.0 installed.
 - The HPE TPM 2.0 is active and not hidden.
 - The server is configured for UEFI Boot Mode (local key management is not supported in Legacy Boot Mode).
- For remote key management:
 - The HPE iLO is enrolled in and connected to a key management server.
 - The server has an HPE iLO Advanced License.

For more information, see [Using a key management server](#).

Procedure

1. From the System Utilities screen, select **System Configuration > BIOS/Platform Configuration (RBSU) > Server Security > Device Encryption Options**.
2. Select the Key Management setting:
 - **Disabled**—The default setting. Key management is disabled.
 - **Local**—Enables local key management. The password used for encryption is stored locally on the server.
HPE TPM 2.0 must be installed to view and select this setting.
 - **Remote**—Enables remote key management. The password used for encryption is stored on a remote key server.
HPE iLO must be enrolled in and connected to a key manager to view and select this setting.
3. Press the **F12** key to save and exit.
4. Reboot the server.
5. During POST, press the **F9** key to enter System Utilities.
6. Do one of the following:
 - [Encrypt persistent memory modules with local key management](#)
 - [Encrypt persistent memory modules with remote key management](#)

Encrypting persistent memory modules with local key management

Prerequisites

Local key management must be enabled. For more information, see [Enabling key management](#).

Procedure

1. During POST, press the F9 key to enter System Utilities.
2. From the System Utilities screen, select System Configuration > BIOS/Platform Configuration (RBSU) > Server Security > Device Encryption Options > Device Encryption Settings > Unencrypted Devices.
3. Make the following selections:

- **Select Device**—Select the specific persistent memory module to encrypt.
- **Select Operation**—Select **Enable Encryption**.

4. Select the Passphrase Type:

- **Auto**—The system automatically generates a 32 byte random password. Hewlett Packard Enterprise recommends using system-generated passwords as a best practice.
- **Manual**—Enter a 32 byte password manually.

5. Select Start Operation.

The persistent memory module is now encrypted.

6. To encrypt another persistent memory module, select it from the Select Device menu.
7. To enable encryption for each individual persistent memory module, repeat the process.
8. View the status of the encrypted persistent memory module.

For more information, see [Viewing the status of persistent memory modules](#).

9. Hewlett Packard Enterprise recommends exporting the password database to a USB device for backup purposes:
 - a. From the System Utilities screen, select System Configuration > BIOS/Platform Configuration (RBSU) > Server Security > Device Encryption Options > Device Encryption Migration Options > Device Encryption Export Options.
 - b. Provide a password in the Transient Passphrase field.

This password protects the exported file and must be entered when recovering the encrypted persistent memory modules after relocation.
 - c. Select Select File, and browse to a location on the USB key.
 - d. Select Export Encryption Settings to create and export the file.

Encrypting persistent memory modules with remote key management

When Remote Key Management is enabled, persistent memory module passwords are automatically generated, stored, and managed on the key management server.

Prerequisites

- The HPE iLO must be enrolled in and connected to a key management server and have an HPE iLO Advanced License. For more information, see [Using a key management server](#).
- Remote key management must be enabled. For more information, see [Enabling key management](#).

Procedure

1. During POST, press the F9 key to enter System Utilities.
2. From the System Utilities screen, select System Configuration > BIOS/Platform Configuration (RBSU) > Server Security > Device Encryption Options > Device Encryption Settings > Unencrypted Devices.
3. Make the following selections:
 - **Select Device**—Select the specific persistent memory module to encrypt.
 - **Select Operation**—Select **Enable Encryption**.
4. Select Start Operation.
The persistent memory module is now encrypted.
5. To encrypt another persistent memory module, select it from the Select Device menu.
6. To enable encryption for each individual persistent memory module, repeat the process.

Using a key management server

iLO 5 supports key managers, which can be used in conjunction with Intel Optane persistent memory for HPE. The UEFI-managed encryption allows data-at-rest encryption for persistent memory modules using 256-bit XTS-AES algorithms.

A key manager generates, stores, serves, controls, and audits access to data encryption keys. It enables you to protect and preserve access to business-critical, sensitive, data-at-rest encryption keys.

iLO manages the key exchange between the key manager and the other products. iLO uses a unique user account based on its own MAC address for communicating with the key manager. For the initial creation of this account, iLO uses a deployment user account that pre-exists on the key manager with administrator privileges. For more information about the deployment user account, see the key manager documentation.

Supported key managers

iLO supports the following key managers:

- Utimaco Enterprise Secure Key Manager (ESKM) 4.0 and later

ESKM 5.0 or later is required when the FIPS security state is enabled.

 **CAUTION:**

If you use ESKM, ensure that you install the software update that includes updated code signing certificates. If you do not install the required update, your ESKM will enter an error state when restarted after January 1, 2019. For more information, see the [ESKM documentation](#).

- Thales TCT KeySecure for Government G350v (previously known as SafeNet AT KeySecure G350v 8.6.0)
- Thales KeySecure K150v (previously known as SafeNet KeySecure 150v 8.12.0)
- Thales CipherTrust Manager 2.2.0, K170v (virtual) and K570 (physical) appliances

 **NOTE:**

Using a key manager is not supported when iLO is configured to use the CNSA security state.

Configuring key manager servers

Prerequisites

- Configure iLO Settings privilege
- A license that supports this feature is installed. For information about the available license types and the features they support, see the licensing documentation at the following website: <https://www.hpe.com/support/ilo-docs>.
- iLO is not configured to use the CNSA security state.

Procedure

1. Click Administration in the navigation tree, and then click the Key Manager tab.
2. Click  in the Key Manager Servers section.

The Edit Key Manager Server Settings page opens.

3. Enter the following information:
 - Primary Key Server Address
 - Primary Key Server Port
 - Secondary Key Server Address
 - Secondary Key Server Port
4. (Optional) To check for server redundancy in configurations with a primary and secondary key server, enable the Require Redundancy option.

Hewlett Packard Enterprise recommends enabling this option.

5. Click OK.

For more information about Thales CipherTrust Manager 2.2.0, see [Remote Key Manager Support for Cipher Trust Manager Configuration guide](#).

Key manager server options

Primary Key Server Address

The primary key server hostname, IP address, or FQDN. This string can be up to 79 characters long.

Primary Key Server Port

The primary key server port.

Secondary Key Server Address

The secondary key server hostname, IP address, or FQDN. This string can be up to 79 characters long.

Secondary Key Server Port

The secondary key server port.

Require Redundancy

When this option is enabled, iLO verifies that the encryption keys are copied to both of the configured key servers.

When this option is disabled, iLO will not verify that encryption keys are copied to both of the configured key servers.

Hewlett Packard Enterprise recommends enabling this option.

Adding key manager configuration details

Prerequisites

- Configure iLO Settings privilege
- A license that supports this feature is installed. For information about the available license types and the features they support, see the licensing documentation at the following website: <https://www.hpe.com/support/ilo-docs>.
- iLO is not configured to use the CNSA security state.
- At least one key manager server is configured.

Procedure

1. Click Administration in the navigation tree, and then click the Key Manager tab.
2. Click  in the Key Manager Configuration section.

The Edit Key Manager Configuration Settings page opens.

3. Enter the following information in the iLO Account on Key Manager section:
 - Account Group
 - (Optional) Key Manager Local CA Certificate Name

The Account Name value is read-only.

4. Enter the following information in the Key Manager Administrator Account section:
 - Login Name
 - Password

5. Click OK.

iLO sends an information request to the key manager server.

- If the ilo-<iLO MAC address> account name does not exist:
 - The user account you entered in the Key Manager Administrator Account section creates the account name and associates it with the key manager local user and its generated password.
 - The account name is added to the account group you entered in step 3.
- If the ilo-<iLO MAC address> account name exists:
 - The user account you entered in the Key Manager Administrator Account section associates the account name with the key manager local user, and a new password is generated.
 - If the user account you entered in the Key Manager Administrator Account section is not a member of the account group associated with the ilo-<iLO MAC address> account, it is added to the account group.
 - If the ilo-<iLO MAC address> is already a member of a key manager local group, the group you entered in step 3 is ignored. The existing group assignment on the key manager is used, and it is displayed in the iLO web interface. If a new group assignment is needed, you must update the key manager before updating the iLO settings.

If you entered the Key Manager Local CA Certificate Name in step 3, certificate information is listed in the Imported Certificate Details section of the Key Manager page.

Key manager configuration details

Account Name

The listed iLO Account on Key Manager account name is ilo-<iLO MAC address>. The account name is read-only and is used when iLO communicates with the key manager.

Account Group

The local group created on the key manager for use with iLO user accounts and the keys iLO imports into the key manager. When keys are imported, they are automatically accessible to all devices assigned to the same group.

See the Secure Encryption installation and user guide for more information about groups and their use with key management.

Key Manager Local CA Certificate Name

To ensure that iLO is communicating with a trusted key manager server, enter the name of the local certificate authority certificate in the key manager. It is typically named Local CA and is listed in the key manager under local CAs. iLO will retrieve the certificate and use it to authenticate the key manager servers for all future transactions.

Secure Encryption does not support using a third-party trusted or intermediate CA.

Login Name

The local user name with administrator permissions that is configured on the key manager. This user name is the key manager deployment user.

The deployment user account must be created before you add key manager configuration details in iLO.

Password

The password for the local user name with administrator permissions that is configured on the key manager.



Testing the key manager configuration

To verify the configured settings, test the key manager configuration. The following tests are attempted:

- Confirm that the key manager software version is compatible with iLO.
- Connect to the primary key manager server (and secondary key manager server, if configured) by using TLS.
- Authenticate to the key manager by using the configured credentials and account.

Prerequisites

- A license that supports this feature is installed. For information about the available license types and the features they support, see the licensing documentation at the following website: <https://www.hpe.com/support/ilo-docs>.
- A key manager is set up and the key manager configuration is complete in iLO.

Procedure

1. Click Administration in the navigation tree, and then click the Key Manager tab.
2. Click .

The test results are displayed in the Key Manager Events table. A success or failure message is displayed at the top of the iLO web interface window.

Viewing key manager events

Prerequisites

A license that supports this feature is installed. For information about the available license types and the features they support, see the licensing documentation at the following website: <https://www.hpe.com/support/ilo-docs>.

Procedure

1. Click Administration in the navigation tree, and then click the Key Manager tab.
2. Scroll to the Key Manager Events section.

Each event is listed with a time stamp and description.

Clearing the key manager log

Prerequisites

- Configure iLO Settings privilege
- A license that supports this feature is installed. For information about the available license types and the features they support, see the licensing documentation at the following website: <https://www.hpe.com/support/ilo-docs>.

Procedure

1. Click Administration in the navigation tree, and then click the Key Manager tab.
2. Click Clear Key Manager Log.
iLO prompts you to confirm the request.
3. Click Yes, clear.



Other BIOS/Platform Configuration (RBSU) options

When persistent memory modules are installed, the following BIOS/Platform Configuration (RBSU) settings are not applicable to persistent memory modules and might not be supported or are supported only when set to their default values:

- **Advanced Memory Protection**—Persistent memory modules are disabled if the configuration is not set to Advanced ECC. When Advanced Memory Protection is set to **Advanced ECC Support**, Advanced Memory Protection is hidden in the menu. Advanced ECC is supported in MM mode (supported in AD before BIOS 1.50). HPE Fast Fault Tolerance (also called ADDDC) is supported in AD mode after BIOS 1.50.
 - UEFI System Utilities: System Configuration > BIOS/Platform Configuration (RBSU) > Memory Options > Advanced Memory Protection
 - iLO RESTful API property name: `AdvancedMemProtection`
- **Maximum Memory Bus Frequency**—This option is enabled by default when persistent memory modules are installed. It enables the system to run memory at a lower maximum speed than what is supported by the installed processor and DIMM configuration.
 - UEFI System Utilities: System Configuration > BIOS/Platform Configuration (RBSU) > Memory Options > Maximum Memory Bus Frequency
 - iLO RESTful API property name: `MaxMemBusFreqMHz`
- **Memory Patrol Scrubbing**—This option is enabled by default when persistent memory modules are installed. This option corrects memory soft errors. Over the length of the system runtime, it reduces the risk of producing multibit and uncorrectable errors.
 - UEFI System Utilities: System Configuration > BIOS/Platform Configuration (RBSU) > Memory Options > Memory Patrol Scrubbing
 - iLO RESTful API property name: `MemPatrolScrubbing`
- **Memory Mirroring Mode**—This option is not supported when persistent memory modules are installed.
 - UEFI System Utilities: System Configuration > BIOS/Platform Configuration (RBSU) > Memory Options > Memory Mirroring Mode
 - iLO RESTful API property name: `MemMirrorMode`
- **Memory Refresh Rate**—This option controls the refresh rate of the memory controller and might affect the performance and resilience of the server memory. Hewlett Packard Enterprise recommends that you leave this setting in the default state unless indicated in other documentation for this server.

For optimal power consumption and performance, Hewlett Packard Enterprise recommends that you select **1x Refresh**.

- UEFI System Utilities: System Configuration > BIOS/Platform Configuration (RBSU) > Memory Options > Memory Refresh Rate
- iLO RESTful API property name: `MemRefreshRate`
- **Sub-NUMA Clustering**—This option is not supported and is automatically set to Disabled when persistent memory modules are installed.
 - UEFI System Utilities: System Configuration > BIOS/Platform Configuration (RBSU) > Power and Performance Options > Sub-NUMA Clustering
 - iLO RESTful API property name: `SubNumaClustering`
- **eADR (enhanced Asynchronous DRAM Refresh)**—This option is disabled by default when persistent memory modules are installed. The eADR is a platform feature with the hardware requirements including the third generation Intel Xeon Scalable processors and Intel Optane PMem 200 Series memory modules. It provides data persistence when power failure occurs by flushing data from IIO Caches + CPU Caches + iMC WPQ to persistent memory modules. The eADR should only be enabled in the RBSU in the events of CF9 warm reset and cold reset.

The `ipmi_poweroff` module, present in all Linux kernels since 2.6.13, has the capability to override the ACPI power control method with the IPMI chassis power off command. This IPMI mechanism may interrupt eADR operations leading to unwanted dirty-shutdown events. Intel recommends disabling the `ipmi_poweroff` module, or otherwise accounting for the fact that `ipmi_poweroff` usage may trigger a dirty-shutdown event.

Disable the `ipmi_poweroff` module in all Linux kernels since 2.6.13 by command `modprobe -r ipmi_poweroff`. See the `modprobe.d` Linux man page for details on how to disable module.

- UEFI System Utilities: System Configuration > BIOS/Platform Configuration (RBSU) > Memory Options > Persistent Memory Options > PMEM Options > Performance Options > eADR
- iLO RESTful API property name: `eADR`
- **Intel Performance Monitoring Support**—Intel processors include performance counters that software can use to measure DRAM performance (including persistent memory module performance). This option is a monitoring tool, and does not impact performance. For example, the Intel Processoe Counter Monitor (PCM) tools can report per-channel bandwidth.

Hewlett Packard Enterprise recommends that you enable **Intel Performance Monitoring Support** so that you can run the persistent memory module performance monitoring tools.

- UEFI System Utilities: System Configuration > BIOS/Platform Configuration (RBSU) > Power and Performance Options > Intel Performance Monitoring Support
- iLO RESTful API property name: `IntelPerfMonitoring`
- **User default options**—After configuring the persistent memory settings for the server, Hewlett Packard Enterprise recommends saving the settings as the user default settings.
 - UEFI System Utilities: System Configuration > BIOS/Platform Configuration (RBSU) > System Default Options > User Default Options
 - iLO RESTful API property name: `SaveUserDefaults`

Managing Intel Optane persistent memory for HPE

There are a number of tools available to manage Intel Optane persistent memory for HPE, including:

- UEFI System Utilities
- RESTful Interface Tool
- Ipmctl, which can be run in the command line or under the UEFI Shell



Changing the goal configuration using UEFI System Utilities

Goal Configuration Options displays the most recent configuration settings, but not necessarily the active configuration. The configuration settings defined in this screen are only applied at the next server reboot.

i IMPORTANT:

Be sure to observe all pop-up messages displayed in UEFI System Utilities that pertain to persistent memory. Failure to follow the instructions in these messages might cause persistent memory data loss.

i IMPORTANT:

Always follow recommendations from your software application provider for high-availability best practices to ensure maximum uptime and data protection.

i IMPORTANT:

When data must be preserved, Hewlett Packard Enterprise strongly recommends that you perform a manual backup of all user data on the persistent memory modules before changing the goal configuration or performing relocation procedures.

Prerequisites

1. If the persistent memory modules are encrypted, you must disable key management before changing the goal configuration.
2. Sanitize all the persistent memory modules in the server using the Overwrite Media method. For more information, see [Sanitizing a persistent memory module](#).

Procedure

1. If persistent memory module encryption is enabled, disable it.
For more information, see [Disabling key management](#).
2. From the System Utilities screen, select **System Configuration > BIOS/Platform Configuration (RBSU) > Memory Options > Persistent Memory Options > PMM Options > Goal Configuration Options**.
3. Update the following selections:
 - **Volatile Memory Capacity**—% of persistent memory module capacity providing volatile memory.
 - Memory mode—Select **100%**.
 - App Direct mode—Select **0%**.
 - **Persistent Memory Interleaving**—Enabled or Disabled.
4. Select **Apply Goal Configuration**.
5. To save your changes, press the **F10** key.
6. To immediately commit the new goal configuration settings, reboot the server.
7. If encryption was disabled to change the goal configuration, enable it.

For more information, see [Enabling key management](#).

Deleting the goal configuration using UEFI System Utilities

Goal Configuration Options displays the most recent configuration settings, but not necessarily the active configuration. The configuration settings defined in this screen are only applied at the next server reboot.

i IMPORTANT:

Be sure to observe all pop-up messages displayed in UEFI System Utilities that pertain to persistent memory. Failure to follow the instructions in these messages might cause persistent memory data loss.

i IMPORTANT:

Always follow recommendations from your software application provider for high-availability best practices to ensure maximum uptime and data protection.

Prerequisites

If the persistent memory modules are encrypted, you must disable key management before deleting the goal configuration.

Procedure

1. If persistent memory module encryption is enabled, disable it.
For more information, see [Disabling key management](#).
2. From the System Utilities screen, select **System Configuration > BIOS/Platform Configuration (RBSU) > Memory Options > Persistent Memory Options > PMM Options > Goal Configuration Options**.
3. Select **Delete Goal Configuration**.
4. To save your changes, press the **F10** key.
5. To immediately delete the goal configuration settings, reboot the server.
6. If encryption was disabled to change the goal configuration, enable it.
For more information, see [Enabling key management](#).

Changing persistent memory module passwords

IMPORTANT:

Be sure to observe all pop-up messages displayed in UEFI System Utilities that pertain to persistent memory. Failure to follow the instructions in these messages might cause persistent memory data loss.

Procedure

1. During POST, press the F9 key to enter System Utilities.
2. From the System Utilities screen, select System Configuration > BIOS/Platform Configuration (RBSU) > Server Security > Device Encryption Options > Device Encryption Settings > Encrypted Devices.
3. Choose the persistent memory module from Select Device.
4. Select Modify Passphrase from Select Operation.
5. Choose the Passphrase Type:

This selection is available only when Local Key Management is enabled. When Remote Key Management is enabled, persistent memory module passwords are automatically generated, stored, and managed on the key management server.

- Auto—The system automatically generates a 32 byte random password. Hewlett Packard Enterprise recommends using system-generated passwords as a best practice.
 - Manual—Enter a 32 byte password manually.
6. Select Start Operation.
The persistent memory module password is changed.
 7. To change the password for each individual persistent memory module, repeat the process.
 8. Hewlett Packard Enterprise recommends exporting the password database to a USB device for backup purposes:
 - a. From the System Utilities screen, select System Configuration > BIOS/Platform Configuration (RBSU) > Server Security > Device Encryption Options > Device Encryption Migration Options > Device Encryption Export Options.
 - b. Provide a password in the Transient Passphrase field.
This password protects the exported file and must be entered when recovering the encrypted persistent memory modules after relocation.
 - c. Select Select File, and browse to a location on the USB key.
 - d. Select Export Encryption Settings to create and export the file.

Viewing the status of persistent memory modules

Procedure

1. From the System Utilities screen, select System Configuration > BIOS/Platform Configuration (RBSU) > Server Security > Device Encryption Options > Device Encryption Status.

The Device Encryption Status screen displays the name, encryption status, and the password of each persistent memory module in the server.

2. Review the status for each persistent memory module:

- Not encrypted—The persistent memory module is not encrypted.
- Local/TPM—The persistent memory module is encrypted with local key management and the password is displayed.

Make note of this password and store it securely. Hewlett Packard Enterprise recommends downloading the password file to a USB drive for backup purposes.

- Unknown key:
 - An encrypted persistent memory module from another server was installed, and not yet migrated.
 - The Restore Manufacturing Default Options was selected in UEFI System Utilities.
 - The HPE TPM has failed.

Changing the key management mode

The key management mode can be changed between local and remote key management. Encrypted persistent memory modules remain encrypted, but the passwords and the storage of those passwords change, based on the key management mode selected.

i IMPORTANT:

Be sure to observe all pop-up messages displayed in UEFI System Utilities that pertain to persistent memory. Failure to follow the instructions in these messages might cause persistent memory data loss.

Procedure

1. From the System Utilities screen, select **System Configuration > BIOS/Platform Configuration (RBSU) > Server Security > Device Encryption Options**.
2. Change the Key Management setting to one of the following:
 - **Local**—Enables local key management. The password used for encryption is stored locally on the server.
HPE TPM 2.0 must be installed to view and select this setting.
 - **Remote**—Enables remote key management. The password used for encryption is stored on a remote key server.
HPE iLO must be enrolled in and connected to a key manager to view and select this setting.
3. Press the **F12** key to save and exit.
4. Reboot the server.

Disabling key management

Disabling key management disables encryption for all encrypted persistent memory modules in the server. To disable encryption only for a single or specific persistent memory modules, see [Disabling encryption for a persistent memory module](#).

i IMPORTANT:

Be sure to observe all pop-up messages displayed in UEFI System Utilities that pertain to persistent memory. Failure to follow the instructions in these messages might cause persistent memory data loss.

Procedure

1. From the System Utilities screen, select **System Configuration > BIOS/Platform Configuration (RBSU) > Server Security > Device Encryption Options**.
2. Select the Key Management setting, and change it to Disabled.
3. Press the **F12** key to save and exit.
4. Reboot the server.

Disabling encryption for a persistent memory module

Use this procedure to disable encryption for a single or specific persistent memory modules.

To disable encryption for all persistent memory modules in the server at once, as might be required for migration or service procedures, see [Disabling key management](#).

i IMPORTANT:

Be sure to observe all pop-up messages displayed in UEFI System Utilities that pertain to persistent memory. Failure to follow the instructions in these messages might cause persistent memory data loss.

Procedure

1. From the System Utilities screen, select **System Configuration > BIOS/Platform Configuration (RBSU) > Server Security > Device Encryption Options > Device Encryption Settings > Encrypted Devices**.

2. Make the following selections:

a. **Select Device**—Select the persistent memory module.

b. **Select Operation**—Disable Encryption.

3. Select **Start Operation**.

If local key management is enabled, enter the passphrase for the persistent memory module.

The selected persistent memory module is now unencrypted.

4. Repeat this process to disable encryption for other persistent memory modules.

Changing Performance Options using UEFI System Utilities

i IMPORTANT:

Be sure to observe all pop-up messages displayed in UEFI System Utilities that pertain to persistent memory. Failure to follow the instructions in these messages might cause persistent memory data loss.

i IMPORTANT:

Always follow recommendations from your software application provider for high-availability best practices to ensure maximum uptime and data protection.

Procedure

1. From the System Utilities screen, select **System Configuration > BIOS/Platform Configuration (RBSU) > Memory Options > Persistent Memory Options > PMM Options > Performance Options.**
2. Based on your server workload and performance requirements, update the following options:
 - **Performance Setting**—Controls the baseline performance setting, depending on the workload behavior:
 - **Bandwidth Optimized**—Default; optimize the performance of App Direct (AD)
 - **Balanced**—Optimize the performance of Memory Mode
 - **FastGo Configuration**—Controls optimization of traffic within the processor:
 - **Auto**—Default
 - **Enabled**
 - **Disabled**—Recommended for improving the bandwidth of the sequential write traffic load in App Direct mode
 - **Snoopy Mode for App Direct**—Enable this option to avoid directory updates to persistent memory modules for non-NUMA (non-uniform memory access) optimized workloads:
 - **Disabled**—Default
 - **Enabled**
 - **Snoopy Mode for Memory mode**—Enable this option to avoid directory updates to persistent memory modules for non-NUMA optimized workloads:
 - **Disabled**—Default
 - **Enabled**
3. To save your changes, press the **F12** key.



HPE iLO RESTful API overview

The HPE iLO RESTful API for server management provides intelligent remote control. Use this single interface to perform remote server provisioning, configuration, inventory, and monitoring. The HPE iLO RESTful API conforms to the DMTF Redfish API standard. For more information about the HPE iLO RESTful API, see the Hewlett Packard Enterprise website (<https://www.hpe.com/us/en/servers/restful-api.html>).

The HPE iLO RESTful API can be accessed using a number of tools. Hewlett Packard Enterprise recommends using the RESTful Interface Tool. Third-party tools such as Postman, curl, and wget are also available.

Data model overview

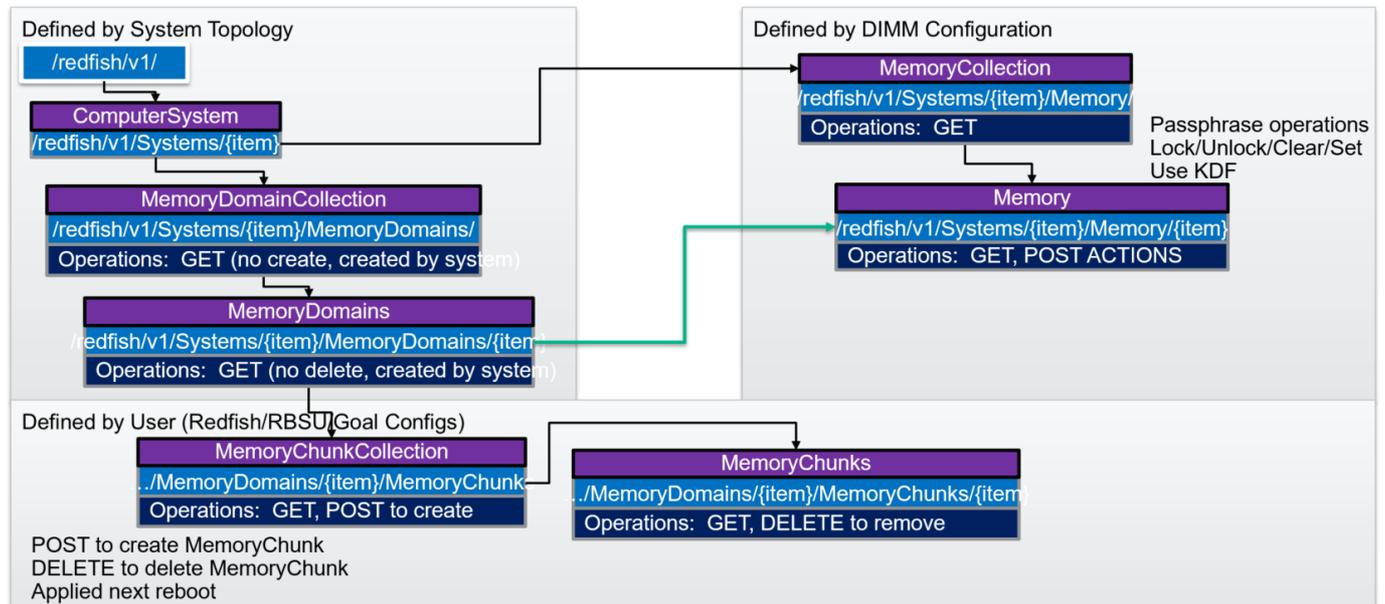
Specific resources describe the physical properties and configuration of persistent memory modules:

- Memory
- Memory Chunks
- Memory Domains
- Memory Regions

Term	Definition
Memory	Memory represents the DIMMs installed in the system.
Memory Chunk	A Memory Chunk is a group of one or more regions. The Memory Chunk represents an interleave set. Memory Domains and Chunks will be reported only for Persistent Regions. Volatile Regions will be treated just like DIMMs with no such data reported.
Memory Domain	Memory Domains are used to indicate to the client which Memory (DIMMs) can be grouped together in Memory Chunks to form interleave sets or otherwise grouped together (informational only, not configurable).
Memory Region	A region is a portion of a persistent memory module of a specific size and mode. A persistent memory module can have one or more regions. Regions can be of the same or different mode on a persistent memory module. For example, a persistent memory module can have a persistent region and a volatile region.
Interleave Set	A group of Memory Regions that are interleaved together. Represented by a Memory Chunk in Redfish.

Data model diagram

The following diagram shows the data model for persistent memory modules. The diagram shows the hierarchy, URI, and supported operations of each resource:



Examples: Retrieving memory resources

RESTful Interface Tool select and list to retrieve memory

The RESTful Interface Tool can be used to retrieve resources. There are several commands available:

- `select`
- `get`
- `list`
- `rawget`

The following is an example Windows batch script to retrieve all the memory resources in the system and print in JSON format:

```
@echo off

set argC=0
for %%x in (*) do Set /A argC+=1
if %argC% LSS 3 goto :failCondition
goto :main

:failCondition
@echo Usage:
@echo ilorest-script-memory-remote.bat [URL] [USERNAME] [PASSWORD]
goto :EOF

:main
@echo Logging in...
ilorest.exe --nologo login %1 -u %2 -p %3
@echo selecting Memory type...
ilorest.exe --nologo select Memory.
@echo list Memory data in JSON format...
ilorest.exe --nologo list -json
```

Using python to retrieve expanded Memory Collection

The following example uses the Python requests library to retrieve the Memory Collection for a given server. The expand query is used to retrieve all members at once.

```
import requests
from requests.auth import HTTPBasicAuth
import sys
import json

# server info
if len(sys.argv) < 4:
    sys.stdout.write("\nPlease supply the URL, username and password:" \
    "\nUsage: python clear_all_tasks.py https://ilourl username password\n")
    exit(-1)

iLO_URL = sys.argv[1]
username = sys.argv[2]
password = sys.argv[3]

# REST info
MEMORY_URI = "/redfish/v1/systems/1/Memory?$expand=.#"

# Get the Memory
sys.stdout.write("Retrieving all Memory...")
with requests.Session() as s:
    get_response = s.get(iLO_URL + MEMORY_URI, \
        auth=HTTPBasicAuth(username, password))
    body = get_response.json()
s.close()
```

```
if get_response.status_code != 200:
    sys.stdout.write("error occurred: {}".format(get_response.status_code))
else:
    sys.stdout.write(json.dumps(body, indent=2, separators=(',', ': ')))
```

Using Postman to retrieve expanded Memory Collection

The following example uses Postman to retrieve the Memory Collection for a given server. The expand query is used to retrieve all members at once.

Operation: GET

Path: /redfish/v1/systems/1/memory?\$expand=.#

The screenshot shows the Postman interface for a GET request. The URL is `https://ilo.fulldomain.com/redfish/v1/systems/1/memory?$expand=.#`. The request is configured with Basic Auth, username 'user', and a masked password. The response body is displayed in JSON format, showing a list of memory members with detailed properties.

```
36 }
37 },
38 "Members": [
39 {
40   "@odata.context": "/redfish/v1/$metadata#Memory.Memory",
41   "@odata.id": "/redfish/v1/Systems/1/Memory/procldimm1",
42   "@odata.type": "#Memory.v1_7_0.Memory",
43   "Id": "procldimm1",
44   "BusWidthBits": 72,
45   "CacheSizeMiB": -1,
46   "CapacityMiB": 0,
47   "DataWidthBits": 64,
48   "DeviceLocator": "PROC 1 DIMM 1",
49   "ErrorCorrection": "MultiBitECC",
50   "LogicalSizeMiB": 0,
51   "MemoryLocation": {
52     "Channel": 6,
53     "MemoryController": 2,
54     "Slot": 1,
55     "Socket": 1
56   }
57 },
58 ]
59 }
```

Managing Intel Optane persistent memory for HPE with the HPE iLO RESTful API

To manage persistent memory modules using the HPE iLO RESTful API, use the associated commands.

Command	System Utilities option
PmmPerformance	Performance Setting
BandwidthOptimized	Bandwidth Optimized (default)
Balanced	Balanced
PmmFastGo	FastGo Configuration
Enabled	Enabled
Disabled	Disabled
Auto	Auto (default)
PmmAppDirectSnoopyMode	Snoopy Mode for App Direct
Enabled	Enabled
Disabled	Disabled (default)
PmmMemModeSnoopyMode	Snoopy Mode for Memory Mode
Enabled	Enabled
Disabled	Disabled (default)
VolatileMemCapacityPercent	Volatile Memory Capacity
PersistentMemoryInterleaving	Persistent Memory Interleaving
Enabled	Enabled
Disabled	Disabled
ApplyDefaultNamespaces	Apply Default Namespaces
Enabled	Enabled
Disabled	Disabled
SecurityFreezeLock	Security Freeze Lock
Enabled	Enabled
Disabled	Disabled
PmmSanitizeOperation	Sanitize/Erase Operation on Reboot
NoAction	No Action
CryptoErase	Cryptographically Erase
Overwrite	Overwrite Media
CryptoEraseOverwrite	Cryptographically Erase and then Overwrite Media
PmmSanitizePolicy	Policy after Sanitize/Erase on Reboot
SanitizeAndRebootSystem	Sanitize/Erase and Reboot System
SanitizeAndShutdownSystem	Sanitize/Erase and Power System Off
SanitizeAndBootToFirmwareUI	Sanitize/Erase and Reboot to System Utilities
SanitizeToFactoryDefaults	Sanitize/Erase to Factory Defaults and Power System Off
SanitizeAllPmm	Perform Sanitize/Erase Operation on: All persistent memory modules in the System
SanitizeProcXPmm¹	Perform Sanitize/Erase Operation on: All persistent memory modules on Processor X
SanitizeProcXPmmY¹	Perform Sanitize/Erase Operation on: Processor X DIMM Y

¹ Where X and Y represent the processor and DIMM slot number, such as `SanitizeProc1Pmm4`.



Persistent memory module provisioning with the HPE iLO RESTful API

The HPE iLO RESTful API provides a mechanism for configuring persistent memory modules. The configuration is modified by creating and deleting the Memory Chunks. Because configuration requires a reboot, the HPE iLO RESTful API uses Redfish Tasks to represent pending and complete configuration operations.

Required REST actions to provision persistent memory modules

Configuration	REST actions
100% App Direct Interleaved (Requires one Memory Chunk per processor)	<p>POST a Memory Chunk to the Memory Domain to configure. Set the type to PMEM and the size/percentage to 100% or the size to the total of the persistent memory module capacity for the Memory Domain.</p> <p>To interleave, include all the persistent memory modules in the Memory Chunk Interleave Set.</p>
100% App Direct Not Interleaved (Requires one Memory Chunk per persistent memory module)	<p>POST Memory Chunks to the Memory Domain to configure. Set the type to PMEM and the size/percentage to either 100 for percentage, or the capacity of the persistent memory module for size.</p> <p>For the not-interleaved case, there is an Interleave Set for each persistent memory module.</p>
100% Volatile (Requires one Memory Chunk per processor)	<p>POST a Memory Chunk to the Memory Domain to configure. Set the type to PMEM and the size/percentage to zero.</p> <p>To interleave, include all the persistent memory modules in the Memory Chunk Interleave Set.</p>
Clear existing configuration	DELETE existing Memory Chunks.
Clear pending configuration	DELETE Tasks where TaskState is New and TargetUri is one of the Memory Chunk collections.

Examples: Provisioning persistent memory modules

Using RESTful Interface Tool rawpost to configure 100% App Direct interleaved

The `rawpost` command takes in a JSON file. The following example displays the JSON file and a batch script to configure a server.

Memorychunk-rawpost.txt

```
{
  "path": "/redfish/v1/Systems/1/MemoryDomains/PROC1MemoryDomain/MemoryChunks",
  "body": {
    "AddressRangeType": "PMEM",
    "Oem": {
      "Hpe": {
        "MemoryChunkSizePercentage": 100
      }
    },
    "InterleaveSets": [{
      "Memory": {
        "@odata.id": "/redfish/v1/Systems/1/Memory/proc1dim6/"
      }
    }, {
      "Memory": {
        "@odata.id": "/redfish/v1/Systems/1/Memory/proc1dim7/"
      }
    }
  ]
}
```

Windows batch script

```
@echo off

set argC=0
for %%x in (*) do Set /A argC+=1
if %argC% LSS 3 goto :failCondition
goto :main

:failCondition
@echo Usage:
@echo ilorest-script-memory-remote.bat [URL] [USERNAME] [PASSWORD]
goto :EOF

:main
@echo Logging in...
ilorest.exe --nologo login %1 -u %2 -p %3
@echo rawpost to Memory Chunk collection...
ilorest.exe --nologo rawpost memorychunk-rawpost.txt
@echo Note: Status of 202 is success.
```

Using python to configure 100% App Direct interleaved

The following example uses the requests library to make a POST request to create a Memory Chunk. Because a reboot is required, a Task will be generated.

i IMPORTANT:

Submitting passwords through command line is not secure. Hewlett Packard Enterprise recommends using best practices for command line use, such as storing passwords in files.

```
import requests
from requests.auth import HTTPBasicAuth
import sys
import json

# server info from command line
if len(sys.argv) < 4:
    sys.stdout.write("\nPlease supply the URL, username and password:" \
    "\nUsage: python post_test.py https://ilourl username password\n")
    exit(-1)

iLO_URL = sys.argv[1]
username = sys.argv[2]
password = sys.argv[3]

# REST info
CHUNKS_URI = "/redfish/v1/Systems/1/MemoryDomains/PROClMemoryDomain/MemoryChunks"
headers = {'Content-type': 'application/json', 'Accept': 'application/json'}

MemoryChunk = {
    "AddressRangeType": "PMEM",
    "Oem": {
        "Hpe": {
            "MemoryChunkSizePercentage": 100
        }
    },
    "InterleaveSets": [
        {
            "Memory" : { "@odata.id": "/redfish/v1/Systems/1/Memory/proclimm6/" }
        },
        {
            "Memory" : { "@odata.id": "/redfish/v1/Systems/1/Memory/proclimm7/" }
        }
    ]
}

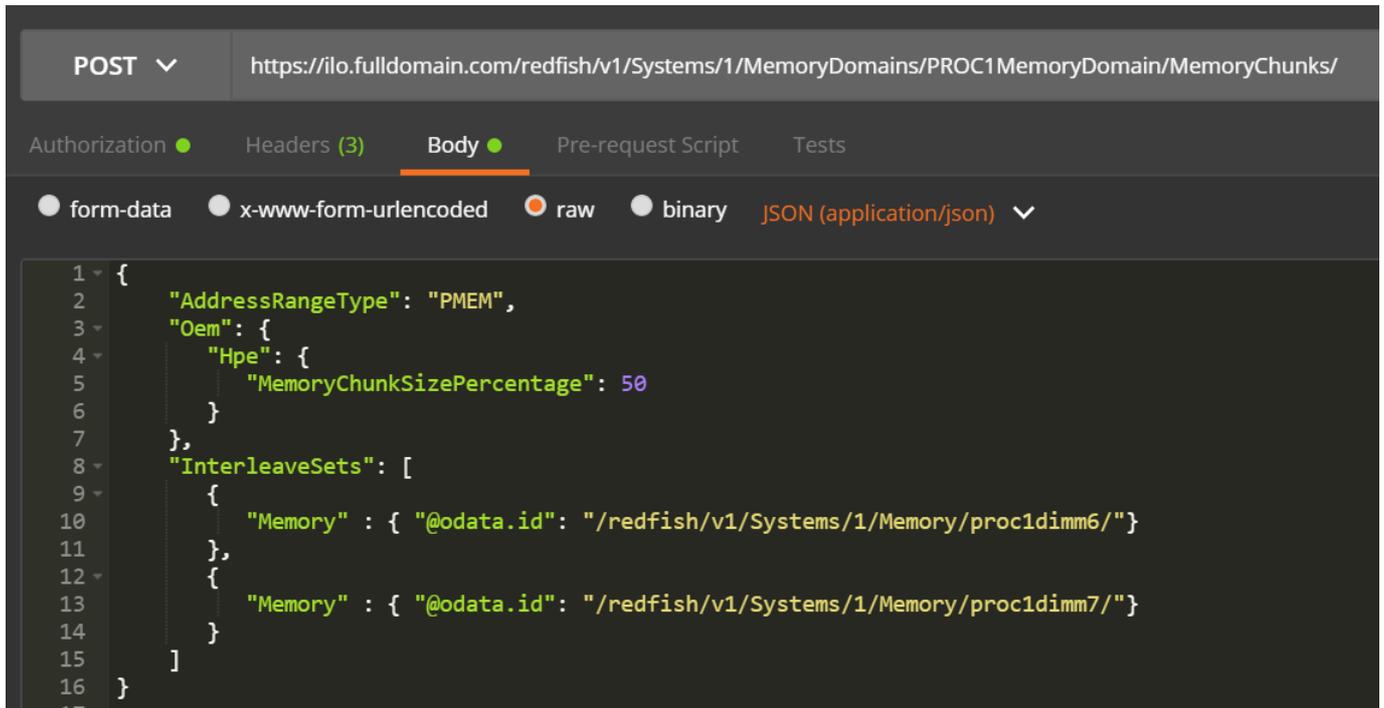
# POST to the URI until there is an error
sys.stdout.write("POST MemoryChunk...")
with requests.Session() as s:
    response = requests.post(iLO_URL + CHUNKS_URI, data=json.dumps(MemoryChunk), headers=headers, \
        auth=HTTPBasicAuth(username, password), verify=False)
s.close()
if response.status_code != 202:
    sys.stdout.write('\n\nREST error; POST unsuccessful. Status={}' \
    '\n'.format(response.status_code))
else:
    output = json.loads(response.text)
    sys.stdout.write("\nPOST successful: {}".format(output.get("Name")))
```

Using Postman to configure 100% App Direct interleaved

Operation: POST

Path and Body (raw JSON) same as the example in [Using RESTful Interface Tool rawpost to configure 100% App Direct interleaved](#).

Headers: Accept: application/json, Content-Type: application/json



Example: Managing Intel Optane persistent memory for HPE using curl

Each DIMM and persistent memory module is represented by a memory object, identified by the processor number and DIMM slot. For a persistent memory module, the attributes appear similar to the following example:

```
curl --insecure --noproxy '*' --location --user 'user:password' --request GET --header 'Content-Type:application/json' --header 'Accept:application/json' http://ilonam
{
  "@odata.context": "/redfish/v1/$metadata#Memory.Memory",
  "@odata.etag": "W/\"AAAA879\"",
  "@odata.id": "/redfish/v1/Systems/1/Memory/proclidimm11/",
  "@odata.type": "#Memory.v1_7_0.Memory",
  "AllocationAlignmentMiB": 1024,
  "AllocationIncrementMiB": 1024,
  "BaseModuleType": "PMM",
  "BusWidthBits": 72,
  "CacheSizeMiB": 0,
  "CapacityMiB": 514624,
  "DataWidthBits": 64,
  "DeviceID": "16721",
  "DeviceLocator": "PROC 1 DIMM 11",
  "ErrorCorrection": "MultiBitECC",
  "FirmwareApiVersion": "01.01.00.5253",
  "FirmwareRevision": "01.01.00.5253",
  "Id": "proclidimm11",
  "LogicalSizeMiB": 0,
  "Manufacturer": "INTEL",
  "MemoryDeviceType": "DDR4",
  "MemoryLocation": {
    "Channel": 3,
    "MemoryController": 1,
    "Slot": 11,
    "Socket": 1
  },
  "MemoryMedia": [
    "Intel3DXPoint"
  ],
  "MemoryType": "IntelOptane",
  "Name": "proclidimm11",
  "NonVolatileSizeMiB": 514048,
  "Oem": {
    "Hpe": {
      "@odata.context": "/redfish/v1/$metadata#HpeMemoryExt.HpeMemoryExt",
      "@odata.type": "#HpeMemoryExt.v2_1_0.HpeMemoryExt",
      "BaseModuleType": "PMM",
      "BlocksRead": 36366041872668,
      "BlocksWritten": 2603586169856,
      "DIMMStatus": "GoodInUse",
      "MinimumVoltageVoltsX10": 12,
      "PredictedMediaLifeLeftPercent": 100,
      "ProductName": "HPE Persistent Memory"
    }
  },
  "OperatingMemoryModes": [
    "Volatile",
    "PMEM"
  ],
  "OperatingSpeedMhz": 2666,
  "PartNumber": "835810-B21",
  "PersistentRegionNumberLimit": 48,
  "PersistentRegionSizeLimitMiB": 514048,
  "PersistentRegionSizeMaxMiB": 0,
  "RankCount": 1,
  "Regions": [
    {
      "MemoryClassification": "Volatile",
      "PassphraseEnabled": false,
      "RegionId": "15",
      "SizeMiB": 576
    },
    {
      "MemoryClassification": "ByteAccessiblePersistent",
      "PassphraseEnabled": false,
      "RegionId": "16",
      "SizeMiB": 514048
    }
  ],
  "SecurityCapabilities": {
    "PassphraseCapable": true
  },
  "SerialNumber": "8089-A2-1834-000026B6",
  "Status": {
    "Health": "OK",
    "State": "Enabled"
  },
  "SubsystemDeviceID": "2426",
  "SubsystemVendorID": "35200",
  "VendorID": "35200",
  "VolatileRegionNumberLimit": 1,
  "VolatileRegionSizeLimitMiB": 576,
  "VolatileRegionSizeMaxMiB": 0,
  "VolatileSizeMiB": 576
}
```

Each processor that has persistent memory modules installed is represented by a MemoryDomains object:

```
{
  "@odata.context": "/redfish/v1/$metadata#MemoryDomainCollection.MemoryDomainCollection",
  "@odata.etag": "W/\"AA6D42B0\"",
  "@odata.id": "/redfish/v1/Systems/1/MemoryDomains/",
  "@odata.type": "#MemoryDomainCollection.MemoryDomainCollection",
  "Description": "Memory Domains Collection",
  "Name": "Memory Domains Collection",
}
```

```
"Members": [  
  {  
    "@odata.id": "/redfish/v1/Systems/1/MemoryDomains/PROCIMemoryDomain/"  
  }  
],  
"Members@odata.count": 1  
}
```

RESTful Interface Tool

Use the RESTful Interface Tool to manage Intel Optane persistent memory 200 series for HPE.

The RESTful Interface Tool is a CLI tool that configures the system using the RESTful API through HPE iLO. The tool can run locally on the server or remotely connect to the server through HPE iLO. The RESTful Interface Tool can run in interactive mode or from command line, which is useful for scripting.

For more information about the iLO RESTful Interface Tool commands, see the iLO RESTful API for HPE iLO 5 documentation at <https://hewlettpackard.github.io/ilo-rest-api-docs/>.

Launching the RESTful Interface Tool

The RESTful Interface Tool supports two modes:

- Interactive mode
- Scripting mode

The examples in this guide are shown in interactive mode. Using scripting mode works in a similar way.

Procedure

Do one of the following:

- To launch the tool in interactive mode, do the following:

1. Locate and run the `ilorest.exe` build file.
2. Log in to the server:

```
iLOrest > login iLO_IP -u username -p password
```

3. Run the commands

```
iLOrest > command [options]
```

- To launch the tool in scripting mode, do the following:

1. Navigate to the folder containing the `ilorest.exe` file.
2. Log in to the server:

```
C:\ilorest>ilorest.exe login iLO_IP -u username -p password
```

3. Run the command:

```
C:\ilorest>ilorest.exe command [options]
```

Discovery commands

This command displays information about the physical view and configuration of persistent memory module and App Direct interleaved sets, based on the flags specified.

```
showpmm [flag] [options]
```

Discovery command flags

The following flags can be used with the command.

Flag	Description
-D, --device	Show information about the physical persistent memory modules.
-C, --config	Show the configuration of the persistent memory modules.
-L, --logical	Display the persistent interleave sets.
-M, --summary	Display the memory summary.

Device discovery

The command shows information about the physical view of the persistent memory modules. If the `showpmm` command is run without any option, then the `--device` flag is the default view.

```
showpmm -D|--device [-I|--dimm=(DimmIDs)] [-j|--json] [-h|--help]
```

Options

The following options can be used with the command.

Option	Description
<code>-h, --help</code>	Display help for the command.
<code>-I, --dimm</code>	To view information about specific persistent memory modules, supply a comma-separated list of DIMM IDs in the format <code>P@S</code> , where P = processor and S = slot. For example: <code>1@1,1@12</code> .
<code>-j, --json</code>	Output data in JSON format. If the json flag is not specified, then the default display format is a table.

Examples

- To display information for all physical persistent memory modules, run:

```
iLOrest > showpmm -D
```

- To display information for the persistent memory modules installed at processor 1, slot 12 and at processor 2, slot 1, run:

```
iLOrest > showpmm -D --dimm=1@12,2@1
```

Return data

The return data displays the following attributes in a table for each persistent memory module.

Attribute	Description
Location	The physical location of the persistent memory module
Capacity	The usable capacity of the persistent memory module.
Status	Overall persistent memory module health.
DIMM Status	Specifies memory module status and whether the module is in use.
Life	The estimated remaining life of the device (in %).
FWVersion	The revision of the active firmware.

Device configuration discovery

The command shows configuration of the individual persistent memory modules using the `--config` flag.

```
showpmm -C|--config [-I|--dimm=(DimmIDs)] [-j|--json] [-h|--help]
```

Options

The following options can be used with the command.

Option	Description
<code>-h, --help</code>	Display help for the command.
<code>-I, --dimm</code>	To view information about specific persistent memory modules, supply a comma-separated list of DIMM IDs in the format <code>P@S</code> , where P = processor and S = slot. For example: <code>1@1,1@12</code> .
<code>-j, --json</code>	Output data in JSON format. If the json flag is not specified, then the default display format is a table.

Examples

- To display configuration details for all persistent memory modules, run:

```
iLOrest > showpmm -C
```

- To display configuration details for the persistent memory modules installed at processor 1, slot 12 and at processor 2, slot 1, run:

```
iLOrest > showpmm -C --dimm=1@12,2@1
```

- To display configuration details in JSON format for the persistent memory module installed at processor 2, slot 12, run:

```
iLOrest > showpmm -C --dimm=2@12 --json
```

Return data

The return data displays the following attributes in a table for each persistent memory module installed in the host server.

Attribute	Description
Location	The physical location of the persistent memory module.
VolatileSize	The total size of volatile regions on the persistent memory module.
PmemSize	The total size of persistent regions on the persistent memory module.
PmemInterleaved	Indicates if the Persistent regions are interleaved.

Persistent interleave regions discovery

The command shows information about the persistent interleave regions, representing a logical view of the persistent memory modules.

```
showpmm -L|--logical [-j|--json] [-h|--help]
```

Options

The following options can be used with the command.

Option	Description
<code>-h, --help</code>	Display help for the command.
<code>-j, --json</code>	Output data in JSON format. If the json flag is not specified, then the default display format is a table.

Examples

- To display information on the persistent interleave regions, run:

```
iLOrest > showpmm --logical
```

- To display information on the persistent interleave regions in JSON format, run:

```
iLOrest > showpmm --logical --json
```

Return data

The return data displays the following attributes in a table for each persistent interleave set.

Attribute	Description
TotalPmemSize	The total size of interleaved persistent regions.
DimmIds	The physical location of the interleaved DIMMs in the format <code>P@S</code> , where P = processor index and S = slot index.

Persistent memory summary

The command displays a configuration summary of the persistent memory modules using the `--summary` flag.

```
showpmm -M|--summary [-j|--json] [-h|--help]
```

Options

The following options can be used with the command.

Option	Description
<code>-h, --help</code>	Display help for the command.
<code>-j, --json</code>	Output data in JSON format. If the json flag is not specified, then the default display format is a table.

Examples

- To display the memory summary, run:

```
iLOrest > showpmm --summary
```

- To display the memory summary in JSON format, run:

```
iLOrest > showpmm --summary --json
```

Return data

The return data displays the following attributes in a table for each persistent interleave set.

Attribute	Description
TotalCapacity	The sum of the usable capacity of all the persistent memory modules.
TotalVolatileSize	The sum of the total size of volatile regions on each module.
TotalPmemSize	The sum of the total size of persistent regions on each module.

Show pending configuration

The command shows the pending configuration tasks related to persistent memory modules that require a reboot to take effect.

```
showpmpendingconfig [-j|--json] [-h|--help]
```

Options

The following options can be used with the command.

Option	Description
<code>-h, --help</code>	Display help for the command.
<code>-j, --json</code>	Output data in JSON format. If the json flag is not specified, then the return data is displayed in a table by default.

Examples

- To display pending configuration details, run:

```
iLOrest > showpmpendingconfig
```

- To display pending configuration details in JSON format, run:

```
iLOrest > showpmpendingconfig --json
```

Return data

The return data displays the following attributes in a table.

Attribute	Description
Operation	The action to be performed.
PmemSize	The total size of all persistent regions.
VolatileSize	The total size of all volatile regions.
DIMMids	The physical location of the interleaved DIMMs in the format <code>P@S</code> , where P = processor index and S = index.

Apply predefined configuration

The command applies a predefined configuration to all the persistent memory modules and deletes any existing or pending configuration.

The command supports three modes:

- 100% Memory mode
- 100% Persistent with interleaving
- 100% Persistent without interleaving

```
Applypmmconfig (-C|--config =(configID) | -L|--list) [-f|--force] [-h|--help]
```

Configuration changes require a reboot to take effect.

Options

The following options can be used with the command.

Option	Description
<code>-h, --help</code>	Display help for the command.
<code>-C, --config</code>	Specify the configID to apply.
<code>--list</code>	List all the available configIDs with a description.
<code>-f, --force</code>	Force the configuration by automatically accepting any prompts. This option bypasses warnings for an existing configuration or pending configuration.

Examples

- To display a list with all the available configIDs with a description, run:

```
iLOrest > applypmmconfig --list
```

- To configure all the persistent memory modules for 100% Memory mode, run:

```
iLOrest > applypmmconfig -C MemoryMode -f
```

- To configure all the persistent memory modules for 100% Persistent with interleaving, run:

```
iLOrest > applypmmconfig -C PmemInterleaved -f
```

- To configure all the persistent memory modules for 100% Persistent without interleaving, run:

```
iLOrest > applypmmconfig -C PmemNotInterleaved -f
```

Return data

The return data displays the following attributes in a table.

Attribute	Description
Operation	The action to be performed.
PmemSize	The total size of all persistent regions.
VolatileSize	The total size of all volatile regions.
DIMMids	The physical location of the interleaved DIMMs in the format <code>P@S</code> , where P = processor index and S = slot index.

Apply a user-defined configuration

The command applies a user-defined configuration to all the persistent memory modules and deletes any existing or pending configuration.

Configurations must adhere to the recommended memory cache ratios. Defining a nonrecommended ratio might impact system performance and will generate messages in the IML.

```
provisionpmm [-m|--memory-mode=(0|%) ] [-i|--pmem-interleave=(On|Off) ] [-p|--proc=(processorID) ]  
[-f|--force] [-h|--help]
```

Configuration changes require a reboot to take effect.

Options

The following options can be used with the command.

Option	Description
<code>-h, --help</code>	Display help for the command.
<code>-m, --memory-mode</code>	Specify the percentage of total capacity to configure as volatile memory. Defaults to 0% volatile memory, and the remaining capacity is configured as persistent memory.
<code>-i --pmem-interleave</code>	Indicate whether the persistent memory regions must be interleaved. Allowed values are <code>on</code> or <code>off</code> .
<code>-p --proc</code>	Specify the processors (comma-separated list of processor numbers) on which the selected configuration will be applied. Defaults to all processors.
<code>-f, --force</code>	Force the configuration by automatically accepting any prompts. This option bypasses warnings for an existing configuration or pending configuration.

Examples

- To configure all persistent memory modules on processors 1 and 3 for 50% volatile memory, with no persistent interleave regions, run:

```
iLOrest > provisionpmm -m 50 -i off -p 1,3
```

- To configure all the persistent memory modules for 25% volatile memory with persistent interleave regions, run:

```
iLOrest > provisionpmm -m 25 -i on
```

Return data

The return data displays the following attributes in a table.

Attribute	Description
Operation	The action to be performed.
PmemSize	The total size of all persistent regions.
VolatileSize	The total size of all volatile regions.
DIMMids	The physical location of the interleaved DIMMs in the format <code>P@S</code> , where P = processor index and S = slot index.

Clear pending configuration

The command clears all pending configuration tasks.

```
clearpmpendingconfig [-h|--help]
```

Options

The following options can be used with the command.

Option	Description
<code>-h, --help</code>	Display help for the command.

Examples

To delete all pending persistent memory configuration tasks, run:

```
iLOrest > clearpmpendingconfig
```

Return data

The return data prints a list of all tasks deleted:

```
Deleted Task #701  
Deleted Task #702  
Deleted Task #703  
Deleted Task #704
```

Show recommended configuration

The command shows the recommended persistent memory configurations.

```
showrecommendedpmmconfig [-h|--help]
```

Options

The following options can be used with the command.

Option	Description
<code>-h, --help</code>	Display help for the command.

Examples

To show the recommended configurations, run:

```
iLOrest > showrecommendedpmmconfig
```

Return data

The return data displays the following attributes in a table.

Attribute	Description
MemoryModeTotalSize	The total size of volatile regions.
PmemTotalSize	The total size of all persistent memory regions.
CacheRatio	The cache ratio.



NOTE:

The ipmctl tool executed from the operating system does not support key management features. To enable key management and enable or disable encryption of persistent memory modules, use the UEFI System Utilities.

Installing ipmctl for Linux

SUSE Linux Enterprise Server 12 SP4

To use ipmctl, Hewlett Packard Enterprise recommends downloading the latest openSUSE pre-build packages from:

- https://build.opensuse.org/package/binaries/home:jhli/ipmctl/SLE_12_SP4
- https://build.opensuse.org/package/binaries/home:jhli/safeclib/SLE_12_SP4

SUSE Linux Enterprise Server 15

To use ipmctl, Hewlett Packard Enterprise recommends downloading the latest openSUSE pre-build packages from:

- https://build.opensuse.org/package/binaries/home:jhli/ipmctl/SLE_15
- https://build.opensuse.org/package/binaries/home:jhli/safeclib/SLE_15

SUSE Linux Enterprise Server 15 SP1

To use ipmctl, Hewlett Packard Enterprise recommends downloading the latest openSUSE pre-build packages from:

- https://build.opensuse.org/package/binaries/home:jhli/ipmctl/SLE_15
- https://build.opensuse.org/package/binaries/home:jhli/safeclib/SLE_15

Red Hat Enterprise Linux 7.6

To use ipmctl, Hewlett Packard Enterprise recommends downloading the CentOS7 pre-build packages from:

- <https://copr.fedorainfracloud.org/coprs/jhli/ipmctl/>
- <https://copr.fedorainfracloud.org/coprs/jhli/safeclib/>

Red Hat Enterprise Linux 8.0

To use ipmctl, Hewlett Packard Enterprise recommends downloading the CentOS7 pre-build packages from:

- <https://copr.fedorainfracloud.org/coprs/jhli/ipmctl/>
- <https://copr.fedorainfracloud.org/coprs/jhli/safeclib/>

Showing persistent memory module configurations with ipmctl

ipmctl can show the current configuration of a persistent memory module:

```
ipmctl show -dimm
DimmID | Capacity | LockState | HealthState | FWVersion
=====
0x0001 | 502.5 GiB | Disabled | Healthy | 2.2.0.1553
0x0011 | 502.5 GiB | Disabled | Healthy | 2.2.0.1553
0x0021 | 502.5 GiB | Disabled | Healthy | 2.2.0.1553

ipmctl show -d Capacity,MemoryCapacity,AppDirectCapacity,UnconfiguredCapacity -dimm 0x1
---DimmID=0x0001---
Capacity=502.5 GiB
MemoryCapacity=0 B
AppDirectCapacity=502.0 GiB
UnconfiguredCapacity=0 B
```

To determine if a goal configuration is already pending, run:

```
ipmctl show -goal
```

Deleting a goal configuration with ipmctl

```
ipmctl delete -goal
```

The system retains the previous goal configuration settings.

Deleting namespace with ipmctl

```
ipmctl delete -namespace
```

Determining the memory mode with ipmctl

Run the following command to determine if the server is in Memory mode:

```
ipmctl show -memoryresources
```

```
Capacity=3015.5 GiB  
MemoryCapacity=0.0 GiB  
AppDirectCapacity=3012.0 GiB  
UnconfiguredCapacity=3.3 GiB  
InaccessibleCapacity=0.0 GiB  
ReservedCapacity=0.2 GiB
```

The server is in Memory mode if the command returns a nonzero amount for MemoryCapacity.

Persistent memory module relocation guidelines

Observe the relocation guidelines when doing the following:

- When relocating persistent memory modules to another DIMM slot on the server.
- When relocating persistent memory modules to another server.
- When reinstalling persistent memory modules after replacing the server system board.

i IMPORTANT:

When data must be preserved, Hewlett Packard Enterprise strongly recommends that you perform a manual backup of all user data on the persistent memory modules before changing the goal configuration or performing relocation procedures.

Requirements for relocating persistent memory modules or a set of persistent memory modules when the data must be preserved

- The destination server hardware must match the original server hardware configuration.
- All System Utilities settings in the destination server must match the original System Utilities settings in the original server.
- If persistent memory modules are used with Persistent Memory Interleaving set to Enabled in the original server, do the following:
 - Install the persistent memory modules in the same DIMM slots in the destination server.
 - Install the entire interleaved set (all the DIMMs and persistent memory modules on the processor) on the destination server.

If any of the requirements cannot be met during relocation, do the following:

- Manually back up the persistent memory data before relocating persistent memory modules to another server.
- Relocate the persistent memory modules to another server.
- Sanitize all persistent memory modules on the new server before using them.

Requirements for relocating encrypted persistent memory modules or a set of persistent memory modules when the data must be preserved

- If persistent memory modules are encrypted with local key management, either manually retrieve the persistent memory module passwords from the server (user-generated passwords only) or export a password file to a USB key.

Hewlett Packard Enterprise recommends exporting the password file to a USB key.

- Follow the requirements for relocating persistent memory modules or a set of persistent memory modules when the data must be preserved.
- Do one of the following:
 - If persistent memory modules are encrypted with local key management, either manually enter the persistent memory module passwords in the System Utilities or import the password file from the USB key.
 - If persistent memory modules are encrypted with remote key management, enroll the HPE iLO in the key management server to provide access to the data on the persistent memory modules.

Requirements for relocating persistent memory modules or a set of persistent memory modules when the data does not have to be preserved

- Install the persistent memory modules in the new location, and then sanitize the persistent memory modules.
- Observe the DIMM and persistent memory module population guidelines.
- Observe the process for removing a persistent memory module.
- Observe the process for installing a persistent memory module.
- Review and configure the system settings for Intel Optane persistent memory for HPE.

Manually backing up persistent memory module data

Hewlett Packard Enterprise recommends backing up data on the persistent memory modules before changing the goal configuration or performing service procedures.

⚠ CAUTION:

Electrostatic discharge can damage electronic components. Be sure you are properly grounded before beginning this procedure.

⚠ CAUTION:

Failure to properly handle DIMMs can cause damage to DIMM components and the system board connector.

For server-specific steps used in this procedure, see the server maintenance and service guide for your product on the Hewlett Packard Enterprise website:

- HPE ProLiant Gen10 Plus servers (<https://www.hpe.com/info/proliantgen10plus-docs>)
- HPE Synergy Gen10 Plus compute modules (<https://www.hpe.com/info/synergy-docs>)

Prerequisites

Before handling or removing a DIMM or persistent memory module, see the [Persistent memory module handling guidelines](#).

Procedure

1. Copy the data from the persistent memory modules to another storage device (such as SSD or HDD).
2. If persistent memory module encryption is enabled, disable it.
For more information, see [Disabling key management](#).
3. Power down the server:
 - a. Shut down the OS as directed by the OS documentation.
 - b. To place the server in standby mode, press the Power On/Standby button.
When the server enters standby power mode, the system power LED changes to amber.
 - c. Disconnect the power cords (rack and tower servers).
4. Do one of the following:
 - Extend the server from the rack.
 - Remove the server from the rack, if necessary.
 - Remove the server or server blade from the enclosure.
5. Place the server on a flat, level work surface.
6. Remove the access panel.
7. Access the DIMM slots.
8. Perform the relocation or replacement procedure.
9. Install any components removed to access the DIMM slots.
10. Install the access panel.
11. Install the server in the rack.
12. Power up the server.
13. Copy the data from the storage device to the persistent memory modules.

Removing a DIMM or persistent memory module

⚠ CAUTION:

Electrostatic discharge can damage electronic components. Be sure you are properly grounded before beginning this procedure.

⚠ CAUTION:

Failure to properly handle DIMMs can cause damage to DIMM components and the system board connector.

For server-specific steps used in this procedure, see the server maintenance and service guide for your product on the Hewlett Packard Enterprise website:

- HPE ProLiant Gen10 Plus servers (<https://www.hpe.com/info/proliantgen10plus-docs>)
- HPE Synergy Gen10 Plus compute modules (<https://www.hpe.com/info/synergy-docs>)

Prerequisites

- Before handling or removing a DIMM or persistent memory module, see the [Persistent memory module handling guidelines](#).
- If the persistent memory modules are encrypted, you must disable encryption before replacing a failed persistent memory module.

Procedure

1. If persistent memory module encryption is enabled, disable it.

For more information, see [Disabling key management](#).

2. Power down the server:

- a. Shut down the OS as directed by the OS documentation.
- b. To place the server in standby mode, press the Power On/Standby button.

When the server enters standby power mode, the system power LED changes to amber.

- c. Disconnect the power cords (rack and tower servers).

3. Do one of the following:

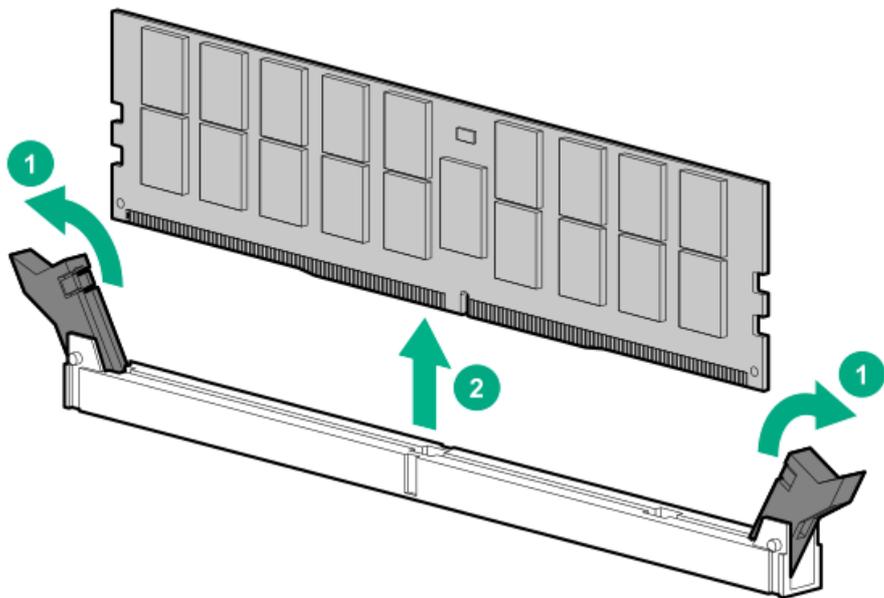
- Extend the server from the rack.
- Remove the server from the rack, if necessary.
- Remove the server or server blade from the enclosure.

4. Place the server on a flat, level work surface.

5. Remove the access panel.

6. Access the DIMM slots.

7. Remove the DIMM or persistent memory module.



Replacing a system board

Review this procedure for an overview of replacing a system board in a server without encrypted persistent memory modules.

For server-specific steps used in this procedure, see the server maintenance and service guide for your product on the Hewlett Packard Enterprise website:

- HPE ProLiant Gen10 Plus servers (<https://www.hpe.com/info/proliantgen10plus-docs>)
- HPE Synergy Gen10 Plus compute modules (<https://www.hpe.com/info/synergy-docs>)

Prerequisites

- Observe the [Persistent memory module relocation guidelines](#).
- If persistent memory modules are encrypted, see one of the following:
 - [Migrating a persistent memory module encrypted with local key management](#)
 - [Migrating a persistent memory module encrypted with remote key management](#)

Procedure

1. Power down the server.
2. Do one of the following:
 - Extend the server from the rack.
 - Remove the server from the rack, if necessary.
 - Remove the server or server blade from the enclosure.
3. Place the server on a flat, level work surface.
4. Access the DIMM slots.
5. Be sure to note the slot locations in which each DIMM and persistent memory module are installed, and then remove the components from the server.
6. Remove the remainder of the components from the system board, and then remove the system board.
7. Install the spare system board.

 **IMPORTANT:** Install all components with the same configuration that was used on the failed system board.

Install all components removed from the failed system board.

Be sure to install the DIMMs and persistent memory modules in the same locations as the old system board.

9. Install the access panel.
10. Power up the server.
11. Ensure all firmware, including option cards and embedded devices, is updated to the same versions to ensure that the latest drivers are being used.
12. If required, re-enter the server serial number and the product ID.

For more information, see the server maintenance and service guide on the Hewlett Packard Enterprise website:

- HPE ProLiant Gen10 Plus servers (<https://www.hpe.com/info/proliantgen10plus-docs>)
- HPE Synergy Gen10 Plus compute modules (<https://www.hpe.com/info/synergy-docs>)

Migrating a persistent memory module

i IMPORTANT:

When data must be preserved, Hewlett Packard Enterprise strongly recommends that you perform a manual backup of all user data on the persistent memory modules before changing the goal configuration or performing relocation procedures.

Migrating a persistent memory module encrypted with local key management

Use this process to migrate persistent memory modules encrypted with local key management. If remote key management is enabled, see [Migrating a persistent memory module encrypted with remote key management](#).

Prerequisites

- Before migrating an encrypted persistent memory module, review and observe the [Persistent memory module relocation guidelines](#).
- To migrate an encrypted persistent memory module, you must obtain the passwords for the encrypted persistent memory modules by doing one of the following:
 - Exporting a password file to a USB key (recommended).
 - Manually retrieving and recording the passwords for each persistent memory module from the server.

Procedure

1. On the server from which the persistent memory module is to be migrated, press the F9 key during POST to access System Utilities.
2. To export the password file to a USB key, do the following:
 - a. From the System Utilities screen, select System Configuration > BIOS/Platform Configuration (RBSU) > Server Security > Device Encryption Options > Device Encryption Migration Options > Device Encryption Export Options.
 - b. Provide a password in the Transient Passphrase field.

This password protects the exported file and must be entered when recovering the encrypted persistent memory modules after relocation.
 - c. Select Select File, and browse to a location on the USB key.
 - d. Select Export Encryption Settings to create and export the file.
3. To manually record the persistent memory module passwords, do the following:
 - a. From the System Utilities screen, select System Configuration > BIOS/Platform Configuration (RBSU) > Server Security > Device Encryption Options > Device Encryption Status.
 - b. Record the password and the location displayed next to each persistent memory module.

You must supply these passwords for the same locations when recovering the encrypted persistent memory modules after relocation.
4. Power down the server.
5. Be sure to note the slot locations in which each DIMM and persistent memory module are installed, and then remove the components from the server.
6. Install the DIMMs and persistent memory modules in the new server or on the new system board.

Be sure to observe the relocation guidelines when installing the DIMMs and persistent memory modules.
7. Power up the server.
8. During POST, press F9 to access System Utilities.
9. From the System Utilities screen, select System Configuration > BIOS/Platform Configuration (RBSU) > Server Security > Device Encryption Options > Device Encryption Migration Options > Device Encryption Recovery Options.
10. To unlock the persistent memory modules using the exported file on a USB key, do the following:
 - a. Connect the USB key.
 - b. Select Restore Passphrases from USB Device.
 - c. Enter the Transient Passphrase created when exporting the file.
 - d. Select Select File, and browse to the location of the password file on the USB key.
 - e. Select Restore Encryption Settings to import the file.
11. To unlock the persistent memory modules for which the passwords were manually recorded, do the following:
 - a. Select Manually Unlock Devices.

- b. Select a persistent memory module and enter the passphrase.
 - c. Press Enter.
 - d. Repeat this process for each encrypted persistent memory module in the list.
12. To save your changes and exit, press the **F12** key.

Migrating a persistent memory module encrypted with remote key management

Use this process to migrate persistent memory modules encrypted with remote key management. If local key management is enabled, see [Migrating a persistent memory module encrypted with local key management](#).

Prerequisites

Before migrating an encrypted persistent memory module, review and observe the [Persistent memory module relocation guidelines](#).

Procedure

1. Log in to HPE iLO.
2. Click Administration in the navigation tree, and then click the Key Manager tab.
3. Make note of the following Key Manager Server entries.

This information is required to complete migration:

- Primary Key Server Address
- Primary Key Server Port
- Secondary Key Server Address (if entered)
- Secondary Key Server Port (if entered)

4. Make note of the Group name under **Key Manager Configuration**.

This information is required to complete migration.

5. Power down the server.
6. Be sure to note the slot locations in which each DIMM and persistent memory module are installed, and then remove the components from the server.
7. Install the DIMMs and persistent memory modules in the new server or on the new system board.

Be sure to observe the relocation guidelines when installing the DIMMs and persistent memory modules.

8. Power up the server.
9. Log in to HPE iLO.
10. Click Administration in the navigation tree, and then click the Key Manager tab.
11. Make the following entries under Key Manager Server entries.

This information must match the information recorded from the old server:

- Primary Key Server Address
- Primary Key Server Port
- Secondary Key Server Address (if entered)
- Secondary Key Server Port (if entered)

12. Enter the Group name under **Key Manager Configuration**.

This entry must match the Group name recorded from the old server.

13. Reboot the server.
14. During POST, press F9 to access System Utilities.
15. From the System Utilities screen, select **System Configuration > BIOS/Platform Configuration (RBSU) > Server Security > Device Encryption Options**.
16. Select Remote from the Key Management menu.
17. Press the **F12** key to save and exit.

Sanitizing a persistent memory module



Sanitization policies

The System Utilities menu item in Sanitize Options for **Policy After Sanitize/Erase on Reboot** provides the following options:

- Sanitize/Erase and Reboot System—Use this policy for the following scenarios:
 - After you add a new persistent memory module to a server.
 - If a persistent memory module is mapped out due to errors and you want to use the persistent memory module again.
 - After you move persistent memory modules previously used in another server to a new server.
Even if everything in the new server is an exact match to the previous server, you must sanitize the persistent memory module.
- Sanitize/Erase and Power System Off—Use this policy for the following scenarios:
 - Decommissioning a persistent memory module.
 - Recommissioning a persistent memory module (move to another server with no requirement to preserve the data).
- Sanitize/Erase and Reboot to System Utilities—Use this policy to change any BIOS/Platform Configuration (RBSU) setting that results in the data on the persistent memory module no longer being interpreted the same way. Examples include Persistent Memory Interleaving.
- Sanitize/Erase to Factory Defaults and Power System Off—Use this policy when retiring a persistent memory module or returning the persistent memory module to Hewlett Packard Enterprise (service replacement).

After selecting a sanitize policy and one or more persistent memory modules to sanitize, the system upgrades all warm reset requests into cold resets. The first cold reset:

1. Flushes any write data still pending in processor write buffers to DRAM.
2. Maps out the persistent memory modules.
3. Sends the sanitize command to the persistent memory modules.

Sanitization policy	After completing the sanitize commands, the system:
Sanitize/Erase and Reboot System	Reboots the server after completing the Sanitize commands.
Sanitize/Erase and Power System Off	Powers off after completing the Sanitize commands.
Sanitize/Erase and Reboot to System Utilities	Performs another cold reset to map in the persistent memory modules again.
Sanitize/Erase to Factory Defaults and Power System Off	Powers off after completing the Sanitize commands.

Sanitization guidelines

All scenarios mentioned assume that the DIMM and persistent memory module population guidelines are followed.

Scenarios in which sanitization is required before using the persistent memory modules

- When a new persistent memory module is added to the system, sanitize the new persistent memory module before the persistent memory module can be used.
- When removing persistent memory modules from a server with Persistent Memory Interleaving set to Enabled, sanitize all the persistent memory modules on the processor where the persistent memory modules were removed.
- When a previously used persistent memory module is added to the system, do one of the following:
 - If the Persistent Memory Interleaving setting is set to Enabled, then sanitize all the persistent memory modules on that processor before using the persistent memory modules.
 - If the Persistent Memory Interleaving setting is set to Disabled, then no sanitization is required.
- When Persistent Memory Interleaving settings are changed, sanitize all the persistent memory modules in the server.

Scenarios in which sanitization might not be required

These scenarios cover migrating persistent memory modules to keep the data and access it in the new server.

- The persistent memory module was already in use in another server that matches the new server in both hardware and System Utilities settings.
- The persistent memory module is installed in the new server in the same DIMM slot as in the original server.
- If persistent memory modules are used when Persistent Memory Interleaving is set to Enabled, install all the persistent memory modules in the Interleave set in the same DIMM slots in the new server.
- If the persistent memory module is used with Persistent Memory Interleaving set to Disabled, install the persistent memory module in any slot on the server.

Sanitization with UEFI System Utilities

Review the sanitization policies and guidelines in this guide before sanitizing a persistent memory module.

Procedure

1. From the System Utilities screen, select **System Configuration > BIOS/Platform Configuration (RBSU) > Memory Options > Persistent Memory Options > PMEM Options > Sanitize Options**, and select the following:
 - **Sanitize/Erase Operation on Reboot:**
 - No Action
 - Cryptographically Erase
 - Overwrite Media
 - Cryptographically Erase then Overwrite Media
 - **Policy after Sanitize/Erase on Reboot:**
 - Sanitize/Erase and Reboot System
 - Sanitize/Erase and Power System Off
 - Sanitize/Erase and Reboot to System Utilities
 - Sanitize/Erase to Factory Defaults and Power System Off
2. Enable the **Perform Sanitize/Erase Operation on** selection.
3. Select the persistent memory modules to sanitize:
 - All PMMs in the System—Sanitizes all persistent memory modules in the server.
 - All PMMs on Processor X—Sanitizes all persistent memory modules on the specified processor.
 - Processor X DIMM Y—Sanitizes only the specified persistent memory module on the processor.
4. To save your changes and exit, press the **F12** key.
5. If required, reboot the server.

Sanitization with the HPE iLO RESTful API

To sanitize a persistent memory module using the HPE iLO RESTful API, use the associated commands.

Command	System Utilities option
PmmSanitizeOperation	Sanitize/Erase Operation on Reboot
NoAction	No Action
CryptoErase	Cryptographically Erase
Overwrite	Overwrite Media
CryptoEraseOverwrite	Cryptographically Erase and Overwrite Media
PmmSanitizePolicy	Policy after Sanitize/Erase on Reboot
SanitizeAndRebootSystem	Sanitize/Erase and Reboot System
SanitizeAndShutdownSystem	Sanitize/Erase and Power System Off
SanitizeAndBootToFirmwareUI	Sanitize/Erase and Reboot to System Utilities
SanitizeToFactoryDefaults	Sanitize/Erase to Factory Defaults and Power System Off
SanitizeAllPmm	Perform Sanitize/Erase Operation on: All persistent memory modules in the System
SanitizeProcXPmm¹	Perform Sanitize/Erase Operation on: All persistent memory modules on Processor X
SanitizeProcXPmmY¹	Perform Sanitize/Erase Operation on: Processor X DIMM Y

¹ Where X and Y represent the processor and DIMM slot number, such as `SanitizeProc1Pmm4`.



Sanitization with ipmctl

The ipmctl tool can be used to sanitize persistent memory modules under the following conditions:

- Persistent memory modules must not be in one of the following states:
 - Unlocked, Frozen
 - Disabled, Frozen
 - Exceeded
- Any namespace associated with the specified persistent memory modules must first be deleted.

To erase the persistent data on all persistent memory modules in the server, run the following command:

```
ipmctl delete -dim
```

Recommissioning a persistent memory module with a lost password

If the passphrase for the persistent memory module is unknown, and you are **NOT** required to retain or access any data stored on it, sanitize the persistent memory module with the Cryptographically Erase option to reuse the module.

This process does **NOT** allow you to retain or access any data previously stored on the persistent memory module. It only allows you to regain use of the hardware.

For more information, see [Sanitizing a persistent memory module](#).

Updating persistent memory module firmware

To update persistent memory module firmware, use one of the following methods:

- Service Pack for ProLiant (SPP)—See the Service Pack for ProLiant Quick Start Guide (<https://www.hpe.com/info/spp/documentation>).

To download the SPP, see the Hewlett Packard Enterprise website (<https://www.hpe.com/servers/spp/download>).

- HPE online flash components

Linux support of Intel Optane persistent memory 200 series for HPE

Linux presents persistent memory module using the device types described in the following table.

Linux device paths

Path	Name	Type	Notes
/dev/pmem*	Persistent memory with file system DAX	Block device	
/dev/pmem*s	Sector atomic	Block device	
/dev/dax*.*	Device DAX	Character device	<ul style="list-style-type: none">• For specialized software• Does not support file systems

The following table shows the interim layers of devices used by the persistent memory drivers.

Linux driver stack devices

Type	Path	Description
nmem	/sys/bus/nd/devices/nmem*	Represents a persistent memory module
Region	/sys/bus/nd/devices/region*	Represents a memory region presented by either an interleaved set of persistent memory modules or a single persistent memory module
Block device	/sys/block/pmem*	Represents filesystem DAX and regular block devices
Device DAX	/sys/class/dax/dax*	Represents device DAX character devices

nmem devices

In Linux, persistent memory modules are represented by nmem devices.

nmem device properties

An nmem device has several properties, including:

- Dev—The Linux driver device name (such as nmem0).
- ID—The serial number printed on the physical label.
- Handle—A system firmware-generated unique identifier for the device.
- Phys_id—The location of the DIMM encoded in hexadecimal.
- Security—The security state of the persistent memory module (disabled, unlocked, locked, frozen, or overwrite).

Listing nmem devices

To show a list of persistent memory modules and their properties, run:

```
ndctl list --human --dimms
[
  {
    "dev": "nmem1",
    "id": "8089-a2-1839-12345678",
    "handle": "0x11",
    "phys_id": "0x27",
    "security": "disabled"
  },
  {
    "dev": "nmem3",
    "id": "8089-a2-1839-87654321",
    "handle": "0x101",
    "phys_id": "0x24",
    "security": "disabled"
  }
]:
]
```

Regions

A region is a portion of system memory presented from one or more persistent memory modules. A region consists of one of the following:

- An interleaved set (with each persistent memory module contributing the same capacity)
- A single persistent memory module

In Linux, regions are set by the goal configuration. Regions are named `regionRR`, where `RR` is any number from 0. The maximum number of regions is the number of persistent memory modules or interleaved sets.

Region properties

A region device has several properties:

- `Dev`—Identifier for the region (for this boot).
- `Size`—Capacity of the persistent memory presented by this region.
- `Available_size`—Size not currently allocated to namespaces.
- `Max_available_extent`—Maximum contiguous size that can be allocated to a namespace.
- `Type`—Always persistent memory.
- `Numa_node`—A `numa_node` ID for the region. This ID can be used to bind processors close to this region using `numactl`.
- `lset_id`—A worldwide unique ID for the region.

If a set of persistent memory modules is moved to another server, this ID remains the same, and ensures that the full set of the persistent memory modules are still together.

- `Persistence_domain`—Always set to memory controller in HPE ProLiant and HPE Synergy Gen10 Plus server products.

 **NOTE:**

SUSE Linux Enterprise Server 15 GA assigns an extra region number for each persistent memory module in the server, and these get the lowest numbers. Hewlett Packard Enterprise recommends installing the kernel update SUSE-SU-2019:0224-1, which corrects this convention.

Listing regions

To show a list of regions and their properties, run:

```
ndctl list --human --regions
[
  {
    "dev":"region1",
    "size":"502.00 GiB (539.02 GB)",
    "available_size":0,
    "max_available_extent":0,
    "type":"pmem",
    "numa_node":0,
    "iset_id":"0x12ccda9021308a22",
    "persistence_domain":"memory_controller"
  },
  {
    "dev":"region3",
    "size":"502.00 GiB (539.02 GB)",
    "available_size":"374.00 GiB (401.58 GB)",
    "max_available_extent":"374.00 GiB (401.58 GB)",
    "type":"pmem",
    "numa_node":0,
    "iset_id":"0x5ed6da900f318a22",
    "persistence_domain":"memory_controller"
  }
]:
]
```

Namespaces

A namespace is a portion of a region. In Linux, a namespace is managed with `ndctl`.

Namespaces are numbered `namespace<regionRR>.<NN>`, in which:

- `regionRR` is the region device name from which the namespace is created.
- `NN` is the number of the namespace, ranging from 0 to 63.

The following table shows the types of namespaces supported by the Linux kernel persistent memory drivers.

Mode	Name	OS	Description
raw	Raw	All	<ul style="list-style-type: none">• Creates a <code>/dev/pmem0</code> block device.• Supports all filesystems without the DAX option• Presents raw persistent memory when <code>Apply Default Namespaces</code> is set to <code>Enabled</code> in BIOS/Platform Configuration (RBSU).
sector	Sector atomic	All	<ul style="list-style-type: none">• Creates a <code>/dev/pmem0s</code> block device.• Supports all filesystems without the DAX option• Block Translation Table used to provide sector (for example, 512 byte or 4,096 byte) atomicity.• If power is lost while in the middle of writing to a block, reverts to the previous contents.
fsdax	Filesystem DAX	Linux	<ul style="list-style-type: none">• Creates a <code>/dev/pmem0</code> block device.• Supports filesystems offering the DAX option - <code>ext4</code> and <code>xfs</code>.• When mounted with <code>-o dax</code> option, applications get direct access to persistent memory by removing the page cache from the I/O path.
devdax	Device DAX	Linux	<ul style="list-style-type: none">• Creates a <code>/dev/dax0.0</code> character device for persistent memory-aware applications for the lowest software overhead.• No filesystem support.• There is no <code>read()</code> or <code>write()</code> support, only <code>mmap()</code>.

Namespace properties

A namespace device has several properties:

- **Dev**—Unique device name for this namespace, based on the region name (such as namespace6.0).
- **Mode**—raw, sector, fsdax, or devdax.
- **Size**—Capacity of this namespace.
- **uuid**—A worldwide unique identifier for the namespace.

Since the namespace device name and region names can change based on the presence of other regions, they are not safe to use in scripts.

- **Sector**—Logical block size.
- **Blockdev**—Name of the `/dev/pmemNN` block device using this namespace, if any.
- **Chardev**—Name of the `/dev/daxNN.MM` character device using this namespace, if any.
- **Numa_node**—A `numa_node` ID for the namespace. This ID can be used to bind processors close to this namespace using `numactl`.

Creating a namespace

When creating namespaces, you can specify options for size and regions. If you do not specify the size, the operation allocates the maximum size.

Example: To create the entire fsdax namespace from region 0, run:

```
$ sudo ndctl create-namespace -m fsdax -r region0
```

Example: To create a 32 GB raw namespace from region 1, run:

```
$ sudo ndctl create-namespace -m raw -s 32G -r region1
```

List all namespaces

To list all namespaces, run:

```
$ ndctl list
```

To show a list of namespaces and their properties, run:

```
# ndctl list --human --namespaces
[
  {
    "dev": "namespace1.0",
    "mode": "fsdax",
    "map": "dev",
    "size": "494.15 GiB (530.59 GB)",
    "uuid": "ff189419-de3d-406d-8f7f-812696a25ca8",
    "raw_uuid": "24841e1f-ab7e-43e5-a2fd-695af39bb682",
    "sector_size": 512,
    "blockdev": "pmem1",
    "numa_node": 0
  },
  {
    "dev": "namespace3.0",
    "mode": "raw",
    "size": "128.00 GiB (137.44 GB)",
    "uuid": "ba1733ea-782a-441a-91a3-e9c0af088752",
    "sector_size": 512,
    "blockdev": "pmem3",
    "numa_node": 0
  },
  :
]
```

Changing namespace mode

Run this command to change the namespace mode of an existing namespace.

Example: To change existing namespace0.0 to "fsdax," run:

```
$ sudo ndctl create-namespace -f -e namespace0.0 -m fsdax
```

Deleting namespaces

To delete a namespace, run:

```
$ sudo ndctl disable-namespace namespace0.0  
$ sudo ndctl destroy-namespace --force namespace0.0
```



Initialization of pmem devices

Depending on the number and capacity of persistent memory modules in the server, persistent memory modules might not have completed initialization when the login prompt appears.

Wait for the initialization to complete. You can use the `list` command to confirm that persistent memory modules have initialized:

```
ndctl list
```

The command returns a list of namespaces when the persistent memory modules have completed initialization. If the command does not immediately return the information, initialization is still in progress.

Showing the amount of memory in the system

The `free` command shows the amount of memory in the system, not counting memory reserved for processor page tables. With large persistent memory module capacities, this becomes quite noticeable.

The `-h` (or `--human`) option reports the capacity in human-readable format including units (the default unit is KiB units):

```
$ free -h
              total        used         free       shared    buff/cache   available
Mem:           62G          423M          59G           2.1M           2.9G           61G
Swap:          7.8G           0B           7.8G
```

The numbers are subject to truncation and rounding, and the units display as G but should be interpreted as Gi.

The `-b` option prints the precise sizes in bytes:

```
$ free -b
              total        used         free       shared    buff/cache   available
Mem:    67403063296  444485632  63883395072  2240512  3075182592  66085310464
Swap:    8388603904           0  8388603904
```

Filesystems

You can place any filesystem (such as, ext4, xfs, btrfs) on a pmem block device.

ext4 and xfs support the DAX mount option (-o dax), which allow applications to perform direct access by removing the page cache from the I/O path. Using this DAX option requires a pmem block device set to the fsdax namespace mode.

The following example creates ext4, xfs, and btrfs filesystems on three pmem block devices, and mounts ext4 and xfs with DAX.



NOTE:

If you are using RHEL8, first disable the `reflink` feature. To disable the feature, run the following command:

```
sudo mkfs.xfs -m reflink=0 /dev/pmem0
```

```
$ sudo mkfs.ext4 -F /dev/pmem0
$ sudo mount -o dax /dev/pmem0 /mnt/pmem0

$ sudo mkfs.xfs -f /dev/pmem1
$ sudo mount -o dax /dev/pmem1 /mnt/pmem1

$ sudo mkfs.btrfs -f /dev/pmem2
$ sudo mount /dev/pmem2 /mnt/pmem2
```

To confirm that the DAX mount option was enabled, review effective mount options. The filesystem might drop the DAX option when a pmem block device is not set to fsdax mode.

```
$ mount | grep pmem
/dev/pmem0 on /mnt/pmem0 type ext4 (rw,relatime,dax,data=ordered)
/dev/pmem1 on /mnt/pmem1 type xfs (rw,relatime,attr2,dax,inode64,noquota)
/dev/pmem2 on /mnt/pmem2 type btrfs (rw,relatime,ssd,space_cache,subvolid=5,subvol=/)
```

I/O statistics

iostats are disabled by default due to performance overhead (for example, 12M IOPS dropping 25% to 9M IOPS). iostats can be enabled in sysfs.

iostats are collected only for the base pmem device, not per-partition. I/Os that go through DAX paths are not counted, so nothing is collected for I/O to files in filesystems that are mounted with `-o dax`.

```
$ echo 1 > /sys/block/pmem0/queue/iostats
$ echo 1 > /sys/block/pmem1/queue/iostats
$ echo 1 > /sys/block/pmem2/queue/iostats
$ echo 1 > /sys/block/pmem3/queue/iostats

$ iostat -mxy 1
avg-cpu:  %user   %nice %system %iowait  %steal   %idle
           21.53    0.00   78.47    0.00    0.00    0.00

Device:            rrqm/s   wrqm/s         r/s         w/s    rMB/s    wMB/s avgrq-sz avgqu-sz   await  r_await  w_await   svctm  %util
pmem0              0.00     0.00 4706551.00    0.00 18384.95     0.00     8.00     6.00    0.00    0.00    0.00    0.00 113.90
pmem1              0.00     0.00 4701492.00    0.00 18365.20     0.00     8.00     6.01    0.00    0.00    0.00    0.00 119.30
pmem2              0.00     0.00 4701851.00    0.00 18366.60     0.00     8.00     6.37    0.00    0.00    0.00    0.00 108.90
pmem3              0.00     0.00 4688767.00    0.00 18315.50     0.00     8.00     6.43    0.00    0.00    0.00    0.00 117.00
```

VMware support of Intel Optane persistent memory 200 series for HPE

Intel Optane persistent memory 200 series for HPE ARS error handling implementation may have significant impact to PMem availability in certain failure scenarios. The first 2MB of every namespace is used by ESXi for storing PMem's critical metadata. VMKernel PMem typically expects ARS clear error DSM to succeed. Uniquely in HPE's designs, ESXi is unable to clear an ARS error, and if an ARS error happens in this 2MB region, then ESXi cannot reformat and create PMem file system on the DIMMs of those namespaces. Once one namespace cannot be formatted, ESXi errors out and cannot create a PMem datastore on the host rendering all PMem in the host unusable.

The module(s) with uncorrectable error in the first 2MB of a namespace must be replaced before ESXi can create a PMem datastore. HPE IML logs will record an event when a uncorrectable error is detected on any memory module. IML logs can be used to determine which PMem is defective and should be replaced. Part replacement could lead to significant system downtime. Contact HPE for additional details.

All PMem implementations are required to perform integrity checks on PMem to ensure all errors are detected and reported before the OS can be allowed access the PMem. PMem implementations differ on when and how such integrity checks are performed. HPE system BIOS will perform a PMem integrity check on boot when necessary, such as when a system reset is caused by a fatal error. Depending on the amount of PMem memory installed in an HPE System, the PMem integrity check may take a long time to complete. For example, the PMem integrity check can take 1.5 hours to complete on a 2-socket system with 6TB PMem.

For information on using Intel Optane persistent memory 200 series for HPE with VMware, see the [VMware Docs website](#).

To find Hewlett Packard Enterprise servers that are VMware PMEM certified with Intel Optane persistent memory 200 series for HPE, see the [VMware Compatibility Guide](#).

Windows Server support of Intel Optane persistent memory 200 series for HPE

For information on using Intel Optane persistent memory 200 series for HPE with Windows Server, see the technical white paper [Deploying HPE Persistent Memory on Microsoft Windows Server 2012 R2, Server 2016, and Server 2019](#) on the [Hewlett Packard Enterprise website](#).





System boot fails due to persistent memory file systems

Symptom

The system boots to an emergency prompt/recovery shell when a large amount of persistent memory is installed and when fsdax is used to create namespaces.

Solution 1

Cause

The PMEM devices do not initialize within the auto-mount time defined in the `/etc/fstab` file in systems running the following versions:

- Red Hat Enterprise Linux 7.x
- Red Hat Enterprise Linux 8.0 (without [RHSA-2019:1959](#))
- SUSE Linux Enterprise Server 12 SPx
- SUSE Linux Enterprise Server 15

Action

To work around the issue, increase the `DefaultTimeoutStartSec` value in the `/etc/systemd/system.conf` file to a sufficiently large value, such as `1200s`.

The system boot will no longer time out.

Solution 2

Cause

In systems running SUSE Linux Enterprise Server 12 SP4, configuring a large amount of PMEM devices can delay loading of the `btrfs` module.

Action

Force the `libnvdimm` module to load after the `btrfs` kernel module by adding the following entry to, for example, "`/etc/modprobe.d/99-local.conf`":

```
# Load btrfs before libnvdimm
softdep libnvdimm pre: btrfs
```

For more information, see <https://www.suse.com/support/kb/doc/?id=7024085>.

Troubleshooting resources

Troubleshooting resources are available for HPE Gen10 and Gen10 Plus server products in the following documents:

- Troubleshooting Guide for HPE ProLiant Gen10 and Gen10 Plus servers provides procedures for resolving common problems and comprehensive courses of action for fault isolation and identification, issue resolution, and software maintenance.
- Integrated Management Log Messages and Troubleshooting Guide for HPE ProLiant Gen10 and Gen10 Plus servers and HPE Synergy provides IML messages and associated troubleshooting information to resolve critical and cautionary IML events.

To access troubleshooting resources for your product, see the [Hewlett Packard Enterprise website](#).



Websites

General websites

Persistent Memory websites

Hewlett Packard Enterprise Information Library for Persistent Memory

www.hpe.com/info/persistentmemory-docs

HPE Persistent Memory Portfolio

www.hpe.com/info/persistentmemory

For additional websites, see [Support and other resources](#).



Accessing Hewlett Packard Enterprise Support

- For live assistance, go to the Contact Hewlett Packard Enterprise Worldwide website:

<https://www.hpe.com/info/assistance>

- To access documentation and support services, go to the Hewlett Packard Enterprise Support Center website:

<https://www.hpe.com/support/hpesc>

Information to collect

- Technical support registration number (if applicable)
- Product name, model or version, and serial number
- Operating system name and version
- Firmware version
- Error messages
- Product-specific reports and logs
- Add-on products or components
- Third-party products or components



Accessing updates

- Some software products provide a mechanism for accessing software updates through the product interface. Review your product documentation to identify the recommended software update method.

- To download product updates:

Hewlett Packard Enterprise Support Center

<https://www.hpe.com/support/hpesc>

Hewlett Packard Enterprise Support Center: Software downloads

<https://www.hpe.com/support/downloads>

My HPE Software Center

<https://www.hpe.com/software/hpesoftwarecenter>

- To subscribe to eNewsletters and alerts:

<https://www.hpe.com/support/e-updates>

- To view and update your entitlements, and to link your contracts and warranties with your profile, go to the Hewlett Packard Enterprise Support Center More Information on Access to Support Materials page:

<https://www.hpe.com/support/AccessToSupportMaterials>

ⓘ IMPORTANT:

Access to some updates might require product entitlement when accessed through the Hewlett Packard Enterprise Support Center. You must have an HPE Passport set up with relevant entitlements.

Remote support

Remote support is available with supported devices as part of your warranty or contractual support agreement. It provides intelligent event diagnosis, and automatic, secure submission of hardware event notifications to Hewlett Packard Enterprise, which initiates a fast and accurate resolution based on the service level of your product. Hewlett Packard Enterprise strongly recommends that you register your device for remote support.

If your product includes additional remote support details, use search to locate that information.

HPE Get Connected

<https://www.hpe.com/services/getconnected>

HPE Pointnext Tech Care

<https://www.hpe.com/services/techcare>

HPE Complete Care

<https://www.hpe.com/services/completecure>



Warranty information

To view the warranty information for your product, see the links provided below:

HPE ProLiant and IA-32 Servers and Options

<https://www.hpe.com/support/ProLiantServers-Warranties>

HPE Enterprise and Cloudline Servers

<https://www.hpe.com/support/EnterpriseServers-Warranties>

HPE Storage Products

<https://www.hpe.com/support/Storage-Warranties>

HPE Networking Products

<https://www.hpe.com/support/Networking-Warranties>



Regulatory information

To view the regulatory information for your product, view the Safety and Compliance Information for Server, Storage, Power, Networking, and Rack Products, available at the Hewlett Packard Enterprise Support Center:

<https://www.hpe.com/support/Safety-Compliance-EnterpriseProducts>

Additional regulatory information

Hewlett Packard Enterprise is committed to providing our customers with information about the chemical substances in our products as needed to comply with legal requirements such as REACH (Regulation EC No 1907/2006 of the European Parliament and the Council). A chemical information report for this product can be found at:

<https://www.hpe.com/info/reach>

For Hewlett Packard Enterprise product environmental and safety information and compliance data, including RoHS and REACH, see:

<https://www.hpe.com/info/ecodata>

For Hewlett Packard Enterprise environmental information, including company programs, product recycling, and energy efficiency, see:

<https://www.hpe.com/info/environment>

Documentation feedback

Hewlett Packard Enterprise is committed to providing documentation that meets your needs. To help us improve the documentation, use the Feedback button and icons (located at the bottom of an opened document) on the Hewlett Packard Enterprise Support Center portal (<https://www.hpe.com/support/hpesc>) to send any errors, suggestions, or comments. All document information is captured by the process.

