

The Lenovo logo is displayed in white text on a black rectangular background.

Implementing Intel Optane Persistent Memory with VMware vSphere

Describes the hardware prerequisites for persistent memory

Explains how to configure PMem operation modes

Shows vSphere support of PMem operation modes

Provide instructions on how to configure and use PMem modules in vSphere

Chengcheng Peng
Skyler Zhang



Abstract

Intel Optane Persistent Memory is an innovative memory technology that redefines traditional architectures, delivering a unique combination of affordable large capacity and support for data persistence. It has the speed characteristics of memory, but it retains data through power cycles. Intel Optane Persistent Memory Modules (PMem modules) are available on Lenovo® ThinkSystem™ servers and currently come in capacities of 128 GB, 256 GB, and 512 GB.

This paper describes how to configure and use persistent memory with different operation modes in VMware vSphere on Lenovo ThinkSystem servers. This document is intended for technical specialists, and IT administrators who are familiar with persistent memory and vSphere products.

At Lenovo Press, we bring together experts to produce technical publications around topics of importance to you, providing information and best practices for using Lenovo products and solutions to solve IT challenges.

See a list of our most recent publications at the Lenovo Press web site:

<http://lenovopress.com>

Do you have the latest version? We update our papers from time to time, so check whether you have the latest version of this document by clicking the **Check for Updates** button on the front page of the PDF. Pressing this button will take you to a web page that will tell you if you are reading the latest version of the document and give you a link to the latest if needed. While you're there, you can also sign up to get notified via email whenever we make an update.

Contents

Introduction	3
Hardware setup	3
Persistent memory configuration	4
Firmware update	5
Check Persistent memory status	5
Creating a goal to configure Persistent memory operation modes	7
vSphere support of Persistent memory modes	8
PMem mode support	9
Using PMem modules in Memory Mode	9
Using PMem modules in App Direct Mode	11
Using PMem modules in Mixed Mode	20
References	21
Authors	22
Change history	22
Notices	23
Trademarks	24

Introduction

Intel Optane Persistent Memory represents a new class of memory and storage technology explicitly architected for data center usage. It's designed to improve the overall performance of a data center system by providing large amounts of persistent storage at near memory speeds.

Persistent memory modules (PMem modules) are DDR4 socket compatible and are available in sizes of 128 GB, 256 GB, and 512 GB per module. For example, a two-socket system can have up to 6 TB of PMem, 3 TB per CPU.

Two generations of Persistent memory are available:

- ▶ Persistent Memory 100 Series, supported on 2nd Gen Intel Xeon Scalable processors
- ▶ Persistent Memory 200 Series, supported on 3rd Gen Intel Xeon Scalable processors

Persistent memory offers performance advantages such as significantly lower latency than fetching data from system storage (SSD or HDD), high capacities and affordable cost.

Persistent memory can be configured in three modes:

- ▶ Memory Mode: The PMem modules act as large capacity DDR4 memory modules.
- ▶ App Direct Mode: The PMem modules provide all persistence features to the operating system and applications that support them.
- ▶ Mixed Mode (only supported in PMem 100 Series): Mixed Mode is a combination of Memory Mode and App Direct Mode, where a portion of the capacity of the PMem modules is used for the Memory Mode operations, and the remaining capacity of the PMem modules is used for the App Direct Mode operations. In this mode, all installed DRAM DIMMs are hidden from the operating system and act as a caching layer for portion of the PMem modules in Memory Mode.

More information about Persistent memory can be found in the Lenovo Press product guides:

- ▶ Persistent Memory 100 Series:
<https://lenovopress.lenovo.com/lp1066-intel-optane-persistent-memory-100-series>
- ▶ Persistent Memory 200 Series:
<https://lenovopress.lenovo.com/lp1380-intel-optane-persistent-memory-200-series>

Hardware setup

There are many hardware requirements when setting up Persistent memory. This section presents the hardware prerequisites needed to use PMem modules in a ThinkSystem server.

- ▶ Server selection

PMem modules are only supported in the second-generation Intel Xeon Scalable processors. PMem modules are not supported in the first-generation Xeon Scalable processors. Refer to the Persistent memory product guides to select a server to support PMem modules:

<https://lenovopress.com/lp1066-intel-optane-persistent-memory-100-series>
<https://lenovopress.com/lp1380-intel-optane-persistent-memory-200-series>

- ▶ Processor selection

All Platinum processors, all Gold processors, and the Silver 4215 processor support PMem modules. Refer to the Processor support section of the product guides to select correct processors to support Persistent memory:

<https://lenovopress.com/lp1066#processor-support>

<https://lenovopress.com/lp1380#processor-support>

- ▶ PMem modules and DRAM DIMMs selection

When selecting PMem modules and DRAM DIMMs for use in the server, they should meet the following requirements:

- All installed PMem modules must be the same size. Mixing PMem modules of different capacities is not supported.

- All installed DRAM DIMMs must be the same size and structure (i.e. same part number). Mixing different DRAM DIMMs is not supported.

Refer to the Memory DIMM support section of the product guides to get more details about Memory DIMM support:

<https://lenovopress.com/lp1066#memory-dimm-support>

<https://lenovopress.com/lp1380#memory-dimm-support>

- ▶ PMem modules and DRAM DIMMs installation

Refer to the following guide to install the PMem modules and memory DRAMs correctly on Lenovo ThinkSystem servers:

https://pubs.lenovo.com/sr650/dimm_installation_dcpmm

Persistent memory configuration

Before configure Persistent memory, make sure that the capacity of installed PMem modules and DRAM DIMMs meets system requirements for the following different operation modes:

- ▶ Memory mode requirements: Refer to the Memory mode requirements section of the product guides to get details about memory mode requirements:

<https://lenovopress.com/lp1066#memory-mode-requirements>

<https://lenovopress.com/lp1380#memory-mode-requirements>

- ▶ App Direct Mode requirements: Refer to App Direct Mode requirements section of the product guides to get details about App Direct Mode requirements:

<https://lenovopress.com/lp1066#app-direct-mode-requirements>

<https://lenovopress.com/lp1380#app-direct-mode-requirements>

- ▶ Mixed Mode requirements (100 Series only): Refer to Mixed Mode requirements section to get details about Mixed Mode requirements:

<https://lenovopress.com/lp1066#mixed-mode-requirements>

Firmware update

It is recommended that you update the firmware of BMC, UEFI and PMem modules to the latest version before configuring persistent memory, and that you make sure that Persistent memory FW and UEFI FW use the same Intel BKC version.

- ▶ BMC and UEFI Firmware Update
- ▶ PMem modules Firmware Update

Refer to the following guide to update the Persistent memory firmware to the latest version:

https://sysmgmt.lenovofiles.com/help/topic/com.lenovo.lxca.doc/update_fw.html

Check Persistent memory status

When you have installed the PMem modules in Lenovo ThinkSystem server and upgraded the BMC, UEFI and Persistent memory firmware, reboot the server to validate the PMem modules. After the validation, you can check the status of the PMem modules in XClarity Controller as shown in Figure 1 on page 5.

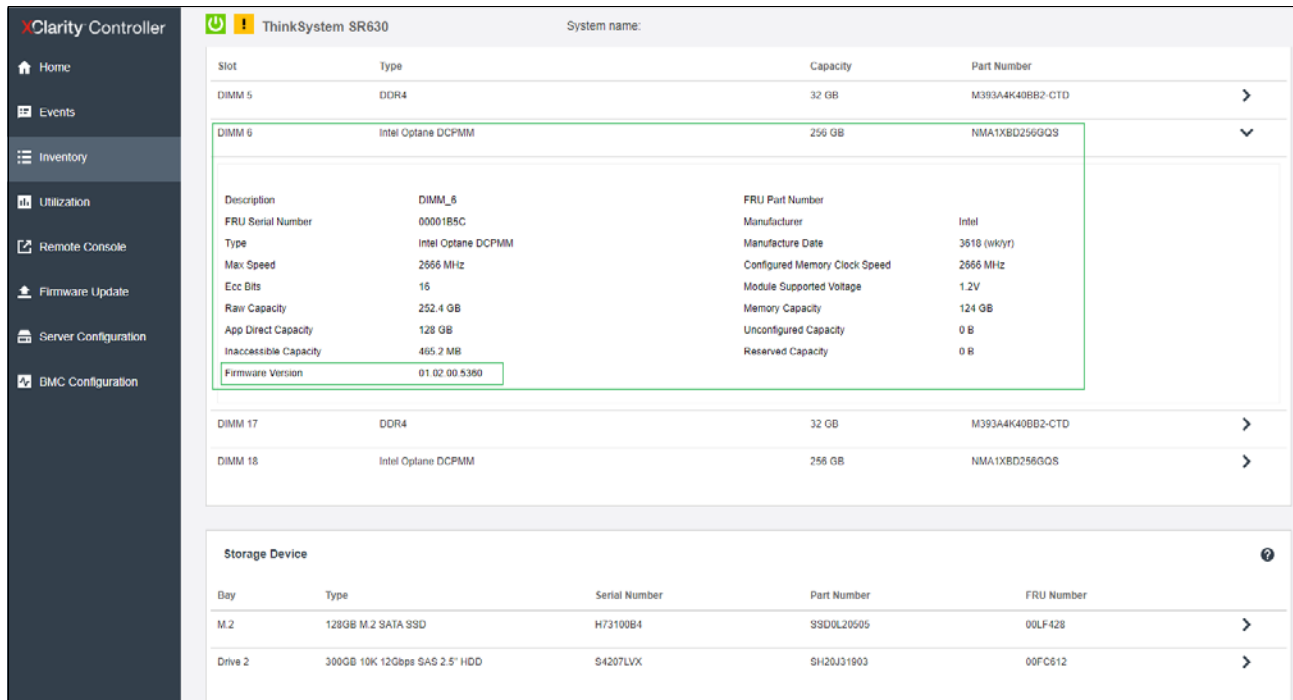


Figure 1 Persistent memory status in XCC

You can also check the Persistent memory detail information including goals, regions, namespaces, security, and configuration in Lenovo UEFI.

For Persistent Memory 100 Series, power on Lenovo ThinkSystem server and then press F1 to enter System Setup. Select **System Settings** → **Intel Optane PMem modules** where you can check PMem modules details as shown in Figure 2 on page 6.

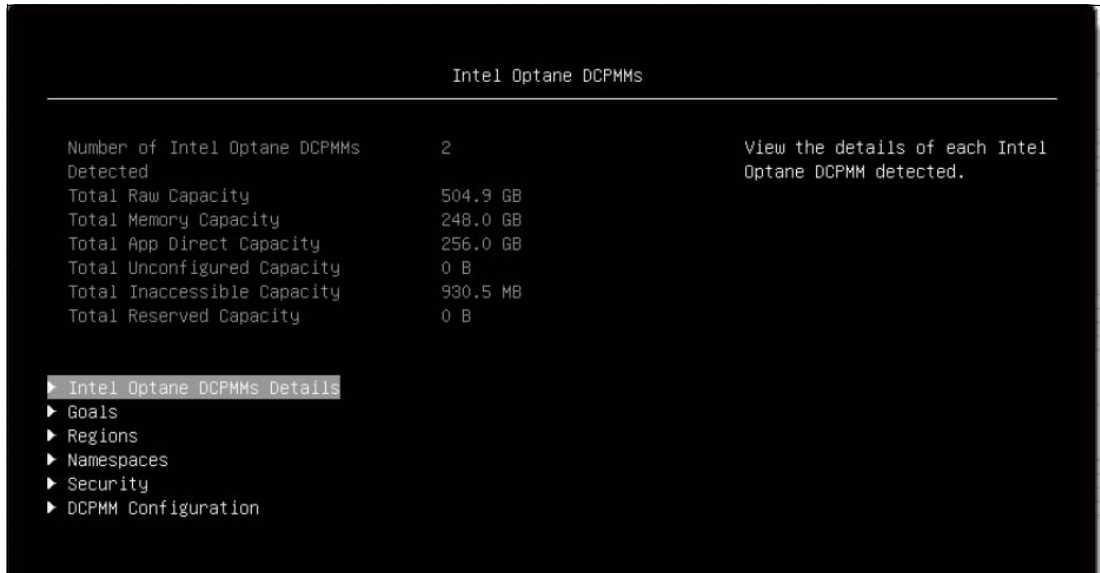


Figure 2 Persistent Memory 100 Series details in UEFI setting

For Persistent Memory 200 Series, power on Lenovo ThinkSystem server and then press F1 to enter System Setup. Select **System Settings** → **Intel Optane PMEMs** where you can check PMem details as shown in Figure 3.

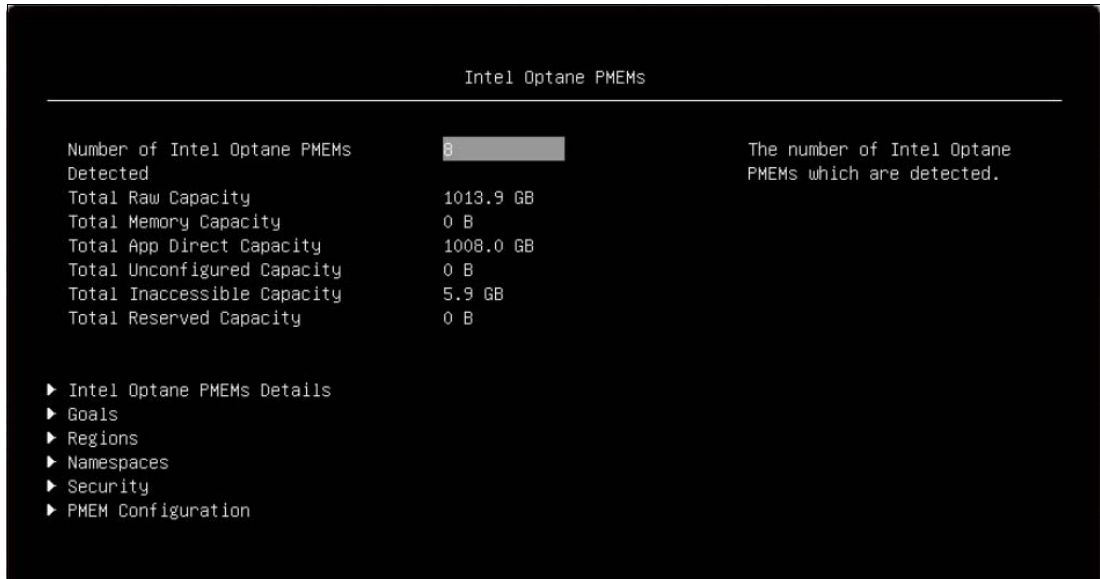


Figure 3 Persistent Memory 100 Series details in UEFI setting

Creating a goal to configure Persistent memory operation modes

A goal is a mechanism to allocate PMem modules capacity for Memory mode and App-direct mode. User must choose platform or processor for the unit and choose a percent to allocate PMem modules capacity in Memory mode. The remaining PMem modules capacity will be in App-direct mode.

It is recommended you create a goal via UEFI settings:

1. Power on the Lenovo server and press F1 when prompted to enter System Setup.
2. Go to **System Settings** → **Intel Optane PMem modules** → **Goals**, as shown in Figure 4 on page 7.

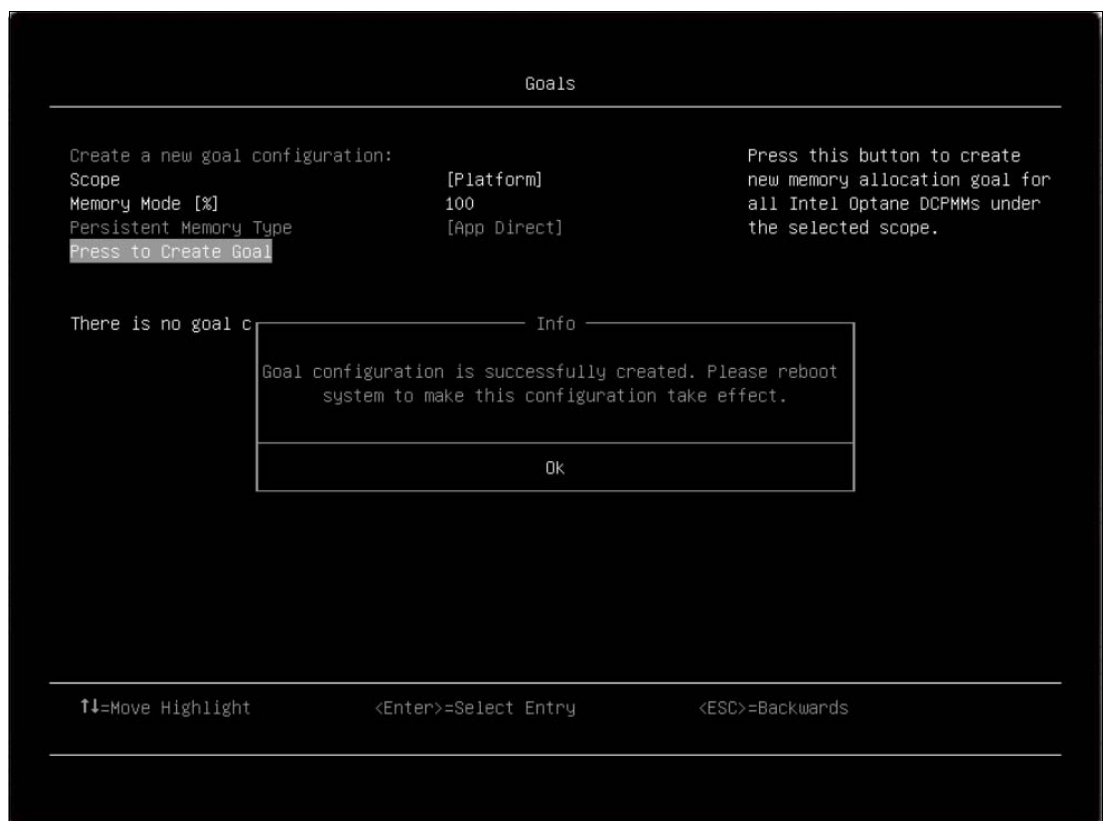


Figure 4 Create Goal in UEFI setting

3. In the "Memory Mode [%]" field, specify the percentage of Persistent memory capacity that is to be allocated to system memory, and hence decide the Persistent memory mode:
 - Memory Mode: Input 100 as Memory Mode [%]
 - App Direct Mode: Input 0 as Memory Mode [%]
 - Mixed Memory Mode: Input 1-99 as Memory Mode [%]
4. Click **Press to create Goal** to create a goal.
5. Reboot host to make goal take effect.

For more details about Persistent memory configuration on Lenovo ThinkSystem servers, you can refer to the following InfoCenter page:

https://thinksystem.lenovofiles.com/help/index.jsp?topic=%2F7X21%2Fmemory_configuration_Persistent_memory.html&cp=2_0_7_4_0

vSphere support of Persistent memory modes

VMware vSphere uses PMEM (persistent memory) to describe Intel PMem modules. This section describes how to configure and use Intel Persistent memory in vSphere 6.7 EP 10 (Build #13981272) and above on a ThinkSystem server.

VMware supports Intel Persistent memory in both Memory Mode and App Direct Mode. vSphere 6.5 U3 enables Intel Persistent memory in Memory Mode. vSphere 6.7 EP10 (Build #13981272) enables Persistent memory in App Direct Mode. This means vSphere 6.7 EP 10 and above supports Memory Mode, App Direct Mode and Mixed Mode.

Table 1 shows the vSphere version and Intel Persistent memory modes support matrix.

Table 1 vSphere version and Persistent memory modes support matrix

Persistent memory mode	Supported vSphere version
Memory Mode	(1) vSphere 6.5 U3 or later (2) vSphere 6.7 U2 + Express Patch 10 (ESXi670-201906002) or later Note: Earlier 6.7 releases are not supported.
App Direct Mode	vSphere 6.7 U2 + Express Patch 10 (ESXi670-201906002) or later
Mixed Mode	vSphere 6.7 U2 + Express Patch 10 (ESXi670-201906002) or later

You can also verify the Persistent memory compatibility of ThinkSystem servers with vSphere using the VMware Compatibility Guide, available at the following address:

<https://www.vmware.com/resources/compatibility/search.php>

Filter the servers by selecting the Persistent Memory feature as shown in Figure 5.

The screenshot shows a web-based server selection tool. The 'Partner Name' dropdown is set to 'Lenovo'. The 'Features' dropdown is open, and 'Persistent Memory' is selected. Other filters include 'Max Cores per Socket', 'CPU Series', and 'Posted Date Range'. Below the filters are buttons for 'Update and View Results' and 'Reset'. A link for 'Read Important Support Information' is provided. The main heading is 'Server Device and Model Information'. A note states that detailed lists show actual vendor devices. A link to view more details and subscribe to RSS feeds is present. The search results table shows 10 results for 'Systems / Servers'.

Partner Name	Model	CPU Series	Supported Releases
Lenovo	ThinkSystem SD530	Intel Xeon Gold 6200/5200 (Cascade-Lake-SP) Series	ESXi 6.7 U3 6.7 U2
Lenovo	ThinkSystem SD530	Intel Xeon Platinum 8200 (Cascade-Lake-SP) Series	ESXi 6.7 U3 6.7 U2
Lenovo	ThinkSystem SR590	Intel Xeon Gold 6200/5200 (Cascade-Lake-SP) Series	ESXi 6.7 U3 6.7 U2
Lenovo	ThinkSystem SR590	Intel Xeon Platinum 8200 (Cascade-Lake-SP) Series	ESXi 6.7 U3 6.7 U2
Lenovo	ThinkSystem SR630	Intel Xeon Gold 6200/5200 (Cascade-Lake-SP) Series	ESXi 6.7 U3 6.7 U2

Figure 5 Persistent memory compatible check on VCG

PMem mode support

There are two different series of Persistent memory: 100 Series and 200 Series. They vary in their support for PMem mode, as shown in Table 2.

Table 2 PMem mode support for 100 Series and 200 Series

DCPMM series	Memory Mode	App Direct Mode	Mixed Mode
Intel Optane PMem 100 Series	Support	Support	Support
Intel Optane PMem 200 Series	Support	Support	No support

Using PMem modules in Memory Mode

When PMem modules are configured in Memory Mode, 100% of Persistent memory capacity acts as system memory. All the PMem modules are seen by the ESXi as second level of volatile memory (2LM). The DRAM DIMMs are hidden from the ESXi and are used as a

high-speed cache for the PMem modules. The total displayed volatile system memory in memory mode is the sum of Persistent memory capacity. ESXi and all the virtual machines can consume the Persistent memory as 2LM.

The following are steps for configuring and using PMem modules in Memory Mode in vSphere 6.7 EP 10 (Build #13981272) on a Lenovo ThinkSystem SR630 server.

1. Refer to “Creating a goal to configure Persistent memory operation modes” on page 7 to configure the PMem modules in Memory Mode.
2. Install vSphere 6.7 EP10 or later on the server
3. After booting to ESXi host, and create virtual machines. Both ESXi and virtual machines can consume the Intel Persistent memory as 2LM.
4. Login to ESXi host web client and click **Host** → **Hardware** to check the 2LM information, Figure 6.

Reserved memory: After configuring Intel Persistent memory in memory mode, it shows a reduction in system memory. It is used for reserving space for metadata. For example, our lab system configuration with 2x 256 GB Intel PMem modules, configured in memory mode provides a system memory of 503.67 GB instead of 512 GB

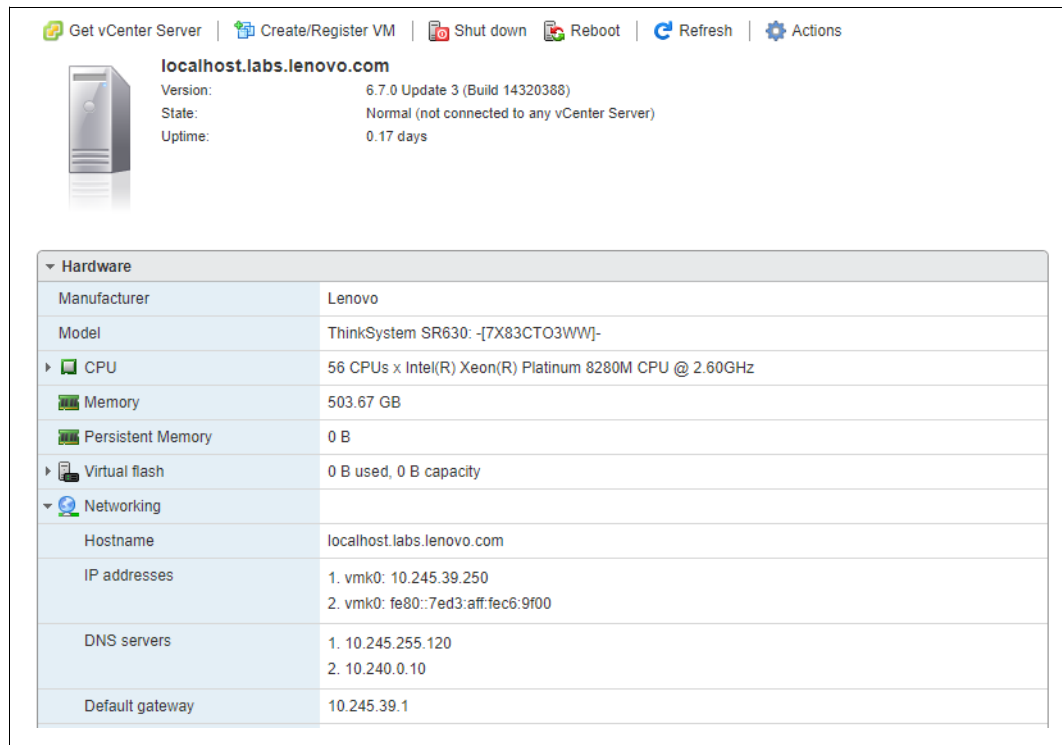


Figure 6 2LM on host web client

5. Run the following ESXCLI command to check physical memory.

```
~# esxcli hardware memory get
```

Figure 7 shows the output of the memory check command on our lab server with 2x 256 GB PMem modules and 2x 32GB DRAM DIMMs, the physical memory size approximately equivalent to sum of Persistent memory capacity. The DRAM DIMMs are hidden from the ESXi and are used as a high-speed cache for the PMem modules.

```
[root@localhost:~] esxcli hardware memory get
Physical Memory: 540808691712 Bytes
Reliable Memory: 0 Bytes
NUMA Node Count: 2
[root@localhost:~] █
```

Figure 7 Output of memory check in ESXi 6.7 U3

Using PMem modules in App Direct Mode

When PMem modules are configured in App Direct Mode, 0% of Persistent memory capacity acts as system memory, all the DRAM DIMMs act as system memory. PMem modules act as independent and persistent memory resources directly accessible by applications.

To use Intel Persistent memory on ESXi host, you must be familiar with the concept of the PMEM datastore and how VMs can access PMem modules:

- ▶ PMEM datastore

In App Direct Mode, ESXi detects all the Intel PMem modules, formats and mounts them as a local PMEM datastore. Only one local PMEM datastore per host is supported. The PMEM datastore is used to store virtual NVDIMM (non-volatile DIMM) devices and traditional virtual disks of a virtual machine. The virtual machine home directory with the vmx and vmware.log files cannot be placed on the PMEM datastore.

The PMEM datastore workflow is as follows:

- a. During boot, ESXi automatically creates namespace
- b. Create VMware PMEM partition on each namespace
- c. Concatenates partitions together into single logically contiguous space
- d. Formats and mounts it as a local PMEM datastore

Tip: In general, there's no need to create namespaces or partitions for ESXi to consume the PMem modules. Even though ESXi provides commands to create namespaces and partitions, it's recommended to leave that responsibility to ESXi which will automatically create namespaces on each interleaved set during boot.

If the namespaces were already created by other means like UEFI or Intel ipmctl, ESXi discovers those namespaces during boot and attempts to create PMEM volume with the free space available in the namespaces.

- ▶ PMem modules Access Modes for virtual machines

When PMem modules are configured in App Direct Mode, ESXi exposes PMem modules to a virtual machine in the following two different modes. PMEM-aware virtual machines can have direct access to PMem modules. Traditional virtual machines can use fast virtual disks stored on the PMEM datastore.

- Direct Access Mode: vSphere presents PMem modules to PMEM-aware virtual machines as a virtual NVDIMM device. The virtual machines can use the virtual NVDIMM module as a standard byte-addressable memory that can persist across power cycles.

- Virtual Disk Mode: vSphere presents PMem modules to traditional virtual machines just as if it were a virtual SCSI device, There's no need to change anything for the guest OS or applications. In this way, Virtual Persistent Memory Disk (vPMEMDisk) allows using PMEM in older operating systems and applications.

The following are steps for configuring and using PMem modules in App Direct Mode in vSphere 6.7 EP 10 (Build #13981272) on a ThinkSystem SR630 server:

1. Refer to "Creating a goal to configure Persistent memory operation modes" on page 7 to configure the PMem modules in App Direct mode.
2. Install vSphere 6.7 EP10 or later on the server
3. After booting to ESXi host, ESXi detects the PMem modules and exposes them as a local PMEM datastore to the virtual machines that run on the host.
4. Login to the ESXi host web client and click **Host** → **Storage** → **Datastores** to check PMEM datastore, as shown in Figure 8.

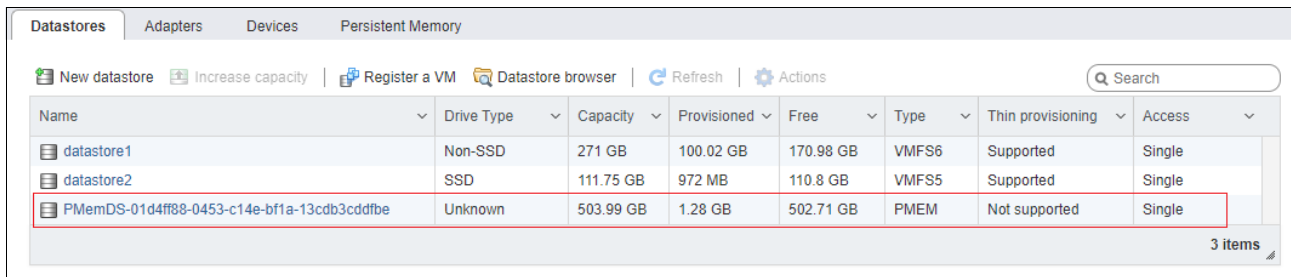


Figure 8 PMEM datastore check on host web client

5. Run the following ESXCLI command to check PMEM datastore, as shown in Figure 9 on page 11.

```
~# esxcli storage filesystem list
```

The output of the command is shown in Figure 9 on page 12.




Figure 9 PMEM datastore check in ESXCLI

6. In the ESXi host web client, click **Host** → **Storage** → **Persistent Memory** to check PMem modules namespaces, as shown in Figure 10 on page 12.

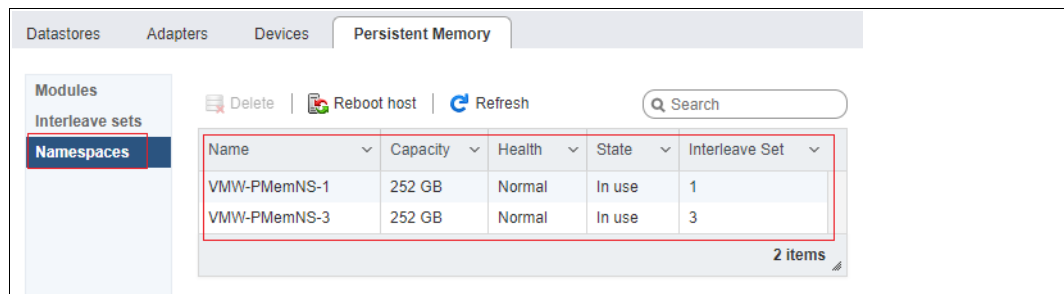


Figure 10 Namespace check in host web client

7. Run the following localcli command to check PMem modules namespace, as shown in Figure 11.

```
~# localcli --plugin-dir /usr/lib/vmware/esxcli/int hardwareinternal nvd namespace details list
```

```
[root@localhost:~] localcli --plugin-dir /usr/lib/vmware/esxcli/int hardwareinternal nvd namespace details list

Namespace VMW-PMemNS-1:
  UUID of Namespace: 5bcae9ef-344b-58db-a0fc-abe2cdfdedaa
  Name of Namespace: VMW-PMemNS-1
  Namespace Size: 270582939648
  Namespace State: InUse
  Health summary: OK
  Interleave set ID: 0x1

Namespace VMW-PMemNS-3:
  UUID of Namespace: 135e8db4-2541-560b-b16d-d2aa9b14418c
  Name of Namespace: VMW-PMemNS-3
  Namespace Size: 270582939648
  Namespace State: InUse
  Health summary: OK
  Interleave set ID: 0x3
```

Figure 11 Namespace check in ESXi localcli

8. Go to vSphere web client, create a virtual machine installed a PMEM-aware guest OS (e.g. RHEL7.6) to consume PMem modules as a virtual NVDIMM device in Direct Access Mode, make sure that the virtual hardware version is 14 or higher.

Tip: In Direct Access Mode, the PMem modules are exposed to a virtual machine as virtual NVDIMMs. It enables the virtual machine to use PMem modules in byte-addressable random mode. The virtual machine must have a PMEM-aware guest OS. The PMem modules are compatible with latest operating systems that support persistent memory, for example, Windows Server 2016 and RHEL 7.6. Each virtual machine can have a maximum of one virtual NVDIMM controller and each NVDIMM controller can have up to 64 virtual NVDIMM devices.

- a. Power off the virtual machine and then right-click the virtual machine in the inventory and select **Edit Settings** to add a new NVDIMM device.
- b. On the Virtual Hardware tab, click **Add other device** and select **NVDIMM** from the drop-down menu, as shown in Figure 12 on page 14.




Figure 12 Add virtual NVDIMM on host web client

- c. In the New NVDIMM text box, enter the size of the NVDIMM device, click **Save** button, as shown in Figure 13.

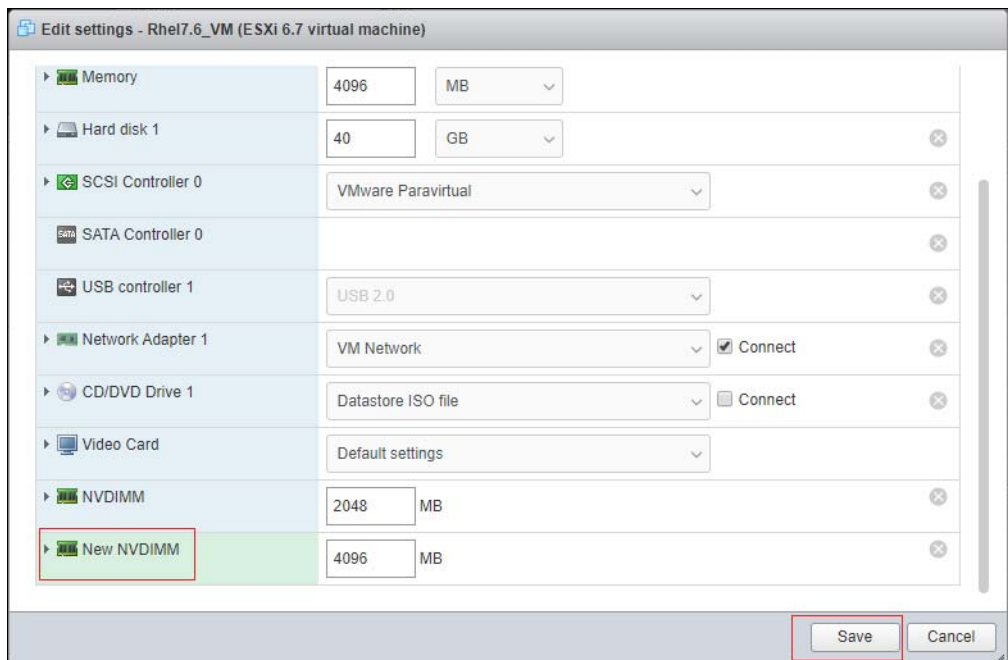


Figure 13 Configure virtual NVDIMM size on host web client

- d. Power on the virtual machine, use the following commands to check this new virtual NVDIMM device (`/dev/pmem0`) and format/mount it as file system (e.g. `ext4`) which support direct-access (DAX).

For the Linux OS:

The following command is to make sure the mode of the current namespace:

```
~# sudo ndctl list -Nu
```

The following command is to change the mode of the namespace to fsdax mode, if the namespace's mode is not fsdax:

```
~# sudo ndctl create-namespace -f -e namespace0.0 --mode=fsdax
```

The following command is to check virtual NVDIMM:

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

The following command is to format the virtual NVDIMM device:

```
~# sudo mkfs.ext4 /dev/pmem0
```

The following command is to mount the virtual NVDIMM device as DAX supported file system:

```
~# sudo mount -o dax /dev/pmem0 /opt/pmem
```

The following command is to check the DAX supported filesystem:

```
~# df -lh
```

Figure 14 shows the output of above commands on RHEL 9.0.

```
[root@localhost ~]# sudo ndctl list -Nu
{
  "dev": "namespace0.0",
  "mode": "raw",
  "size": "1024.00 MiB (1073.74 MB)",
  "sector_size": 512,
  "blockdev": "pmem0"
}
[root@localhost ~]# sudo ndctl create-namespace -f -e namespace0.0 --mode=fsdax
{
  "dev": "namespace0.0",
  "mode": "fsdax",
  "map": "dev",
  "size": "1006.00 MiB (1054.87 MB)",
  "uuid": "42809c7c-2296-44fb-b659-add52fbd6fdc",
  "sector_size": 512,
  "align": 2097152,
  "blockdev": "pmem0"
}
[root@localhost ~]# ls /dev/pmem*
/dev/pmem0
[root@localhost ~]# sudo mkfs.ext4 /dev/pmem0
mke2fs 1.46.5 (30-Dec-2021)
Creating filesystem with 257536 4k blocks and 64384 inodes
Filesystem UUID: 329f8126-73db-4eb2-bf9e-2c9f78f5daf8
Superblock backups stored on blocks:
    32768, 98304, 163840, 229376

Allocating group tables: done
Writing inode tables: done
Creating journal (4096 blocks): done
Writing superblocks and filesystem accounting information: done

[root@localhost ~]# sudo mount -o dax /dev/pmem0 /opt/pmem
[root@localhost ~]# df -lh
Filesystem                Size      Used Avail Use% Mounted on
devtmpfs                  1.8G         0  1.8G   0% /dev
tmpfs                      1.9G         0  1.9G   0% /dev/shm
tmpfs                      745M    9.4M  736M   2% /run
/dev/mapper/rhel_xcc--40fee95e805b-root 13G    6.8G  6.1G  54% /
/dev/sda2                  1014M    260M  755M  26% /boot
/dev/sda1                   599M    7.0M  592M   2% /boot/efi
tmpfs                      373M    96K  373M   1% /run/user/0
/dev/pmem0                 972M    24K  906M   1% /opt/pmem
[root@localhost ~]#
```

Figure 14 Command output on RHEL 9.0 virtual machine

For the Windows OS:

The following screenshot is to check the status of DCPMM via Device Manager in Figure 15.

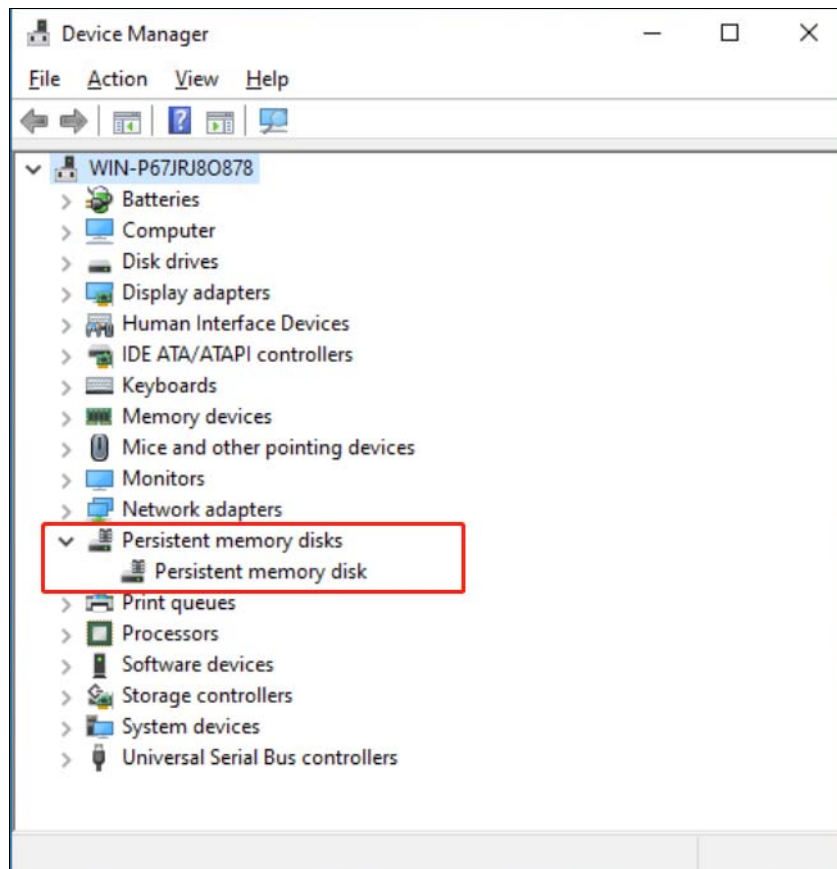


Figure 15 Checking the status of DCPMM via Device Manager

The following command is to verify the driver enumerates the PMEM Disk and get the DiskNumber for further initialization work.

```
~# Get-PmemDisk
```

The following command is to partition as GPT.

```
~# get-disk -Number 2 | Initialize-Disk -PartitionStyle GPT
```

The following command is to mount as DAX NTFS File systems.

```
~# get-disk -Number 2 | New-Volume -FriendlyName DAX-VOL -DriveLetter F |  
Format-volume -Filesystem NTFS -IsDAX $true
```

The following command is to check if the above DCPMM disk is DAX mode.

```
~# fsutil.exe fsinfo volumeinfo f:
```

Figure 16 on page 17 shows the output of above commands on Windows Server 2022.


```

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

Install the latest PowerShell for new features and improvements! https://aka.ms/PSWindows

PS C:\Users\Administrator> Get-PmemDisk

DiskNumber Size HealthStatus AtomicityType CanBeRemoved PhysicalDeviceIds UnsafeShutdownCount
-----
1          1 GB Healthy      None          False         {0}           0

PS C:\Users\Administrator> Get-Disk -Number 1 | Initialize-Disk -PartitionStyle GPT
PS C:\Users\Administrator> Get-Disk -Number 1 | New-Volume -FriendlyName DAX-VOL -DriveLetter F | Format-Volume -FileSystem NTFS -IsDAX $true

DriveLetter FriendlyName FileSystemType DriveType HealthStatus OperationalStatus SizeRemaining      Size
-----
F            DAX-VOL      NTFS          Fixed      Healthy      OK              1000.42 MB 1005.94 MB

PS C:\Users\Administrator> fsutil.exe fsinfo volumeinfo f:
Volume Name :
Volume Serial Number : 0x2e08f560
Max Component Length : 255
File System Name : NTFS
Is ReadWrite
Not Thinly-Provisioned
Supports Case-sensitive filenames
Preserves Case of filenames
Supports Unicode in filenames
Preserves & Enforces ACL's
Supports Disk Quotas
Supports Reparse Points
Returns Handle Close Result Information
Supports POSIX-style Unlink and Rename
Supports Object Identifiers
Supports Named Streams
Supports Hard Links
Supports Extended Attributes
Supports Open By FileID
Supports USN Journal
Is DAX volume
BTT (Block Translation Table) enabled
PS C:\Users\Administrator>

```

Figure 16 Command output on Windows Server 2022 virtual machine

9. Go to the vSphere web client and configure the guest OS virtual machine to consume PMem modules as a Virtual Persistent Memory Disk in Virtual Disk Mode.

Tip: In Virtual Disk Mode, the PMem modules are accessed by the virtual machine as a virtual SCSI device and the virtual disk is stored in a PMEM datastore. This mode is available to any traditional virtual machine and supports any hardware version, including all legacy versions. Virtual machines are not required to be PMEM-aware.

- a. Power off the virtual machine and then right-click the virtual machine in the inventory and select **Edit Settings** to add a new persistent memory disk device.
- b. On the Virtual Hardware tab, click **Add hard disk** and select **New persistent memory disk** from the drop-down menu, as shown in Figure 17.




Figure 17 Add new virtual persistent memory disk on host web client

- c. In the New hard disk text box, enter the size of the new virtual persistent memory disk, click Save button, as shown in Figure 18.

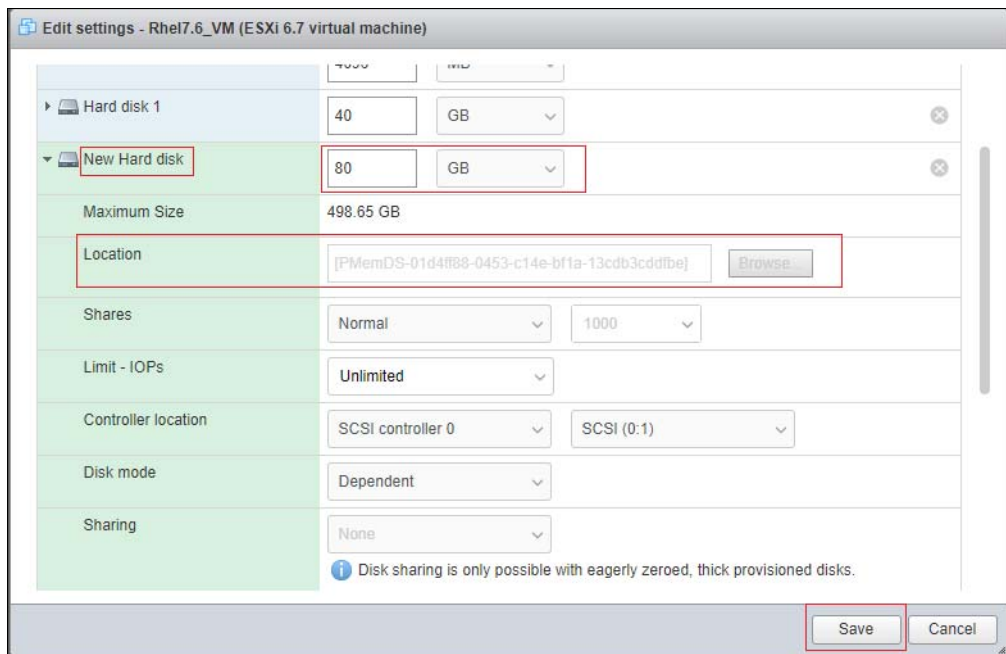


Figure 18 Configure virtual persistent memory disk size on host web client

- d. Power on the virtual machine, use the following commands to check the new added virtual persistent memory disk, format/mount it as file system.

For the Linux OS:

The following command is to check virtual persistent memory disk.

```
~# ls /dev/sdb*
```

The following command is to format the virtual persistent memory disk device:

```
~# sudo mkfs.ext4 /dev/sdb
```

The following command is to mount the virtual persistent memory disk as file system.

```
~# sudo mount /dev/sdb /root/pmem_test
```

The following command is to check the filesystem:

```
~# df -lh
```

Figure 19 shows the output of above commands.

```
[root@localhost ~]# ls /dev/sdb*
/dev/sdb
[root@localhost ~]# sudo mkfs.ext4 /dev/sdb
mke2fs 1.42.9 (28-Dec-2013)
/dev/sdb is entire device, not just one partition!
Proceed anyway? (y,n) y
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
Stride=0 blocks, Stripe width=0 blocks
5242880 inodes, 20971520 blocks
1048576 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=2168455168
640 block groups
32768 blocks per group, 32768 fragments per group
8192 inodes per group
Superblock backups stored on blocks:
    32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632, 2654208,
    4096000, 7962624, 11239424, 20480000

Allocating group tables: done
Writing inode tables: done
Creating journal (32768 blocks): done
Writing superblocks and filesystem accounting information: done

[root@localhost ~]# mount /dev/sdb /root/pmem_test
[root@localhost ~]# df -lh
Filesystem      Size  Used Avail Use% Mounted on
/dev/mapper/rhel-root  36G  1.1G   35G   3% /
devtmpfs        1.9G     0  1.9G   0% /dev
tmpfs           1.9G     0  1.9G   0% /dev/shm
tmpfs           1.9G  8.9M  1.9G   1% /run
tmpfs           1.9G     0  1.9G   0% /sys/fs/cgroup
/dev/sdal       1014M  145M  870M  15% /boot
tmpfs           379M     0  379M   0% /run/user/0
/dev/sdb        79G   57M   75G   1% /root/pmem_test
```

Figure 19 Command output on RHEL 7.6 virtual machine

For the Windows OS:

Open disk manager console “diskmgmt.msc”, initialize the disks and format them.

If the disk is offline, right click on the disk and select “online” to bring this disk online.

Check the DCPMM disk via Device Manager and Disk Manager.

Figure 20 shows the output of above steps.

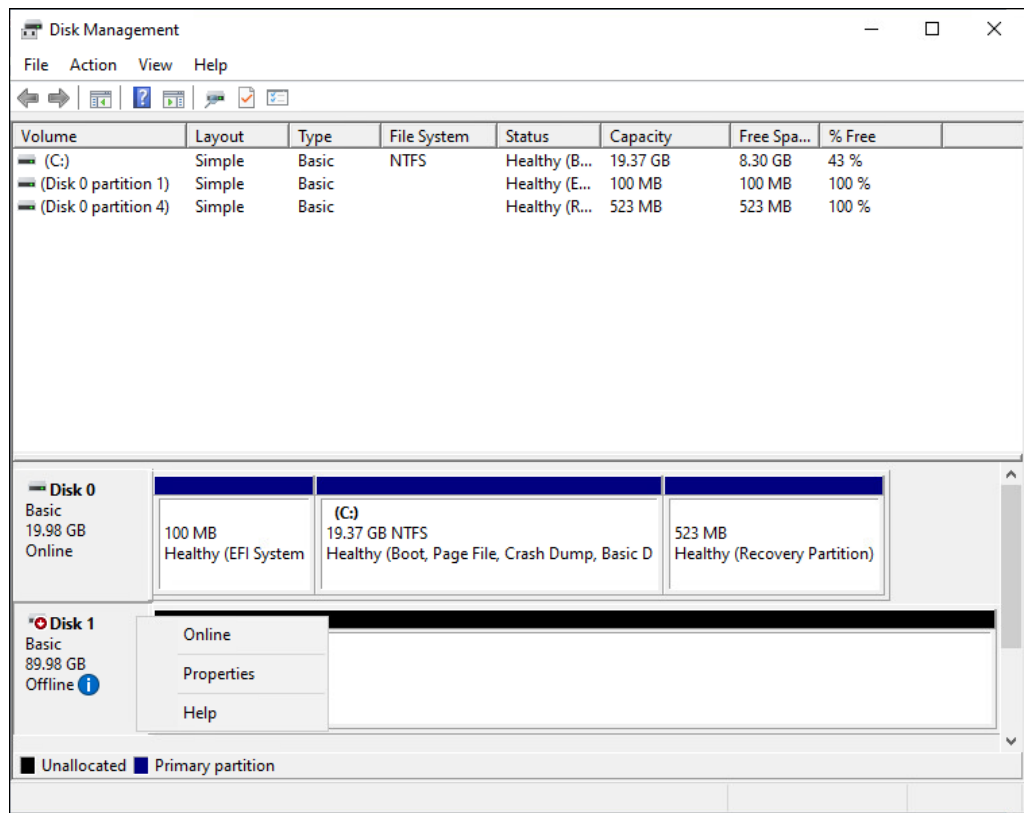


Figure 20 Disk Manager on Windows Server 2022

Using PMem modules in Mixed Mode

Note: Mixed Mode is only supported with Persistent Memory 100 Series.

When PMem modules are configured in Mixed Mode, 1-99% of Persistent memory capacity acts as system memory. In this mode, some percentage of Persistent memory capacity is directly accessible to ESXi and virtual machines as persistent memory in App Direct Mode operations, while the rest serves as system memory in Memory Mode operations. DRAM DIMMs are hidden from the ESXi and act as cache. The total displayed volatile system memory in this mode is the Persistent memory capacity that is assigned as volatile system memory.

The following are steps for configuring and using PMem modules in Mixed Mode in vSphere 6.7 EP 10 (Build #13981272) on a ThinkSystem SR630 server:

1. Refer to “Creating a goal to configure Persistent memory operation modes” on page 7 to configure the PMem modules in Mixed Mode.

In our lab setup, we installed our SR630 server with 2x 256 GB PMem modules and 2x 32 GB DRAM DIMMs and configure the PMem modules with 50% Memory Mode and 50% App Direct Mode.

Tip: The suggested ratios for Mixed Mode are 25%, 50% and 75%.

2. Install vSphere 6.7 EP10 or later on the server
3. Boot to the ESXi host and create virtual machines. Both ESXi and virtual machines can consume about 50% of the Persistent memory as 2LM, and can consume 50% of Persistent memory as persistent memory, all the DRAM DIMMs act as cache.
4. Login to ESXi host web client, and click **Host** → **Hardware** to check the 2LM and persistent memory as shown in Figure 21 on page 21.

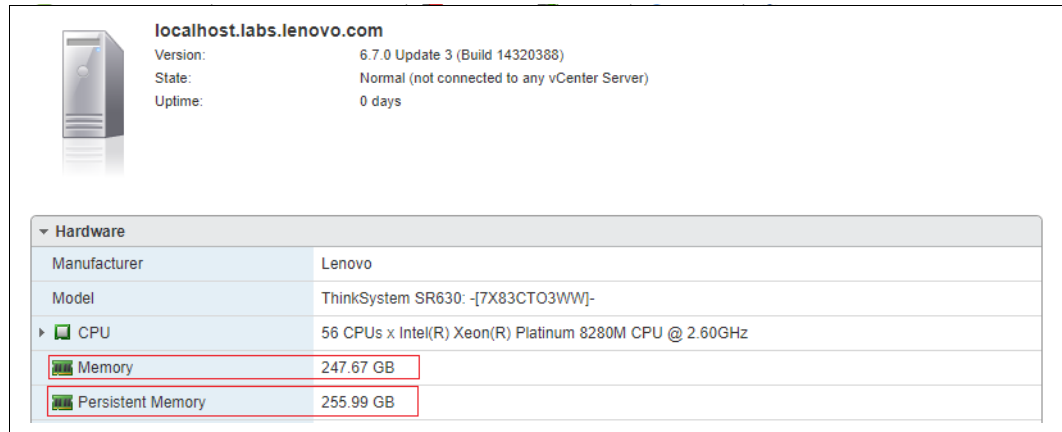


Figure 21 Check 2LM and PMEM on host web client

5. Refer to the “Using PMem modules in Memory Mode” on page 9 and “Using PMem modules in App Direct Mode” on page 11 to configure and use PMem modules in Mixed mode.

References

Review the following web pages for more information:

- ▶ Intel Optane Persistent Memory 100 Series Product Guide
<https://lenovopress.lenovo.com/lp1066-intel-optane-persistent-memory-100-series>
- ▶ Intel Optane Persistent Memory 200 Series Product Guide
<https://lenovopress.lenovo.com/lp1380-intel-optane-persistent-memory-200-series>
- ▶ Intel Optane DC Persistent Memory: A Major Advance in Memory and Storage Architecture
<https://software.intel.com/en-us/articles/intel-optane-dc-persistent-memory-a-major-advance-in-memory-and-storage-architecture>
- ▶ Intel Optane Persistent Memory
<https://www.intel.com/content/www/us/en/products/docs/memory-storage/optane-persistent-memory/overview.html>
- ▶ VMware vSphere Documentation
<https://docs.vmware.com/en/VMware-vSphere/6.7/com.vmware.vsphere.resmgmt.doc/GUID-EB72D358-9C2C-4FBD-81A9-A145E155CE31.html>

Authors

Chengcheng Peng is a VMware Engineer at the Lenovo Infrastructure Solutions Group in Beijing, China. She joined the OS team in 2018. Her main interests are vSphere security and storage. She has 5 years' experience as a VMware Engineer.

Skyler Zhang is a VMware Engineer at the Lenovo Infrastructure Solutions Group, based in Beijing, China. She joined the OS team in 2022.

Thanks to the following specialists for their contributions and suggestions:

- ▶ Boyong Li, Lenovo OS Technical Leader
- ▶ Yuepei, Lenovo OS PM
- ▶ David Watts, Lenovo Press

Change history

April 3, 2023:

- ▶ Added details about implementing Intel Optane Persistent Memory 200 Series
- ▶ Updated ESXi information to ESXi 7.0 U3
- ▶ Optimized the steps of checking for DAX support on Linux.
- ▶ Added steps on how to check PMem on Windows.

Notices

Lenovo may not offer the products, services, or features discussed in this document in all countries. Consult your local Lenovo representative for information on the products and services currently available in your area. Any reference to a Lenovo product, program, or service is not intended to state or imply that only that Lenovo product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any Lenovo intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any other product, program, or service.

Lenovo may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

Lenovo (United States), Inc.
1009 Think Place - Building One
Morrisville, NC 27560
U.S.A.
Attention: Lenovo Director of Licensing

LENOVO PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. Lenovo may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

The products described in this document are not intended for use in implantation or other life support applications where malfunction may result in injury or death to persons. The information contained in this document does not affect or change Lenovo product specifications or warranties. Nothing in this document shall operate as an express or implied license or indemnity under the intellectual property rights of Lenovo or third parties. All information contained in this document was obtained in specific environments and is presented as an illustration. The result obtained in other operating environments may vary.

Lenovo may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any references in this publication to non-Lenovo Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this Lenovo product, and use of those Web sites is at your own risk.

Any performance data contained herein was determined in a controlled environment. Therefore, the result obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

This document was created or updated on April 3, 2023.

Send us your comments via the **Rate & Provide Feedback** form found at <http://lenovopress.com/lp1225>

Trademarks

Lenovo and the Lenovo logo are trademarks or registered trademarks of Lenovo in the United States, other countries, or both. These and other Lenovo trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by Lenovo at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of Lenovo trademarks is available from <https://www.lenovo.com/us/en/legal/copytrade/>.

The following terms are trademarks of Lenovo in the United States, other countries, or both:

Lenovo®

Lenovo(logo)®

ThinkSystem™

The following terms are trademarks of other companies:

Intel, Intel Optane, Xeon, and the Intel logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Windows, Windows Server, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.