

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

# Implementing Intel Optane DC Persistent Memory with VMware vSphere

---

**Describes the hardware prerequisites for DCPMMs**

---

**Explains how to configure DCPMM operation modes**

---

**Shows vSphere support of DCPMM operation modes**

---

**Provide instructions on how to configure and use DCPMMs in vSphere**

**Chengcheng Peng**



# Abstract

Intel Optane DC Persistent Memory represents a ground-breaking 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 DC Persistent Memory Modules (DCPMMs) 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 DCPMM with different operation modes in VMware vSphere 6.7 EP 10 (Build #13981272) and above 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
DCPMM configuration .....	4
Firmware update .....	4
Check DCPMM status .....	4
Creating a goal to configure DCPMM operation modes .....	6
vSphere support of DCPMM modes .....	7
Using DCPMMs in Memory Mode .....	9
Using DCPMMs in App Direct Mode .....	10
Using DCPMMs in Mixed Mode .....	17
References .....	17
Authors .....	18
Notices .....	19
Trademarks .....	20

# Introduction

Intel Optane DC 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.

DC Persistent Memory Modules (DCPMMs) 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 DCPMM, 3 TB per CPU.

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

DCPMM can be configured in three modes:

- ▶ **Memory Mode:** The DCPMMs act as large capacity DDR4 memory modules.
- ▶ **App Direct Mode:** The DCPMMs provide all persistence features to the operating system and applications that support them.
- ▶ **Mixed Mode:** Mixed Mode is a combination of Memory Mode and App Direct Mode, where a portion of the capacity of the DCPMMs is used for the Memory Mode operations, and the remaining capacity of the DCPMMs 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 DCPMMs in Memory Mode.

More information about DCPMM can be found in the Lenovo Press product guide:

<https://lenovopress.com/lp1066-intel-optane-dc-persistent-memory>

## Hardware setup

There are many hardware requirements when setting up DCPMM experimental. This section presents the hardware prerequisites needed to use DCPMMs in a ThinkSystem server.

### ▶ Server selection

DCPMMs are only supported in the second-generation Intel Xeon Scalable processors. DCPMMs are not supported in the first-generation Xeon Scalable processors. Refer to the DCPMM product guide to select a server to support DCPMMs:

<https://lenovopress.com/lp1066-optane-dc-persistent-memory#server-support>

### ▶ Processor selection

All Platinum processors, all Gold processors, and the Silver 4215 processor support DCPMMs. Refer to the following DCPMM Product Guide to select correct processors to support DCPMMs:

<https://lenovopress.com/lp1066-optane-dc-persistent-memory#processor-support>

### ▶ DCPMMs and DRAM DIMMs selection

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

- All installed DCPMMs must be the same size. Mixing DCPMMs 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 following DCPMM Product Guide to get more details about Memory DIMM support:

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

- ▶ DCPMMs and DRAM DIMMs installation

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

[https://thinksystem.lenovofiles.com/help/index.jsp?topic=%2F7X21%2Fmemory\\_module\\_installation\\_order\\_aep.html](https://thinksystem.lenovofiles.com/help/index.jsp?topic=%2F7X21%2Fmemory_module_installation_order_aep.html)

## DCPMM configuration

Before configure DCPMM, make sure that the capacity of installed DCPMMs and DRAM DIMMs meets system requirements for the following different operation modes:

- ▶ Memory mode requirements: Refer to following DCPMM Product Guide to get details about memory mode requirements:

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

- ▶ App Direct Mode requirements: Refer to following DCPMM Product Guide to get details about App Direct Mode requirements:

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

- ▶ Mixed Mode requirements: Refer to following DCPMM Product Guide 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 DCPMMs to the latest version before configuring persistent memory, and that you make sure that DCPMM FW and UEFI FW use the same Intel BKC version.

- ▶ BMC and UEFI Firmware Update

Refer to the following guide to update BMC and UEFI firmware to the latest version:

[https://thinksystem.lenovofiles.com/help/index.jsp?topic=%2F7X21%2Fupdate\\_the\\_firmware.html](https://thinksystem.lenovofiles.com/help/index.jsp?topic=%2F7X21%2Fupdate_the_firmware.html)

- ▶ DCPMMs Firmware Update

Refer to the following guide to update the DCPMM firmware to the latest version:

[https://sysmgmt.lenovofiles.com/help/index.jsp?topic=%2Fcom.lenovo.lxca.doc%2Fupdate\\_fw.html](https://sysmgmt.lenovofiles.com/help/index.jsp?topic=%2Fcom.lenovo.lxca.doc%2Fupdate_fw.html)

## Check DCPMM status

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

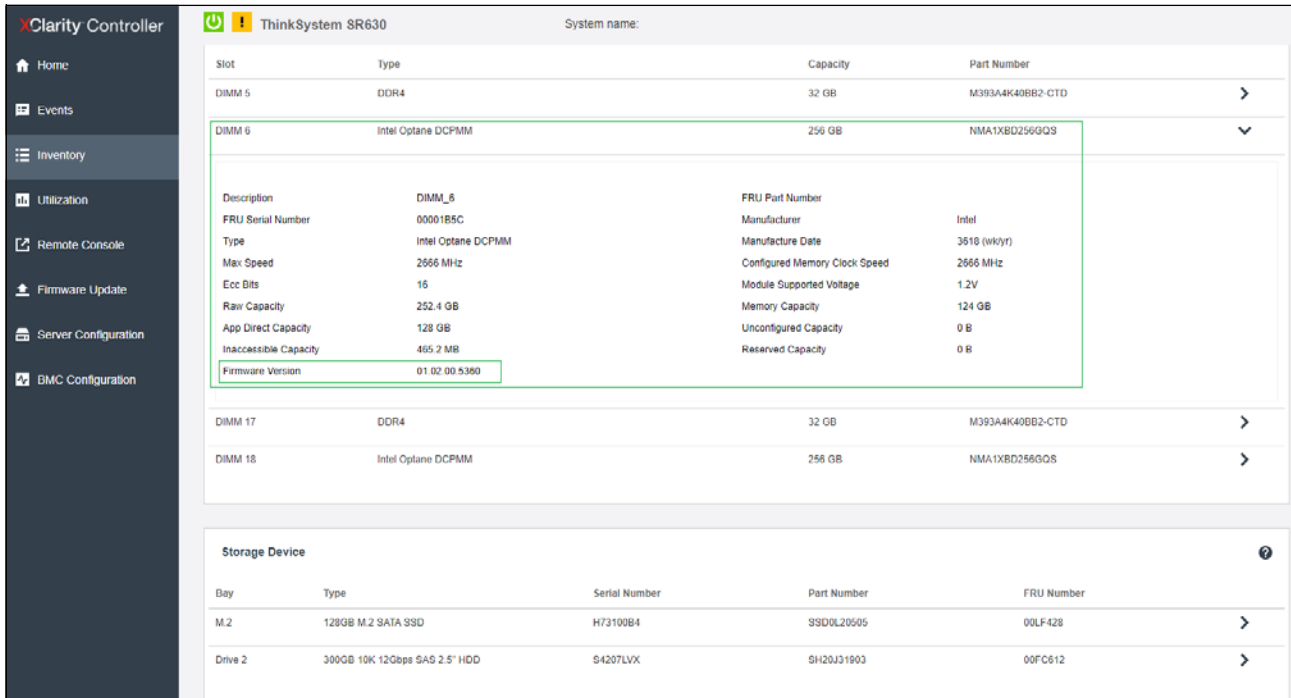


Figure 1 DCPMM status in XCC

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

Power on Lenovo ThinkSystem server and then press F1 to enter System Setup. Select **System Settings** → **Intel Optane DCPMMs** where you can check DCPMMs details as shown in Figure 2 on page 6.

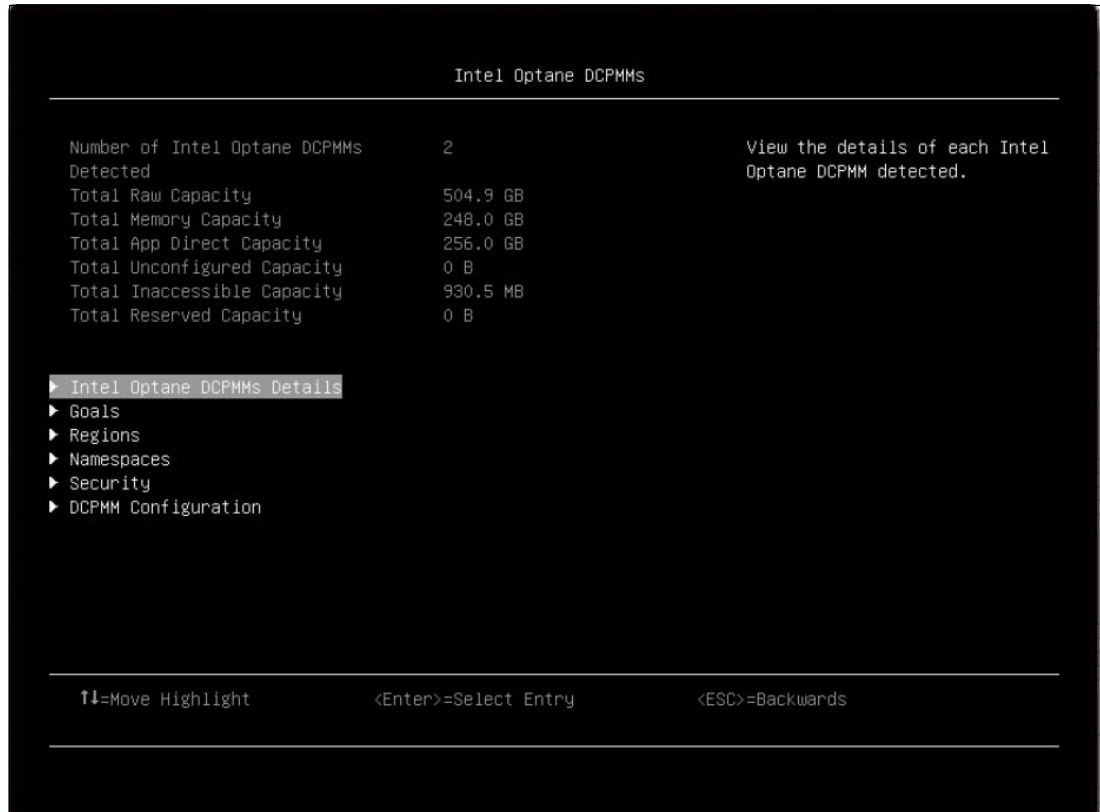


Figure 2 DCPMM details in UEFI setting

## Creating a goal to configure DCPMM operation modes

A goal is a mechanism to allocate DCPMMs capacity for Memory mode and App-direct mode. User must choose platform or processor for the unit and choose a percent to allocate DCPMMs capacity in Memory mode. The remaining DCPMMs 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 DCPMMs** → **Goals**, as shown in Figure 3 on page 7.



Figure 3 Create Goal in UEFI setting

3. In the “Memory Mode [%]” field, specify the percentage of DCPMM capacity that is to be allocated to system memory, and hence decide the DCPMM 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 DCPMM configuration on Lenovo ThinkSystem servers, you can refer to the following InfoCenter page:

[https://thinksystem.lenovofiles.com/help/index.jsp?topic=%2F7X21%2Fmemory\\_configuration\\_dcpmm.html&cp=2\\_0\\_7\\_4\\_0](https://thinksystem.lenovofiles.com/help/index.jsp?topic=%2F7X21%2Fmemory_configuration_dcpmm.html&cp=2_0_7_4_0)

## vSphere support of DCPMM modes

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

VMware supports Intel DCPMM in both Memory Mode and App Direct Mode. vSphere 6.5 U3 enables Intel DCPMM in Memory Mode. vSphere 6.7 EP10 (Build #13981272) enables DCPMM 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 DCPMM modes support matrix.

Table 1 vSphere version and DCPMM modes support matrix

DCPMM 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 DCPMM 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 4.

Click here to [Read Important Support Information](#)

### Server Device and Model Information

The detailed lists show actual vendor devices that are either physically tested or are similar to the devices tested by VMware or VMware partners. VMware provides support only for the devices that are listed in this document.

Click on the 'Model' to view more details and to subscribe to RSS feeds.

[Bookmark](#) | [Print](#) | [Export to CSV](#)

Search Results: Your search for " Systems / Servers " returned 10 results. [Back to Top](#) [Turn Off Auto Scroll](#) Display: 10

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 4 DCPMM compatible check on VCG



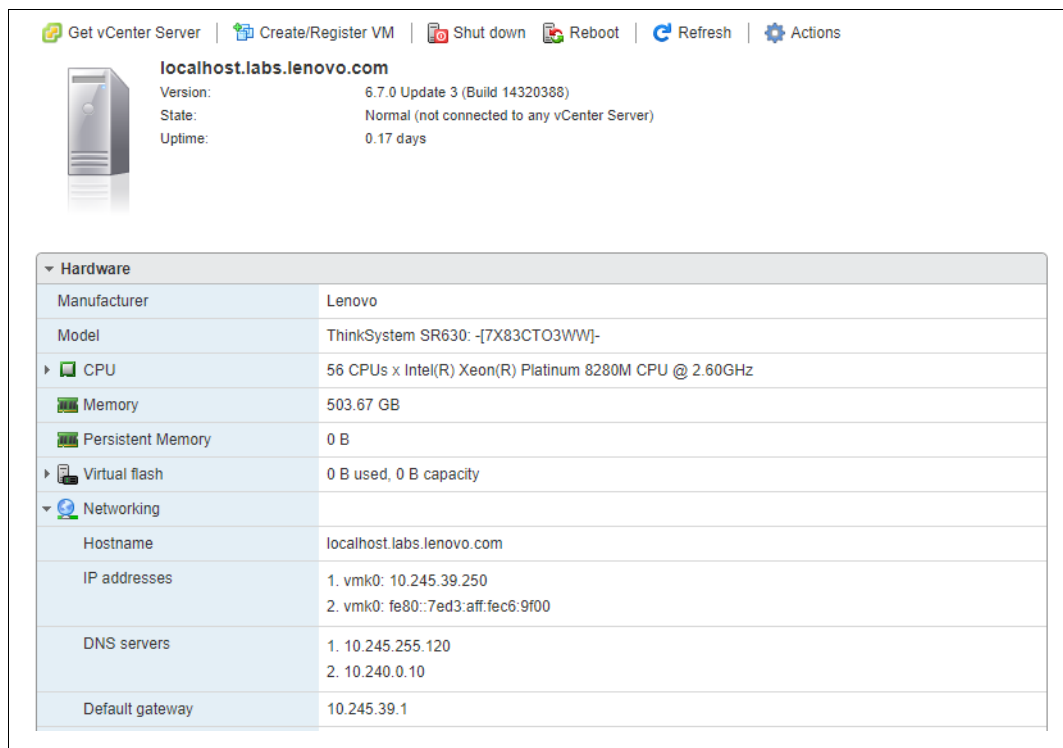
# Using DCPMMs in Memory Mode

When DCPMMs are configured in Memory Mode, 100% of DCPMM capacity acts as system memory. All the DCPMMs 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 DCPMMs. The total displayed volatile system memory in memory mode is the sum of DCPMM capacity. ESXi and all the virtual machines can consume the DCPMM as 2LM.

The following are steps for configuring and using DCPMMs 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 DCPMM operation modes” on page 6 to configure the DCPMMs 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 DCPMM as 2LM.
4. Login to ESXi host web client and click **Host** → **Hardware** to check the 2LM information, Figure 5.

**Reserved memory:** After configuring Intel DCPMM 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 DCPMMs, configured in memory mode provides a system memory of 503.67 GB instead of 512 GB



localhost.labs.lenovo.com	
Version:	6.7.0 Update 3 (Build 14320388)
State:	Normal (not connected to any vCenter Server)
Uptime:	0.17 days

Hardware	
Manufacturer	Lenovo
Model	ThinkSystem SR630- [7X83CTO3WW]-
CPU	56 CPUs x Intel(R) Xeon(R) Platinum 8280M CPU @ 2.60GHz
Memory	503.67 GB
Persistent Memory	0 B
Virtual flash	0 B used, 0 B capacity
Networking	
Hostname	localhost.labs.lenovo.com
IP addresses	1. vmk0: 10.245.39.250 2. vmk0: fe80::7ed3:aff:fec6:9f00
DNS servers	1. 10.245.255.120 2. 10.240.0.10
Default gateway	10.245.39.1

Figure 5 2LM on host web client

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

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

Figure 6 shows the output of the memory check command on our lab server with 2x 256 GB DCPMMs and 2x 32GB DRAM DIMMs, the physical memory size approximately equivalent to sum of DCPMM capacity. The DRAM DIMMs are hidden from the ESXi and are used as a high-speed cache for the DCPMMs.

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

Figure 6 Output of memory check in ESXi 6.7 U3

## Using DCPMMs in App Direct Mode

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

To use Intel DCPMM on ESXi host, you must be familiar with the concept of the PMEM datastore and how VMs can access DCPMMs:

- ▶ PMEM datastore

In App Direct Mode, ESXi detects all the Intel DCPMMs, 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 DCPMMs. 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.

- ▶ DCPMMs Access Modes for virtual machines

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

- Direct Access Mode: vSphere presents DCPMMs 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 DCPMMs 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 DCPMMs 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 DCPMM operation modes” on page 6 to configure the DCPMMs in App Direct mode.
2. Install vSphere 6.7 EP10 or later on the server
3. After booting to ESXi host, ESXi detects the DCPMMs 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 7.

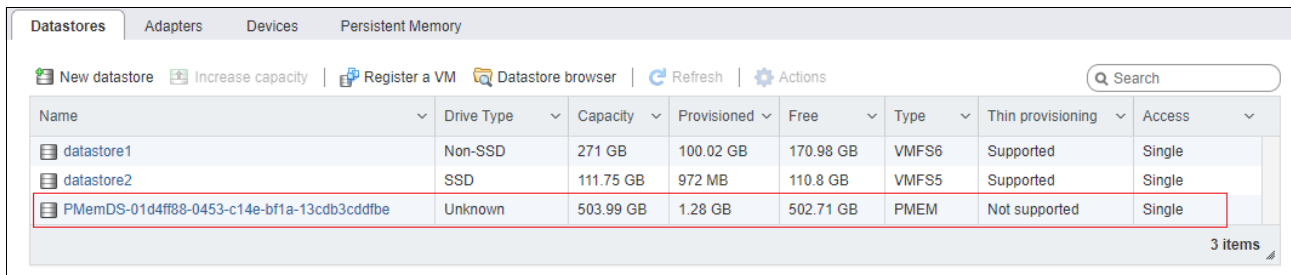


Figure 7 PMEM datastore check on host web client

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 8 on page 11.




Figure 8 PMEM datastore check in ESXCLI

6. In the ESXi host web client, click **Host** → **Storage** → **Persistent Memory** to check DCPMMs namespaces, as shown in Figure 9 on page 11.

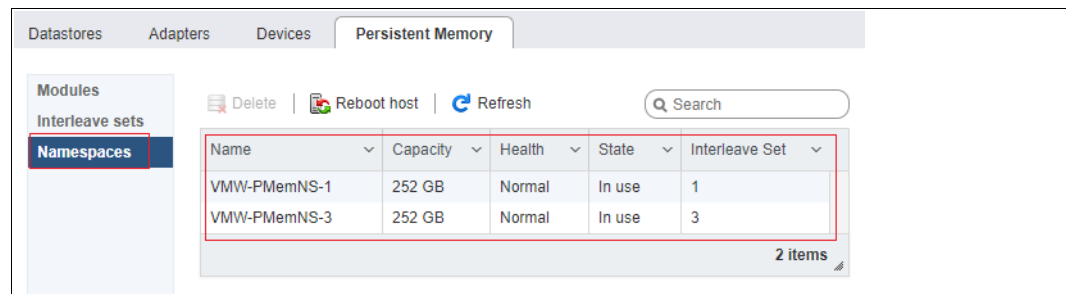


Figure 9 Namespace check in host web client

7. Run the following localcli command to check DCPMMs namespace, as shown in Figure 10.

```
~# 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 10 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 DCPMMs 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 DCPMMs are exposed to a virtual machine as virtual NVDIMMs. It enables the virtual machine to use DCPMMs in byte-addressable random mode. The virtual machine must have a PMEM-aware guest OS. The DCPMMs 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 11 on page 13.

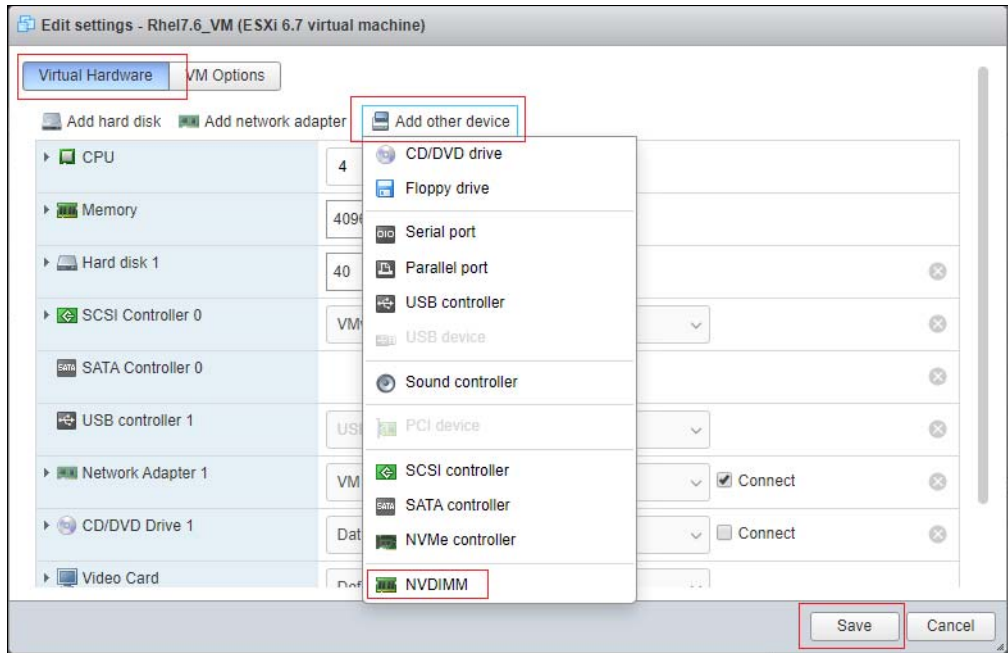


Figure 11 Add virtual NVDIMM on host web client

- c. In the New NVDIMM text box, enter the size of the NVDIMM device, click **Save** button, as shown in Figure 12.

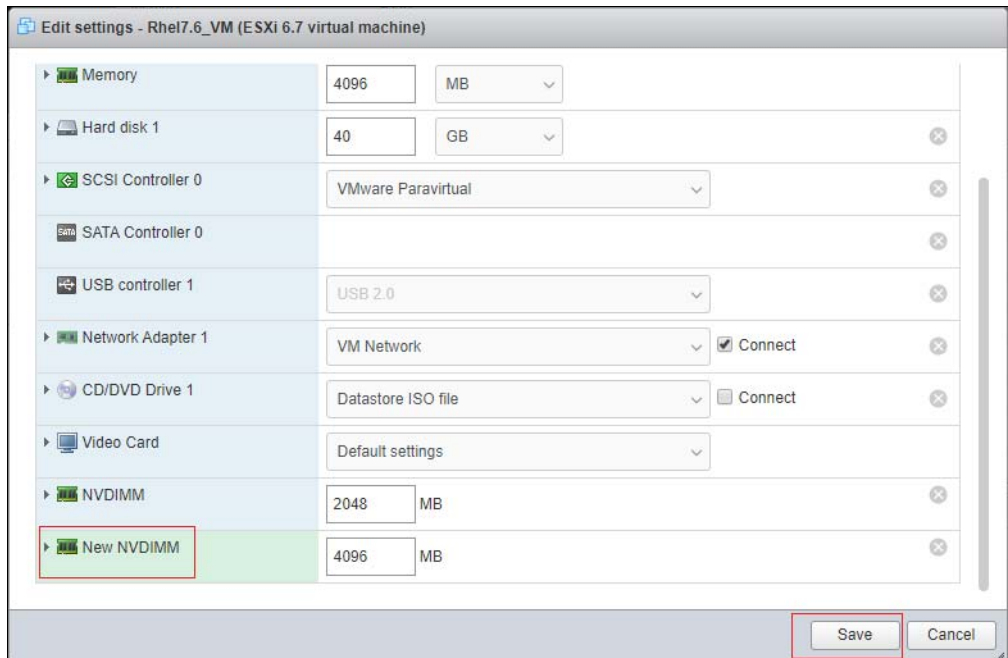


Figure 12 Configure virtual NVDIMM size on host web client

- 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).

The following command is to check virtual NVDIMM in the RHEL 7.6:

```
~# 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 13 shows the output of above commands.

```
[root@localhost dev]# ls /dev/pmem*
/dev/pmem0
[root@localhost dev]# sudo mkfs.ext4 /dev/pmem0
mke2fs 1.42.9 (28-Dec-2013)
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
Stride=0 blocks, Stripe width=0 blocks
262144 inodes, 1048576 blocks
52428 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=1073741824
32 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

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

[root@localhost dev]# sudo mount -o dax /dev/pmem0 /opt/pmem
[root@localhost dev]# 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     9.0M   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/pmem0                 3.9G      16M   3.6G   1% /opt/pmem
```

Figure 13 Command output on RHEL 7.6 virtual machine

9. Go to the vSphere web client and configure the RHEL 7.6 virtual machine to consume DCPMMs as a Virtual Persistent Memory Disk in Virtual Disk Mode.

**Tip:** In Virtual Disk Mode, the DCPMMs 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 14.

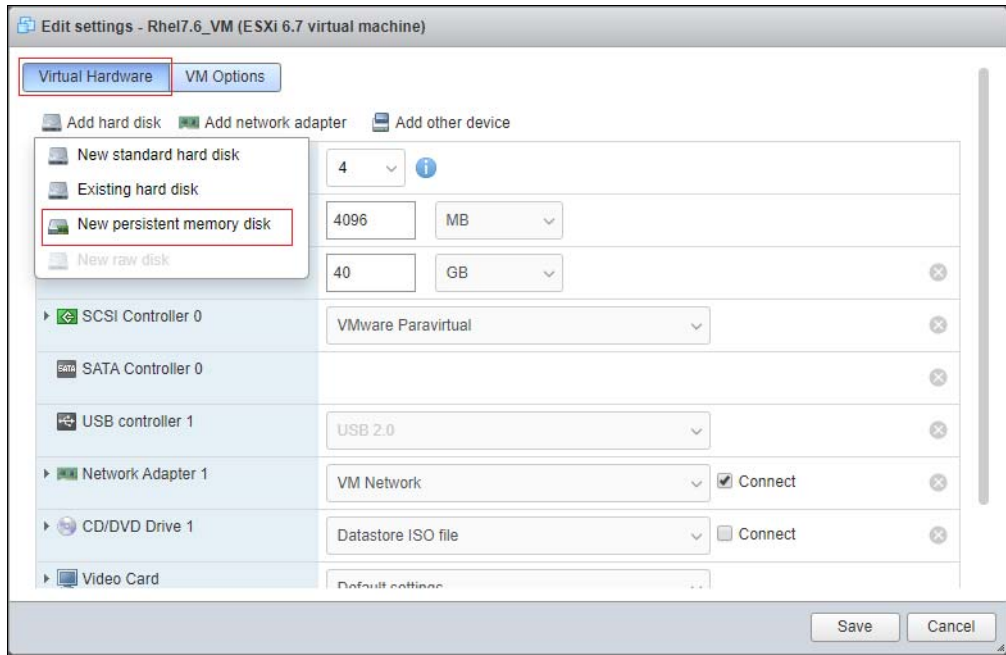


Figure 14 Add new virtual persistent memory disk on host web client

- 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 15.

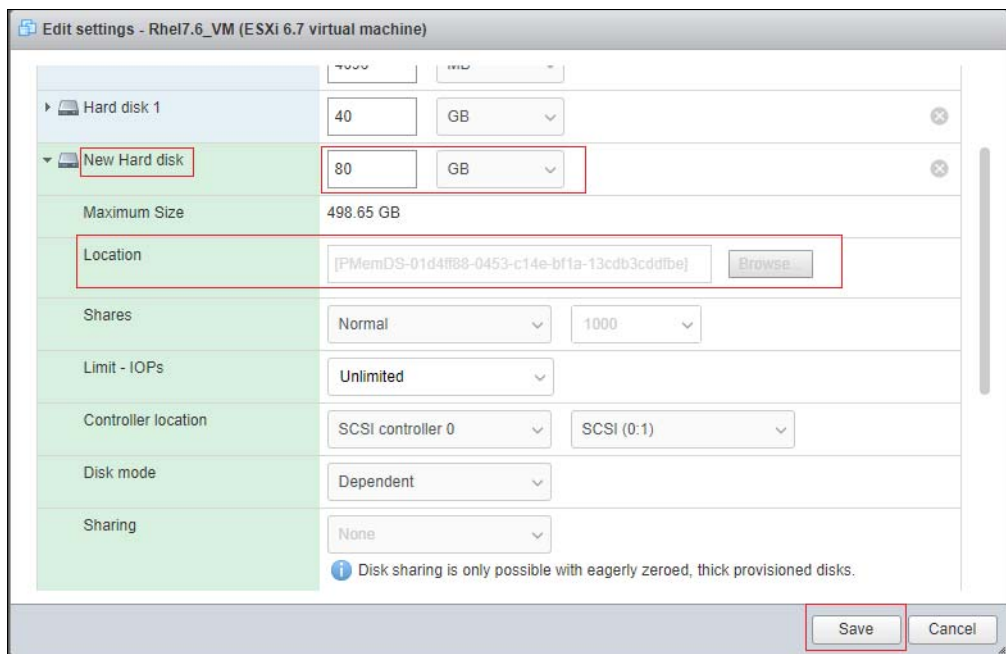


Figure 15 Configure virtual persistent memory disk size on host web client

- 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.

The following command is to check virtual persistent memory disk in the RHEL7.6:

```
~# 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 16 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 16 Command output on RHEL7.6 virtual machine



# Using DCPMMs in Mixed Mode

When DCPMMs are configured in Mixed Mode, 1-99% of DCPMM capacity acts as system memory. In this mode, some percentage of DCPMM 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 DCPMM capacity that is assigned as volatile system memory.

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

1. Refer to “Creating a goal to configure DCPMM operation modes” on page 6 to configure the DCPMMs in Mixed Mode.

In our lab setup, we installed our SR630 server with 2x 256 GB DCPMMs and 2x 32 GB DRAM DIMMs and configure the DCPMMs 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 DCPMM as 2LM, and can consume 50% of DCPMM 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 17 on page 17.

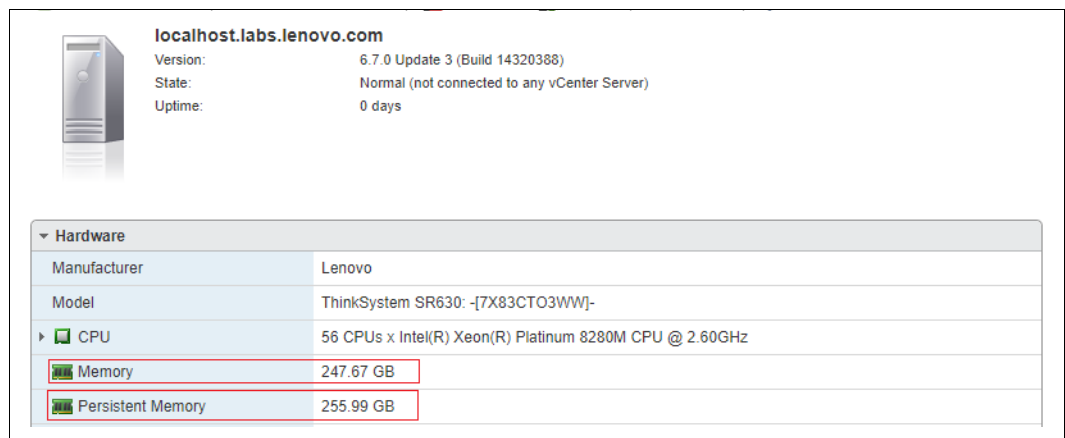


Figure 17 Check 2LM and PMEM on host web client

5. Refer to the “Using DCPMMs in Memory Mode” on page 9 and “Using DCPMMs in App Direct Mode” on page 10 to configure and use DCPMMs in Mixed mode.

## References

Review the following web pages for more information:

- ▶ Intel Optane DC Persistent Memory on ThinkSystem Servers  
<https://lenovopress.com/lp1066-intel-optane-dc-persistent-memory>

- ▶ 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 DC Persistent Memory  
<https://www.intel.com/content/www/us/en/architecture-and-technology/optane-dc-persistent-memory.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 Data Center 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.

Thanks to the following specialists for their contributions and suggestions:

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

## 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 January 10, 2020.

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.