

Intel® Optane™ Persistent Memory and SAP HANA® Platform Configuration

Technology overview and deployment guidelines for using Intel Optane persistent memory with SAP HANA.

SAP HANA altered the database landscape by providing an in-memory database-management system (DBMS). While main memory is fast and provides SAP HANA a performance benefit, it is also volatile. It also doesn't retain its contents during a server reboot, scheduled downtime, or power outage. When a server reboots, all SAP HANA data must be reloaded from storage, which can be time consuming. Intel is changing the game for in-memory databases with Intel Optane persistent memory (PMem). Intel Optane PMem is a higher density, non-volatile memory technology that retains its contents like traditional solid state drives (SSDs) and spinning disks, yet provides speeds that are similar to main memory.

SAP HANA is the first major database platform to support Intel Optane PMem. This gives SAP HANA the ability to expand beyond the density limitations of traditional memory while benefitting from the non-volatile capabilities of persistent memory. This new class of memory technology, combined with the power of SAP HANA 2.0 support package stack (SPS) 04 (or later), is ushering in a new generation of non-volatile in-memory technologies that promise to revolutionize the intelligent business.

This paper provides an overview of the benefits of combining servers equipped with Intel Optane PMem with SAP HANA 2. It then provides a high-level overview of the physical server requirements and steps to configure Intel Optane PMem for use with SAP HANA.

Application Performance: A Key Driver for Success

Application performance is a requirement for succeeding in today's business climate. Organizations rely on sophisticated applications, analytics, and DBMSs to provide a positive customer experience and to increase employee productivity. Customers demand constant availability and instantaneous responses, whereas employees rely on fast application performance from corporate databases and applications like enterprise resource planning (ERP) tools to get their jobs done as efficiently as possible.

Intel continually pushes the performance envelope with each new generation of Intel® Xeon® processors, system memory, and Intel® SSDs. As a result, Intel Optane PMem, a completely new class of memory, promises not only to continue the tradition of pushing the performance envelope, but also to revolutionize application performance and capabilities.

Intel Optane Persistent Memory: A Revolution for Enterprise Computing

Enterprises with large data-storage requirements often use a tiered memory and storage architecture, or heat map, for different types of data and workloads. These tiers include:

- **“Hot tier” DRAM-based memory:** Among the main memory and storage tiers, DRAM-based main memory typically is the fastest memory available, but it is also the most costly, and it provides relatively low density when compared with other storage technologies. Data in DRAM does not persist when the server is rebooted or powered off. For in-memory DBMSs, such as SAP HANA, enterprises use DRAM to provide the fastest speeds possible. But DRAM capacities have not scaled as quickly as other storage technologies, such as flash-based SSDs and spinning disks, limiting the size of in-memory databases and the amount of memory available per processor. Access speed is measured in nanoseconds.
- **“Warm tier” flash-based storage:** Based on NAND memory, flash-based storage is less costly, but much slower than DRAM. Unlike DRAM, flash-based storage retains data when it is turned off, or when the device that contains the flash-based storage restarts. Data in this tier is frequently accessed, but not as often as hot-tier data. Examples of warm data include the financial data that drives monthly or quarterly reports. Access speed is measured in microseconds.
- **“Cold tier” spinning disk-based storage:** Based on hard disk drives (HDDs), this storage tier is the lowest cost storage, but it is slow compared to flash-based storage and DRAM. Spinning disk-based storage is most often used for archival data or big data that is infrequently accessed. For example, long-term data storage from Internet of Things (IoT) devices stored on Apache Hadoop clusters often uses cold-tier storage. Access speed is measured in milliseconds.

Intel Optane PMem changes the data-storage landscape because it lets enterprises keep more data in the hot tier. Persistent memory provides near-DRAM performance in the same dual in-line memory module (DIMM) slot as DRAM.

Intel Optane PMem is also byte-addressable, just like DRAM, but it provides the same persistence as flash- and disk-based block storage. This persistence means that data is retained across server reboots and power outages. In addition, DRAM has traditionally had less capacity than flash- and disk-based storage. In contrast, persistent memory allows larger memory configurations at a cost that can drive down platform total cost of ownership (TCO).

Persistent Memory Can Revolutionize Application Capabilities and Speed

As data volumes grow and the need for fast analytics capabilities continues to increase, so does the need for larger memory capacities for in-memory databases. But with the size limitations of DRAM, enterprises have been unable to easily scale up their in-memory database deployments. The alternative has been scaling out using clustering and shared storage, but scale-out solutions can increase management complexity, introduce a performance penalty, and increase TCO.

Enterprises are also embracing technologies such as artificial intelligence (AI) and machine learning (ML). These technologies often require fast access to large datasets. Additionally, processing high-velocity data generated by technologies such as IoT places increased demands on in-memory database capacity.

Persistent memory brings new opportunities to in-memory databases and the applications that rely on them by providing higher memory density over DRAM. 2nd Generation Intel Xeon processors can provide up to 4.5 TB of byte-addressable capacity per socket. This means up to 18 TB of capacity in a four-socket system and up to 36 TB of capacity in an eight-socket system that

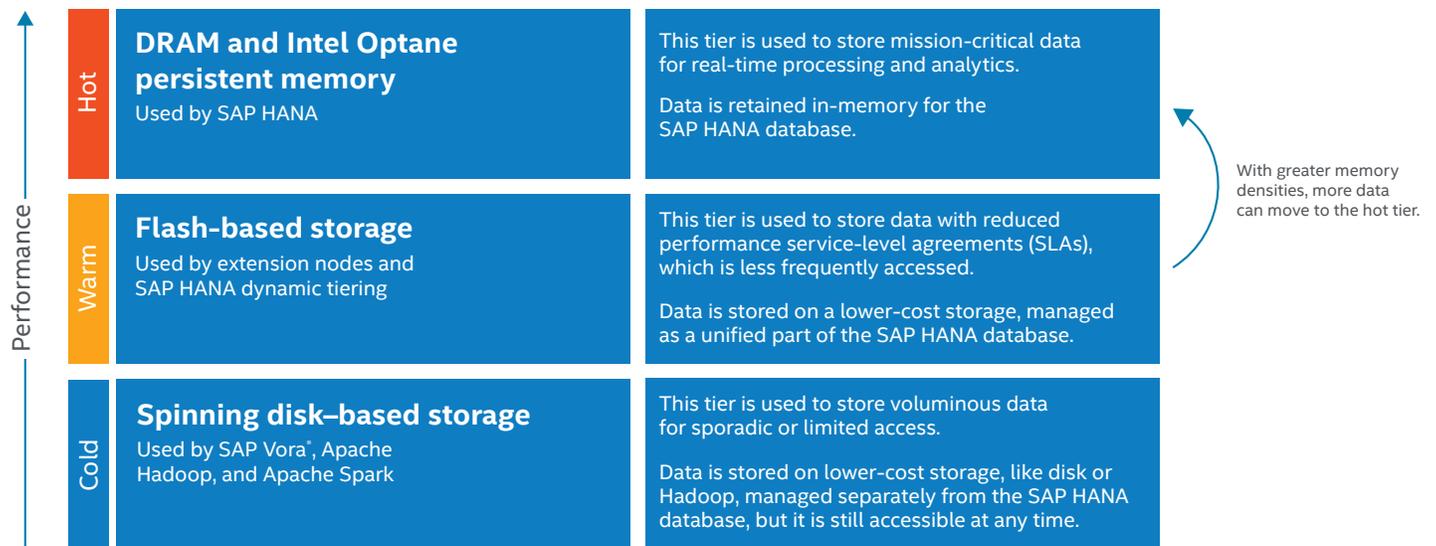


Figure 1. Intel Optane persistent memory provides new opportunities to expand the performance capabilities of in-memory databases by providing near-DDR4 DRAM speeds with the same persistence as flash- and disk-based storage

could run at near DRAM speeds. See [Tables 3 and 4](#) for examples of socket- and memory-slot configurations.

SAP HANA 2.0 SPS 04 (or Later) Support for Intel Optane Persistent Memory

Intel and SAP have collaborated to bring Intel Optane PMem support to SAP HANA. By providing more memory for SAP HANA, enterprises can keep older data, typically stored in the warm data tier, in the hot data tier, which can help reduce data-management complexity and TCO. In addition, the non-volatile nature and higher memory densities of Intel Optane PMem enable better business-continuity solutions and provide a platform for emerging memory-intensive technologies.

Double the Database Capacity for Similar Price and Performance

SAP HANA and application suites such as SAP S/4HANA® and the SAP® BW/4HANA® solution are ideal platforms for addressing the needs of the data-driven enterprise. However, organizations have experienced constraints due to DRAM's stagnant memory density and higher cost; multiple SAP HANA nodes and complex data tiering strategies are often required to address the growing data needs. Intel Optane PMem addresses these issues by providing increased density, up to 36 TB of total memory per server, which increases the memory scalability of SAP HANA, SAP S/4HANA, and SAP BW/4HANA.

Higher memory densities also enable more efficient use of hardware through system consolidation. Enterprises can scale up instead of scaling out to run larger transactional and analytics systems in the same physical hardware space. A reduced data center hardware footprint of SAP HANA lowers power and cooling costs. For example, Intel IT tested scaling up versus scaling out. The scale-up system showed up to 2.4 times query-performance improvement over the scale-out system, and could provide 52 percent more memory capacity.^{1,2}

With more memory available, organizations can also dramatically simplify their SAP HANA deployment architecture by consolidating smaller DRAM-only SAP HANA nodes onto fewer nodes configured with DRAM and Intel Optane PMem. Organizations can also upgrade servers that support Intel Optane PMem with higher

capacity memory modules as database sizes grow. With the ability to store more data in the hot and warm tiers, enterprises can simplify SAP HANA, SAP S/4HANA, and SAP BW/4HANA data-tier management, which helps lower management and infrastructure costs.

Organizations can also combine transactional and analytical workloads onto single SAP HANA nodes, further reducing deployment complexity. With Intel Optane PMem, organizations can realize a 39 percent cost savings per database terabyte.³ Larger memory configurations also let enterprises run memory-intensive workloads, such as predictive analytics, in less time, which leads to greater operational efficiencies.

More Memory Provides Efficient Business-Continuity Solutions

A robust business-continuity solution is key to any business's survival. With hardware solutions that use the expanded memory densities of Intel Optane PMem, enterprises can create more efficient and cost-effective business-continuity solutions.

The larger memory capacity of Intel Optane PMem expands the architectural options available to enterprises that rely on SAP HANA for critical business needs. Smaller DRAM-based SAP HANA nodes at a primary site can replicate data to a larger, offsite node that contains Intel Optane PMem and DRAM. The offsite node can also be used for multiple purposes, such as quality assurance (QA), development, or active-active replication, for less cost than an equivalent DRAM-only system. In the event of a primary site failure, the larger node can take over the demands of the primary-site SAP HANA nodes until the primary site is restored.

SAP HANA deployments that use Intel Optane PMem no longer require data to be loaded from slower storage systems when the node is restarted. This dramatically lowers data load times at startup and enables a quicker return to operations. Faster loading times reduce recovery time, which can cost enterprises more than \$100,000 for every hour of downtime.⁴ For example, Intel and SAP engineers demonstrated how a server, equipped with DRAM only and the SAP HANA 2 platform with 6 terabytes of data, required 50 minutes to start, including data loading.⁵ By contrast, a server equipped with a combination of Intel Optane PMem and DRAM reduced the start time of the SAP HANA 2 platform by 12.5x, to only 4 minutes.⁵

	Before (DRAM + NAND SSD)	After (DRAM + Intel Optane Persistent Memory)
System Architecture		
Data Access	Block Driver	Application Managed (App Direct)
Data Size	4 Kilobyte Block	64 Byte Cache Line
Read Latency (99.99%)	 Range: 25,000–650,000 nanoseconds Typical: 70,000 nanoseconds	 Range: 100s–2,000 nanoseconds Typical: 100s nanoseconds

Figure 2. Intel Optane persistent memory improves performance and quality of service (QoS) by eliminating data paging between DRAM and SSDs

Increased Memory Capacity Enables Emerging Technologies

Larger datasets and the higher velocity data being driven by new technologies place additional strains on the traditional data-deployment model. DRAM's lower memory density limits the amount of data SAP HANA and applications such as SAP S/4HANA and SAP BW/4HANA have available to them in the hot tier. As a result, more data must be stored in the slower warm and cold tiers.

The increased memory density and near-DRAM speed of Intel Optane PMem mean that SAP HANA can store more data in the hot tier, making low-latency access to greater amounts of data a reality for use cases such as AI and ML. Larger data models and faster response times help lower the time required for training AI, which can help lower the time to value of AI deployments.

Larger memory architectures are also ideal for capturing and processing high-velocity data, such as the data generated by large IoT deployments. Higher capacity also means more data can be maintained in memory to meet legal and compliance requirements while driving more robust real-time analytics across larger datasets.

Simplify the SAP HANA Landscape with Scale-up Architecture

For organizations that require scale-out capabilities, SAP HANA provides native support for extension nodes. SAP HANA extension nodes, combined with the additional capacity of Intel Optane PMem, can provide seamless integration of analytical and transactional platform architectures. Hot-tier data within extension nodes is stored in persistent memory modules (Intel Optane PMem) and DRAM. More data can be stored in persistent memory and accessed at near-DRAM speeds, which translates into less disk swapping and increased performance.

SAP HANA backup, recovery, and system replication function exactly the same with Intel Optane PMem. Backup and recovery operations work transparently with persistent memory in that backups save all persistence volume data without accessing persistent memory. Data can then be restored to the persistence volume. SAP systems engineers might consider deploying servers equipped Intel Optane PMem as replicated backup or disaster-recovery nodes.

System replication also works transparently with Intel Optane PMem, and it can be used in a mixed environment with systems configured with or without persistent memory.

Intel Optane Persistent Memory Configuration Considerations for SAP HANA

This section discusses the following persistent memory configuration steps:

- Platform support and operating modes for persistent memory
- Hardware sizing for SAP HANA 2.0 SPS 04 and higher
- DRAM-to-persistent memory ratios
- Memory-slot configuration and Intel Optane PMem placement
- BIOS configuration
- Operating system support and configuration

- Enabling persistent memory support in SAP HANA

Platform Support and Operating Modes for Persistent Memory

Intel Optane PMem is supported on servers equipped with 2nd Generation Intel Xeon Gold processors and Intel Xeon Platinum processors. Two primary modes are supported: App Direct Mode, including Block over App Direct, and Memory Mode. App Direct Mode is the only mode that is currently supported by SAP HANA 2. In App Direct Mode, Intel Optane PMem modules appear as byte-addressable memory resources that are controlled by SAP HANA 2. In this mode, the persistent memory space is controlled directly by SAP HANA.

Memory-Sizing Recommendations for SAP HANA

Sizing for SAP HANA deployments can be accomplished using a fixed core-to-memory ratio based on workload type, or by doing a self-assessment using the SAP HANA Tailored Datacenter Integration (TDI) approach and sizing tools provided by SAP.

The utilization of Intel Optane PMem requires an appropriate sizing to help ensure the best performance of an SAP HANA installation. SAP provides a decision tree at <https://sap.com/sizing> that can help provide the right tools and procedures.

Green-Field Sizing

In case of a brand new installation, Quick Sizer enables customers to define the key performance indicators (KPIs) and criteria for an SAP HANA system. These defined parameters drive an SAP recommendation for the required hardware, memory, and storage. Quick Sizer recommends a ratio between DRAM and Intel Optane PMem that supports the defined KPIs.

For more information about Quick Sizer (including registration for training), visit sap.com/about/benchmark/sizing.quick-sizer.html.

Migrations/Brown-Field Sizing

In contrast to green-field sizing, migration projects are based on existing installations, and a sizing can be done differently. SAP Sizing Reports provides the tools and procedures to determine the hardware requirements from an existing system (see Table 1). This also supports migrations from third-party databases to SAP HANA, by which a projection of the anyDB data into SAP HANA is made in order to achieve appropriate sizing.

The following SAP notes are also relevant to working with persistent memory:

- [SAP Note 2700084](#)—FAQ Persistent Memory
- [SAP Note 2618154](#)—SAP HANA Persistent Memory
- [SAP Note 2813454](#)—Persistent Memory Configuration with SAP BW/4HANA

Note: Accessing SAP Notes Sizing requires signing in with an SAP account.

SAP offers additional expert sizing through the SAP MaxAttention™ program, which can analyze an existing environment and its requirements to design and size an SAP HANA landscape, helping to drive the best performance and best utilization of Intel Optane PMem.

Table 1. SAP sizing resources based on deployment

Starting Point	Sizing for SAP Business Warehouse on SAP HANA and SAP BW/4HANA	Sizing for SAP Business Suite on SAP HANA and SAP S/4HANA	Sizing on Platforms Other Than the SAP NetWeaver® Technology Platform
Green-Field/New Implementation	Quick Sizer	Quick Sizer	Sizing guide available: sap.com/sizing
Migration/Brown Field	SAP Note 2296290 (Sizing Report for SAP BW4/HANA)	SAP Note 1872170 (Sizing Report for SAP S/4HANA)	SAP Note 1514966 (SAP HANA: Sizing SAP In-Memory Database)

DRAM-to-PMem Ratios

Intel Optane PMem modules must be installed with DRAM DIMMs in the same system. Intel Optane PMem modules do not function without any DRAM DIMMs installed.

In any multi-socket (multi-CPU) configuration, each socket contains two integrated memory controllers (IMCs). Each memory controller is connected to three double data rate (DDR) memory channels that are then connected to two physical DIMM/persistent memory slots. In this configuration, a maximum of 12 memory slots per CPU socket can be configured with a combination of Intel Optane PMem modules and DRAM DIMMs. Table 2 lists the maximum number of memory slots available in various CPU socket configurations.

Table 2. Maximum number of memory slots available in common CPU socket configurations

Number of CPU Sockets	Maximum Number of Memory Slots
One	12
Two	24
Four	48
Eight	96
Sixteen	192

SAP HANA 2 currently supports various capacity ratios between DRAM DIMMs and Intel Optane PMem modules. Ratio examples include:

- **1:1 ratio:** On each 2-slot channel, a single 128 GB DDR4 DIMM is matched with a single 128 GB Intel Optane PMem module, or a 256 GB DRAM DIMM is matched with a single 256 GB Intel Optane PMem module.
- **1:2 ratio:** On each 2-slot channel, a 128 GB DRAM DIMM is matched with a 256 GB Intel Optane PMem module, or a 64 GB DDR4 DIMM is matched with a 128 GB Intel Optane PMem module.
- **1:4 ratio:** On each 2-slot channel, a 128 GB DDR4 DIMM is matched with a 512 GB Intel Optane PMem module, or a

64 GB DRAM DIMM is matched with a 256 GB Intel Optane PMem module.

Different-sized Intel Optane PMem modules and DRAM DIMMs can be used together as long as supported ratios are maintained. Tables 3 and 4 illustrate some common DRAM and Intel Optane PMem allocations for SAP HANA analytics and transactional systems, respectively.

Table 3. Common DRAM and Intel Optane persistent memory allocations for SAP HANA analytics systems

Memory Configuration (System, DRAM, and Intel Optane PMem)	Size	DB Size	DRAM: PMem Ratio
2 CPU: 1.5 TB DRAM (128 GB DDR) + 1.5 TB PMem (128 GB PMem)	M	3 TB	1:1
4 CPU: 3 TB DRAM (128 GB DDR) + 3 TB PMem (128 GB PMem)	L	6 TB	1:1
8 CPU: 6 TB DRAM (128 GB DDR) + 6 TB PMem (128 GB PMem)	XL	12 TB	1:1
8 CPU: 12 TB DRAM (256 GB DDR) + 12 TB PMem (256 GB PMem)	XXL	24 TB	1:1
16 CPU: 12 TB DRAM (128 GB DDR) + 12 TB PMem (128 GB PMem)	XXL	24 TB	1:1

Table 4. Common DRAM and Intel Optane persistent memory allocations for SAP HANA transactional systems

Memory Configuration (System, DRAM, and Intel Optane PMem)	Size	DB Size	DRAM: PMem Ratio
2 CPU: 768 GB DRAM (64 GB DDR) + 1.5 TB PMem (128 GB PMem)	S	2.25 TB	1:2
2 CPU: 1.5 TB DRAM (128 GB DDR) + 3 TB PMem (256 GB PMem)	M	4.5 TB	1:2
4 CPU: 1.5 TB DRAM (64 GB DDR) + 3 TB PMem (128 GB PMem)	M	4.5 TB	1:2
4 CPU: 3 TB DRAM (128 GB DDR) + 6 TB PMem (256 GB PMem)	L	9 TB	1:2
8 CPU: 3 TB DRAM (64 GB DDR) + 6 TB PMem (128 GB PMem)	L	9 TB	1:2
4 CPU: 3 TB DRAM (128 GB DDR) + 12 TB PMem (512 GB PMem)	XL	15 TB	1:4
8 CPU: 3 TB DRAM (64 GB DDR) + 12 TB PMem (256 GB PMem)	XL	15 TB	1:4
8 CPU: 6 TB DRAM (128 GB DDR) + 12 TB PMem (256 GB PMem)	XL	18 TB	1:2
16 CPU: 6 TB DRAM (64 GB DDR) + 12 TB PMem (128 GB PMem)	XL	18 TB	1:2
16 CPU: 12 TB DRAM (128 GB DDR) + 24 TB PMem (256 GB PMem)	XXL	36 TB	1:2

Memory-Slot Configuration and Intel Optane Persistent Memory Module Placement

While servers configured with Intel Optane PMem accept a number of different memory configurations, SAP HANA 2 supports fully populated servers only. This means that all DIMM slots in a server need to be used in a 1:1 ratio of DRAM to PMem modules (for example, 6 DRAM and 6 PMem modules per socket). PMem modules should be placed in the slots closest to the CPU.

BIOS Configuration

Intel Optane PMem is enabled through BIOS settings. Because each server vendor's BIOS layout is different, this section discusses general BIOS settings and pre-boot UEFI command-line utilities. To configure servers from a specific OEM, refer to the OEM's server documentation for persistent memory-specific information.

Enabling Intel Optane PMem support in the BIOS generally follows these steps:

1. During the server boot process, enter the BIOS by pressing the key that activates the server's BIOS settings.
2. Navigate to the BIOS option that correlates to the Intel Optane PMem configuration.
3. Set **Volatile Mode** to **Auto**. Do not select **1LM (AppDirect)** or **2LM (Memory Mode)**.
4. Save the changes and exit the BIOS.

Once Intel Optane PMem is enabled in the BIOS, the system can be further configured using the UEFI shell or operating system tools.

Operating System Support and Configuration

Intel Optane PMem is supported by the following SAP HANA and operating system versions:⁶

- SAP HANA versions:
 - SAP HANA 2.0 SPS 03, revision 35 (2.00.035)
 - SAP HANA 2.0 SPS 04, any revision
- Operating systems:
 - SUSE Linux Enterprise Server (SLES)/SLES for SAP applications
 - SLES 12 Service Pack 4 (SP4) or higher
 - SLES 15 GA or higher
 - Red Hat Enterprise Linux (RHEL)/RHEL for SAP solutions
 - RHEL 7.5 or higher

For information on the most current list of SAP HANA supported operating systems, please consult SAP Notes [2235581](#) and [2618154](#).

Once Intel Optane PMem is enabled in the server BIOS, it is then managed using either a UEFI pre-boot shell or from a supported operating system using the Intel Optane PMem command-line interface (CLI). Both the UEFI and operating system persistent memory CLIs provide similar management capabilities, including persistent memory discovery, provisioning, maintenance, and monitoring. The following sections describe the general procedure to configure Intel Optane PMem for use with SAP HANA. For more details on configuring Intel Optane PMem on Linux, please see <https://software.intel.com/en-us/articles/quick-start-guide-configure-intel-optane-dc-persistent-memory-on-linux>.

General Steps: Configure Intel Optane Persistent Memory for SAP HANA⁶

1. Install the management utilities.
2. Create app direct regions (goal)—use interleaving.
3. Reboot the server—required to enable new configuration.
4. Create app direct namespaces.
5. Create a file system on the namespace device.
6. Configure SAP HANA to use the persistent memory file system.
7. Restart SAP HANA to activate and start using Intel Optane PMem.

Install the Management Utilities

To manage Intel Optane PMem modules, the `ipmctl` and `ndctl` utilities must be installed. Follow the instructions for RHEL or SUSE detailed below. For additional information, refer to the [NDCTL User Guide](#) and [IPMCTL User Guide](#).

RHEL for SAP HANA**IPMCTL**

The `ipmctl` package is available on RHEL for SAP HANA 7.6 or later in the Extra Packages for Enterprise Linux (EPEL) repository.

Step 1

Verify the EPEL repository is active with the following command:

```
$ yum repolist
```

For example:

```
$ sudo yum repolist
repo id      repo name          status
epel/x86_64  Extra Packages for Enterprise Linux 7 - x86_64  13,217
```

If the EPEL repository is not listed, install and activate it using the following command:

```
$ sudo yum install epel-release
```

Step 2

Query the package repository to confirm that `ipmctl` is available with the following command:

```
sudo yum info ipmctl
```

Example:

```
$ sudo yum info ipmctl
Available Packages
Name       : ipmctl
Arch      : x86_64
Version   : 01.00.00.3474
Release   : 2.el7
Size      : 70 k
Repo      : epel/x86_64
Summary   : Utility for managing Intel Optane persistent memory modules
URL       : https://github.com/intel/ipmctl
License   : BSD
Description : Utility for managing Intel Optane DC persistent memory modules
            : Supports functionality to:
            : Discover DCPMMs on the platform.
            : Provision the platform memory configuration.
            : View and update the firmware on DCPMMs.
            : Configure data-at-rest security on DCPMMs.
            : Monitor DCPMM health.
            : Track performance of DCPMMs.
            : Debug and troubleshoot DCPMMs.
```

Step 3

Install the ipmctl package with the following command:

```
sudo yum install ipmctl
```

NDCTL

The ndctl package is available on RHEL for SAP HANA 7.6 or later.

Step 1

Query the repository to identify if the ndctl package is available with the following command:

```
yum search ndctl
```

Step 2

Install the ndctl package with the following command:

```
yum install ndctl
```

SLES for SAP HANA**IPMCTL**

The ipmctl package is available in the default package repository on SUSE for SAP HANA 12.4 (v12sp4) or later.

Step 1

Query the package repository to confirm that ipmctl is available with the following command:

```
sudo zypper info ipmctl
```

For example:

```
$ sudo zypper info ipmctl
Information for package ipmctl:
-----
Repository      : SLES12-SP4-Updates
Name            : ipmctl
Version        : 01.00.00.3440-3.8.2
Arch           : x86_64
Vendor         : SUSE LLC <https://www.suse.com/>
Support Level  : Level 3
Installed Size : 3.2 MiB
Installed      : No
Status        : not installed
Source package : ipmctl-01.00.00.3440-3.8.2.src
Summary       : Utility for managing Intel Optane DC persistent memory modules
Description   : Utility for managing Intel Optane DC persistent memory modules Supports
functionality to:
  * Discover PMMs on the platform.
  * Provision the platform memory configuration.
  * View and update the firmware on PMMs.
  * Configure data-at-rest security on PMMs.
  * Monitor PMM health.
  * Track performance of PMMs.
  * Debug and troubleshoot PMMs.
```

Step 2

Install the ipmctl package with the following command:

```
sudo zypper install ipmctl
```

NDCTL

The ndctl package is available for SLES for SAP HANA v12.4 (v12sp4) or later.

Step 1

Query the repository to identify if ndctl is delivered with the following command:

```
sudo zypper search ndctl
```

For example:

```
$ sudo zypper search ndctl
SocketID | ISetID | PersistentMemoryType | Capacity | FreeCapacity | HealthState
=====
0x0000 | 0x37c27f48425c2ccc | AppDirect | 3012.0 GiB | 3012.0 GiB | Healthy
0x0001 | 0xb96c7f4873572ccc | AppDirect | 3012.0 GiB | 3012.0 GiB | Healthy
0x0002 | 0xdda47f4897592ccc | AppDirect | 3012.0 GiB | 3012.0 GiB | Healthy
0x0003 | 0xeb087f4806592ccc | AppDirect | 3012.0 GiB | 3012.0 GiB | Healthy

S | Name | Summary | Type
--|-----|-----|-----
| libndctl16 | Management library for "libnvdimm" subsystem devices (Non-volatile Memory) | package
| ndctl | Manage "libnvdimm" subsystem devices (Non-volatile Memory) | package
```

Step 2

Install the ndctl package with the following command:

```
sudo zypper install ndctl
```

Create App Direct Regions

For SAP HANA, Intel Optane PMem must be provisioned in App Direct Mode with interleaving enabled. Interleaving increases the throughput of reads and writes to persistent memory. The default Create Goal command creates an interleaved region configured for App Direct Mode:

```
ipmctl create -goal PersistentMemoryType=AppDirect
```

Useful Tip: Intel Optane PMem Discovery

Prior to provisioning Intel Optane PMem, users can invoke the Intel Optane PMem commands shown in Table 5 from either the UEFI shell or the Linux version of the Intel Optane PMem shell to display the individual Intel Optane PMem modules and DRAM DIMMs installed on the system and to view information about each module.

Table 5. Intel Optane PMem discovery commands provide information specific to the physical persistent memory and DRAM configuration

Intel Optane PMem Discovery Command	Result
<code>ipmctl show -topology</code>	Displays both the Intel Optane PMem and DRAM modules installed in the system.
<code>ipmctl show -dimm</code>	Displays the Intel Optane PMem modules installed in the system.
<code>ipmctl show -a -dimm</code>	Displays detailed information about the Intel Optane PMem modules.
<code>ipmctl show -memoryresources</code>	Displays how the Intel Optane PMem capacity is provisioned for the App Direct (AD) Mode and Memory Mode (MM). For example: Capacity=12062.3 GiB MemoryCapacity=0.0 GiB AppDirectCapacity=12048.0 GiB UnconfiguredCapacity=0.0 GiB InaccessibleCapacity=14.3 GiB ReservedCapacity=0.0 GiB

Complete the memory provisioning (goal) process by rebooting the server to enable the new settings.

After the reboot, the newly created DIMM-interleave-sets are represented as persistent memory “regions” of App Direct Mode capacity. To view the region setup, use the List Regions command:

```
ndctl list -regions
```

This command returns output similar to the following:

```
[
  {
    "dev": "region1",
    "size": "756.00 GiB (811.75 GB)",
    "available_size": "756.00 GiB (811.75 GB)",
    "max_available_extent": "756.00 GiB (811.75 GB)",
    "type": "pmem",
    "iset_id": "0x6f3e7f4888992ccc",
    "persistence_domain": "memory_controller"
  },
  {
    "dev": "region0",
    "size": "756.00 GiB (811.75 GB)",
    "available_size": "756.00 GiB (811.75 GB)",
    "max_available_extent": "756.00 GiB (811.75 GB)",
    "type": "pmem",
    "iset_id": "0x1c9e7f487b952ccc",
    "persistence_domain": "memory_controller"
  }
]
```

Alternatively, use the following Intel Optane PMem command:

```
ipmctl show -region
```

This command will display similar information in table format and verify that Intel Optane PMem is set to App Direct Mode. This command returns output similar to the following:

To determine which Intel Optane PMem modules are part of a particular region, use the List command with the Regions and DIMMs modifiers. Each module will be listed within the mappings group:

```
ndctl list --regions --dimms
```

This command returns output similar to the following:

```
{
  "dimms": [
    {
      "dev": "nmem1",
      "id": "8680-a2-1730-000006a5",
      "handle": 257,
      "phys_id": 45
    },
    {
      "dev": "nmem3",
      "id": "8680-a2-1730-00000632",
      "handle": 4353,
      "phys_id": 69
    },
    {
      "dev": "nmem0",
      "id": "8680-a2-1730-0000059f",
      "handle": 1,
      "phys_id": 33
    }
  ]
}
```

```

    "dev": "nmem2",
    "id": "8680-a2-1730-000004f3",
    "handle": 4097,
    "phys_id": 57
  }
],
"regions": [
  {
    "dev": "region5",
    "size": 268435456000,
    "available_size": 268435456000,
    "max_available_extent": 268435456000,
    "type": "pmem",
    "iset_id": 8743142107817513552,
    "mappings": [
      {
        "dimm": "nmem3",
        "offset": 268435456,
        "length": 134217728000,
        "position": 1
      },
      {
        "dimm": "nmem2",
        "offset": 268435456,
        "length": 134217728000,
        "position": 0
      }
    ]
  },
  {
    "dev": "region4",
    "size": 268435456000,
    "available_size": 0,
    "max_available_extent": 0,
    "type": "pmem",
    "iset_id": 8015984336272174284,
    "mappings": [
      {
        "dimm": "nmem1",
        "offset": 268435456,
        "length": 134217728000,
        "position": 1
      },
      {
        "dimm": "nmem0",
        "offset": 268435456,
        "length": 134217728000,
        "position": 0
      }
    ]
  }
]
"persistence_domain": "memory_controller"
}
]
}

```

Create App Direct Namespaces

Regions can now be divided into one or more namespaces in order for the capacity to be surfaced to the Linux operating system and used by applications. Just as an SSD can be carved into namespaces, persistent memory namespaces represent the unit of storage that appears as a device that can be used for input/output (I/O). To create a namespace, use the Create Namespace command:

```
ndctl create-namespace
```

The default syntax will create a namespace with direct access (DAX) support, mode=fsdax, which SAP HANA requires to properly utilize Intel Optane PMem. The following example is equivalent syntax to explicitly set the mode; it also provides the option to specify a target region (in this example, the first region is specified, region0).

```
ndctl create-namespace --mode fsdax --region 0
```

This command needs to be executed once for each CPU in the server. For example, in a four-socket server, the command needs to be executed four times:

```
ndctl create-namespace --mode fsdax --region 0
ndctl create-namespace --mode fsdax --region 1
ndctl create-namespace --mode fsdax --region 2
ndctl create-namespace --mode fsdax --region 3
```

For servers with many sockets (such as 8-socket or 16-socket servers), this can be onerous. A generic solution is to run the following command, which creates one namespace per region on the system, no matter how many CPUs and regions there are. This creates an internal FOR loop and will terminate on error:

```
ndctl create-namespace --continue
```

Note: This command requires ndctl v67 or later.

For a server with two sockets, this command returns output similar to the following when creating two namespaces:

```
{
  "dev": "namespace1.0",
  "mode": "fsdax",
  "map": "dev",
  "size": "1488.37 GiB (1598.13 GB)",
  "uuid": "4601589f-667e-4775-8d05-dbb0dc6af6cc8",
  "sector_size": 512,
  "align": 2097152,
  "blockdev": "pmem1"
}
{
  "dev": "namespace0.0",
  "mode": "fsdax",
  "map": "dev",
  "size": "1488.37 GiB (1598.13 GB)",
  "uuid": "dce96530-58fc-4864-869b-8f7fbffa1bed",
  "sector_size": 512,
  "align": 2097152,
  "blockdev": "pmem0"
}
```

Note that if a system has jq installed, it can extract the number of regions with the command "ndctl list -regions." However, jq is not commonly installed by default. The following script will work on any system:

```
# for region in `ndctl list --regions | grep -i region | cut -f4 -d '`
> do
> ndctl create-namespace --mode fsdax --region ${region}
> done
{
  "dev": "namespace1.0",
  "mode": "fsdax",
  "map": "dev",
  "size": "1488.37 GiB (1598.13 GB)",
  "uuid": "027090d6-0cee-4f0d-b387-2b966f456f1c",
  "sector_size": 512,
  "align": 2097152,
  "blockdev": "pmem1"
```

```

}
{
  "dev": "namespace0.0",
  "mode": "fsdax",
  "map": "dev",
  "size": "1488.37 GiB (1598.13 GB)",
  "uuid": "ee64754f-98c6-4a0a-bef6-f2f849020ae8",
  "sector_size": 512,
  "align": 2097152,
  "blockdev": "pmem0"
}

```

To display namespaces, use the List command. Namespaces are shown by default without additional arguments:

```
ndctl list
```

This command returns output similar to the following:

```

[
  {
    "dev": "namespace1.0",
    "mode": "fsdax",
    "map": "dev",
    "size": "744.19 GiB (799.06 GB)",
    "uuid": "56ca2d82-36fc-4bf6-b434-b5315957f5ca",
    "blockdev": "pmem1"
  },
  {
    "dev": "namespace0.0",
    "mode": "fsdax",
    "map": "dev",
    "size": "744.19 GiB (799.06 GB)",
    "uuid": "5ec652fa-3642-4e93-ad76-90d70872f3a3",
    "blockdev": "pmem0"
  }
]

```

Create a File System Using the Namespace

Once created, the "fsdax" namespace is assigned a device name by the operating system, usually /dev/pmemX, where X denotes the NUMA node mapping of the underlying hardware. Use either CLI or graphical user interface (GUI) tools to create an ext4 or XFS file system on each of the /dev/pmemX devices—typically one device per NUMA node or socket. For example, on a four-socket system:

```

mkfs.xfs /dev/pmem0
mkfs.xfs /dev/pmem1
mkfs.xfs /dev/pmem2
mkfs.xfs /dev/pmem3

```

Note: If you are using xfsprogs 5.1 or later, disable the reflink feature. Refer to Known Issue 2 in the [Troubleshooting](#) section for more information.

The file systems created on the Intel Optane PMem device would then need to be mounted to an operating system directory with the DAX option.

Users can create new directories or use existing ones. Multiple SAP HANA systems (SIDs) can share the same persistent memory device folders, as long as the proper permissions are assigned (see Configure SAP HANA to Use Intel Optane PMem). For example:

```

mkdir /mnt/pmem0
mkdir /mnt/pmem1
mkdir /mnt/pmem2
mkdir /mnt/pmem3
mount -t xfs -o dax /dev/pmem0 /mnt/pmem0
mount -t xfs -o dax /dev/pmem1 /mnt/pmem1
mount -t xfs -o dax /dev/pmem2 /mnt/pmem2
mount -t xfs -o dax /dev/pmem3 /mnt/pmem3
mount | grep dax
/dev/pmem0 on /mnt/pmem0 type xfs (rw,relatime,attr2,dax,inode64,noquota)
/dev/pmem1 on /mnt/pmem1 type xfs (rw,relatime,attr2,dax,inode64,noquota)

```

```
/dev/pmem2 on /mnt/pmem2 type xfs (rw,relatime,attr2,dax,inode64,noquota)
/dev/pmem3 on /mnt/pmem3 type xfs (rw,relatime,attr2,dax,inode64,noquota)
```

Users can ensure the file system mounts are available after the server reboots by adding entries to `/etc/fstab`. For example, the following shows the entries added to `/etc/fstab` to automatically mount the file systems at boot time:

```
/dev/pmem0 /mnt/pmem0 xfs dax 0 0
/dev/pmem1 /mnt/pmem1 xfs dax 0 0
/dev/pmem2 /mnt/pmem2 xfs dax 0 0
/dev/pmem3 /mnt/pmem3 xfs dax 0 0
```

It is recommended that you use the UUID naming convention in `/etc/fstab`, rather than `/dev/pmem`, because the devices can enumerate differently at boot time, leading to data being mounted in the wrong location.

Identify the UUID for each file system on persistent memory:

```
# blkid | grep /dev/pmem
/dev/pmem0: UUID="b376b4a9-eabb-4fd9-a606-6b6a33474b09" BLOCK_SIZE="4096" TYPE="xfs"
/dev/pmem1: UUID="2db6d5b5-df59-4635-a14d-61c242b22ac9" BLOCK_SIZE="4096" TYPE="xfs"
/dev/pmem2: UUID="5655d53f-a623-4997-a46c-1552b518d84c" BLOCK_SIZE="4096" TYPE="xfs"
/dev/pmem3: UUID="FKSAAm-s0fM-kFAE-I7J7-raEA-Di6E-0ydFuu" BLOCK_SIZE="4096" TYPE="xfs"
```

Use the UUID for the first column of the `/etc/fstab`. For example:

```
UUID="b376b4a9-eabb-4fd9-a606-6b6a33474b09" /mnt/pmem0 xfs dax 0 0
UUID="2db6d5b5-df59-4635-a14d-61c242b22ac9" /mnt/pmem1 xfs dax 0 0
UUID="5655d53f-a623-4997-a46c-1552b518d84c" /mnt/pmem2 xfs dax 0 0
UUID="FKSAAm-s0fM-kFAE-I7J7-raEA-Di6E-0ydFuu" /mnt/pmem3 xfs dax 0 0
```

Confirm the `/etc/fstab` entries are correct to avoid any boot issues. Unmount the file systems, and then remount them.

```
# umount /dev/pmem*
# mount -a
# mount -v | grep pmem
```

See Known Issue 1 in the [Troubleshooting](#) section for potential boot issues on large memory systems.

Configure SAP HANA to Use Intel Optane PMem

Assign the SAP HANA admin user (<sid>adm) proper permission to the mounted directories. For example:

```
chown db1adm:sapsys /mnt/pmem*
```

With multiple systems on the same server, this might look like the following:

```
mkdir /mnt/pmem0/DB1
...
mkdir /mnt/pmem3/DB1
chown db1adm:sapsys /mnt/pmem*/DB1
...
mkdir /mnt/pmem0/DB2
...
mkdir /mnt/pmem3/DB2
chown db2adm:sapsys /mnt/pmem*/DB2
```

Set the SAP HANA Base Path

The directory that SAP HANA uses as its base path must point to the journaled ext4 or XFS file system created in the previous steps. The base path location is defined as a configuration parameter in the persistence section of the SAP HANA `global.ini` file. Users can change the value of the following parameter using the corresponding user interface (UI) in SAP HANA Studio (see Figure 3) or SAP HANA cockpit:

global.ini > persistence > basepath_persistent_memory_volumes

This section can contain multiple directory locations separated by semicolons, and it can be set at the system (SID) or host level with a mix of Intel Optane PMem and DRAM-only servers and SIDs.

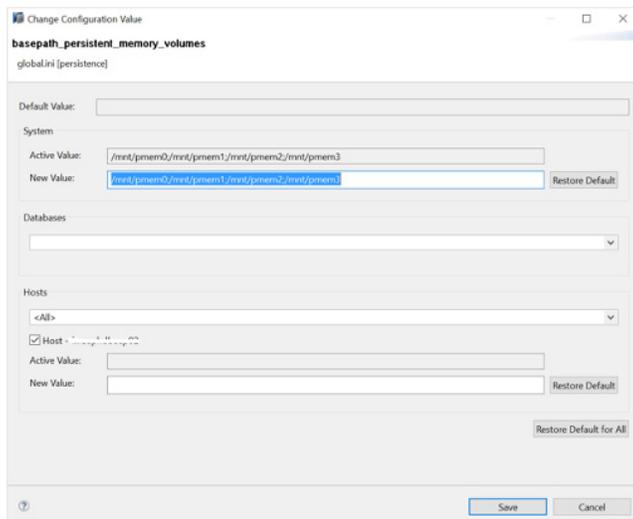


Figure 3. Setting the SAP HANA Base Path via SAP HANA Studio

Users can also directly edit the text of the corresponding global.ini config file shown in the following example for a four-socket server (DB1 is the SID in this example):

```
Filename: /hana/shared/DB1/global/hdb/custom/config/global.ini
# global.ini last modified 2020-01-01 00:00:46.413346 by hdbnameserver [persistence]
basepath_persistent_memory_volumes=/mnt/pm0;/mnt/pm1;/mnt/pm2;/mnt/pm3
```

For a multi-system server (in which DB1 and DB2 are the SIDs), this would look like the following:

```
Filename: /hana/shared/DB1/global/hdb/custom/config/global.ini
# global.ini last modified 2020-01-01 00:00:46.413346 by hdbnameserver
[persistence]
basepath_persistent_memory_volumes=/mnt/pm0/DB1;/mnt/pm1/DB1;/mnt/pm2/DB1;/mnt/pm3/DB1

Filename: /hana/shared/DB2/global/hdb/custom/config/global.ini
# global.ini last modified 2020-01-01 00:00:46.413346 by hdbnameserver
[persistence]
basepath_persistent_memory_volumes=/mnt/pm0/DB2;/mnt/pm1/DB2;/mnt/pm2/DB2;/mnt/pm3/DB2
```

Restart SAP HANA in order to activate the new parameter settings. At startup, SAP HANA tests for a DAX-enabled file system at the location defined in the base path. Once SAP HANA verifies that the file system is DAX-enabled, all tables will use persistent memory by default. This default behavior can be overridden if necessary. (See <https://help.sap.com/>—[SAP HANA Administration Guide](#) or [SAP HANA SQL Reference Guide for SAP HANA Platform](#) for more details.) Save points help ensure that the contents of data in persistent memory are consistent with the persistence and data log volumes.

Troubleshooting

Known Issue 1

Systems with a large amount of persistent memory and with automount entries in /etc/fstab might experience device timeouts, which could cause the boot process to fail and enter emergency mode.

Applies To

- Systems with large amounts of memory and multiple physical CPUs
- Red Hat Enterprise Linux 7.x
- Red Hat Enterprise Linux 8.0 (without Red Hat Security Advisor [RHSA] 2019:1959)
- SUSE Linux Enterprise Server 12 SPx
- SUSE Linux Enterprise Server 15

Issue

During boot, persistent memory devices might time out, causing the system to drop to emergency mode. This issue can occur on systems with a large amount of persistent memory, with two or more physical CPUs, and when automount entries exist in the /etc/fstab for the persistent memory devices.

The console, dmesg, and journalctl will report errors similar to the following for each persistent memory device.

```
[ TIME ] Timed out waiting for device dev-pmem1.device.
[DEPEND] Dependency failed for /hana/pmem1.
[DEPEND] Dependency failed for Local File Systems.
[DEPEND] Dependency failed for Migrate local SELinux policy changes from the old store structure to
the new structure.
[DEPEND] Dependency failed for Relabel all filesystems, if necessary.
[DEPEND] Dependency failed for Mark the need to relabel after reboot.
[DEPEND] Dependency failed for File System Check on /dev/pmem1.
```

Cause

The persistent devices do not initialize within the default auto-mount timeout value of 90 seconds. Systems with a large amount of persistent memory take a long time to boot because the initialization of the memory is serialized. Consequently, if there are persistent memory file systems listed in the `/etc/fstab` file, the system might time out while waiting for devices to become available. To work around this problem, configure the `DefaultTimeoutStartSec` option in the `/etc/systemd/system.conf` file to a sufficiently large value.

Solution 1

To work around the issue, increase the `DefaultTimeoutStartSec` value in the `/etc/systemd/system.conf` file to a sufficiently large value, such as 1,200 (seconds). Use a value of 0 seconds to wait indefinitely. This applies to all devices and mount points on the system.

Solution 2

Add `"x-systemd.device-timeout=0"` to the mount options for the PMem partition in `fstab`. This will control each mount point. Use a value of 0 to wait indefinitely, or use a suitably large value, such as 1,200 (seconds). For example:

```
/dev/pmem0 /mnt/pmem0 xfs defaults,x-systemd.device-timeout=0,dax 0 2
/dev/pmem1 /mnt/pmem1 xfs defaults,x-systemd.device-timeout=0,dax 0 2
/dev/pmem2 /mnt/pmem2 xfs defaults,x-systemd.device-timeout=0,dax 0 2
/dev/pmem3 /mnt/pmem3 xfs defaults,x-systemd.device-timeout=0,dax 0 2
```

Solution 3

Create a custom `systemd.mount` service with a suitable `x-systemd.device-timeout` value. Refer to the `systemd.mount(5)` man page for more information.

Solution 4

For systems running SUSE Linux Enterprise Server 12 SP4, configuring a large number of persistent memory devices can delay loading of the `btrfs` module.

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

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

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

Known Issue 2

Mounting an XFS file system with the `dax` option can return "wrong fs type, bad option, bad superblock, missing code page or helper program, or other error."

Applies To

- Red Hat Enterprise Linux and SUSE Linux Enterprise Server
- XFS file system
- `xfsprogs` 5.1 or later
- File systems mounted with the `dax` option (Direct Access)
- `mount(1)` command

Issue

After creating an XFS file system on an `fsdax` persistent memory namespace, mounting the file system returns the following error:

```
// Create an XFS file system
# mkfs.xfs /dev/pmem0
meta-data=/dev/pmem0          isize=512    agcount=4, agsize=97542016 blks
```

```

=                sectsz=4096  attr=2, projid32bit=1
=                crc=1        finobt=1, sparse=1, rmapbt=0
=                reflink=1
data            =                bsize=4096  blocks=390168064, imaxpct=5
=                sunit=0      swidth=0 blks
naming         =version 2      bsize=4096  ascii-ci=0, ftype=1
log            =internal log   bsize=4096  blocks=190511, version=2
=                sectsz=4096  sunit=1 blks, lazy-count=1
realtime       =none          extsz=4096  blocks=0, rtextents=0

// Mount the file system using the '-o dax' option
# mkdir /mnt/pmemfs0
# mount -o dax /dev/pmem0 /mnt/pmemfs0
mount: /pmemfs0: wrong fs type, bad option, bad superblock on /dev/pmem0, missing codepage or helper
program, or other error.

```

Note: If this issue is not noticed before the host is rebooted with entries for persistent memory devices in `/etc/fstab`, the boot will drop to emergency mode due to mount failures.

Cause

Starting with xfsprogs version 5.1, the default behavior is to create XFS file systems with the reflink option enabled. Previously, this option was disabled. The reflink and dax options are mutually exclusive, which causes the mount to fail. To check the installed version of xfsprogs, use:

```

RHEL> # yum list installed | grep xfsprogs
SUSE> # zipper search -i xfsprogs

```

An error in dmesg reporting "DAX and reflink cannot be used together!" can be seen when the mount fails:

```

# dmesg -T | tail
[758549.018746] pmem0: detected capacity change from 0 to 1598128390144
[758619.823070] XFS (pmem0): DAX enabled. Warning: EXPERIMENTAL, use at your own risk
[758619.823870] XFS (pmem0): DAX and reflink cannot be used together!

```

Solution

The EXT4 file system can be used as an alternative because it does not implement the reflink feature and supports DAX.

For XFS, explicitly disable the reflink feature at file-system-creation time if dax is required. Reflink cannot be disabled after the file system has been created.

For example:

```

# mkfs.xfs -m reflink=0 -f /dev/pmem0
meta-data=/dev/pmem0          isize=512    agcount=4, agsize=97542016 blks
=                sectsz=4096  attr=2, projid32bit=1
=                crc=1        finobt=1, sparse=1, rmapbt=0
=                reflink=0
data            =                bsize=4096  blocks=390168064, imaxpct=5
=                sunit=0      swidth=0 blks
naming         =version 2      bsize=4096  ascii-ci=0, ftype=1
log            =internal log   bsize=4096  blocks=190511, version=2
=                sectsz=4096  sunit=1 blks, lazy-count=1
realtime       =none          extsz=4096  blocks=0, rtextents=0

// Mount the file system with the 'dax' option
# mount -o dax /dev/pmem0 /pmemfs0

// Verify the mount was successful
# df -h /pmemfs0
Filesystem      Size  Used Avail Use% Mounted on
/dev/pmem0      1.5T  1.6G  1.5T   1% /pmemfs0

```

Intel and SAP: The Next Generation of In-Memory Database Performance

Intel Optane PMem changes the landscape of data storage by blurring the line between DRAM and persistent storage. SAP HANA 2.0 SPS 04 (or later) takes advantage of the capabilities of Intel Optane PMem by natively supporting memory persistence while benefitting from greater memory capacities.

This paper describes the configuration steps needed to configure a server running Intel Optane PMem and SAP HANA 2. As more persistent memory-compliant server systems become available from leading OEMs, this overview can help inform configuration of proof-of-concept systems that demonstrate the value of Intel Optane PMem and SAP HANA 2 for customers. For more information, visit intel.com/sap or sap.com/persistent-memory.



¹ Pricing guidance as of March 15, 2019. Scale-up configuration: eight-node SAP HANA 2 landscape. Per-node: eight-socket Intel Xeon Platinum 8276M processors. Memory capacity per socket: 6 x 128 GB DDR4 2,133 MHz and 6 x 128 GB Intel Optane persistent memory. Cost per terabyte on the scale-up system is \$24,682.

Scale-out configuration: 21-node SAP HANA 2 landscape. Per-node: four-socket Intel Xeon Platinum 8276 processor. Memory capacity per socket: 12 x 64 GB DDR4 2,133 MHz. Cost per terabyte on the scale-out system is \$44,991.

² Performance results are based on testing as of March 12, 2019 and may not reflect the publicly available security updates. Baseline: three-node SAP HANA 2 scale-out configuration (per node): 4 x Intel Xeon processor E7-8880 v3 (2.3 GHz, 150 W, 18 cores), CPU sockets: 4; RAM capacity: 64 x 32 GB DIMM, RAM model: DDR4 2,133 megabits per second (Mbps); storage: GPFS, approximately 21.8 TB of formatted local storage per node, SAN storage for backup space only; network: redundant 10 gigabit Ethernet (GbE) network for storage and access, redundant 10G network for node-to-node; OS: SUSE 12 SP2, SAP HANA: 2.00.035, GPFS: 4.2.3.10. Average time of 50 individual test queries executed 30–50 times each, for a total of approximately 25,000 steps: 2.81 seconds.

New configuration, one-node SAP HANA 2 scale-up configuration: CPU: 4 x Intel Xeon Platinum 8260 processor (2.2 GHz, 165 W, 24 cores), CPU sockets: 4; RAM capacity: 24 x 64 GB DIMM, RAM model: DDR4 2,133 Mbps; Intel Optane persistent memory: 24 x 126 GB PMM; storage: XFS, 21 TB; network: redundant 10 GbE network; OS: SUSE 15, SAP HANA: 2.00.035, Intel BKC: WW06. Average time of 50 individual test queries executed 30–50 times each, for a total of approximately 25,000 steps: 1.13 seconds.

³ Based on Intel testing as of March 1, 2019: base configuration: 10 systems with four-socket Intel Xeon processor E7-8894 v4, 768 GB (12 x 64 GB) memory. Compared to: five systems with four-socket Intel Xeon Platinum 8280L (28 cores), 2,304 GB (6 x 256 GB Intel Optane persistent memory + 6 x 128 GB DRAM, 2-2-2, App Direct Mode). Base system included \$35,592 on CPU, \$33,994 on memory, \$24,000 on storage, \$7,603 on RBOM, and \$0 on software, for a total of \$101,189 (or \$1,011,891 for 10 systems; \$67,459 per TB of storage). Comparison configuration included \$71,624 on CPU, \$123,163 on memory, \$54,000 on storage, \$7,603 on RDOM, and \$0 on software, for a total of \$256,390 (or \$1,281,950 for 10 systems; \$42,732 per TB storage).

⁴ The Rand Group. "How Much Does 1 Hour of Downtime Cost the Average Business?" randgroup.com/insights/cost-of-business-downtime/

⁵ SAP HANA simulated workload for SAP BW edition for SAP HANA Standard Application Benchmark Version 2 as of 30 May 2018. Baseline configuration with traditional DRAM: Lenovo ThinkSystem SR950 server with 8 x Intel Xeon Platinum 8176M processors (28 cores, 165 watt, 2.1 GHz). Total memory consists of 48 x 16 GB TruDDR4 2,666 MHz RDIMMs, and 5 x ThinkSystem 2.5" PM1633a 3.84 TB capacity SAS 12 Gb hot swap SSDs for SAP HANA storage. The operating system is SUSE Linux Enterprise Server (SLES) 12 SP3 and uses SAP HANA 2.0 SPS 03 with a 6 TB dataset. Start time: 50 minutes.

New configuration with a combination of DRAM and Intel Optane persistent memory: Lenovo ThinkSystem SR950 server with 8 x Intel Xeon Platinum 8176M processors (28 cores, 165 watt, 2.1 GHz). Total memory consists of 48 x 16 GB TruDDR4 2,666 MHz RDIMMs and 48 x 128 GB Intel Optane persistent memory modules (PMMs), and 5 x ThinkSystem 2.5" PM1633a 3.84 TB capacity SAS 12 Gb hot swap SSDs for SAP HANA storage. The operating system is SLES 12 SP3 and uses SAP HANA 2.0 SPS 03 with a 6 TB dataset. Start time: 4 minutes.

⁶ Refer to SAP Note 2618154—SAP HANA Persistent Memory—Release Information (<https://launchpad.support.sap.com/#/notes/2618154>, requires an SAP support account) for the latest information on supported versions.

Performance varies by use, configuration and other factors. Learn more at www.intel.com/PerformanceIndex.

Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See backup for configuration details. No product or component can be absolutely secure.

Your costs and results may vary.

Intel technologies may require enabled hardware, software or service activation.

No part of this publication may be reproduced or transmitted in any form or for any purpose without the express permission of SAP SE or an SAP affiliate company.

Some software products marketed by SAP SE and its distributors contain proprietary software components of other software vendors. National product specifications may vary.

These materials are provided by SAP SE or an SAP affiliate company for informational purposes only, without representation or warranty of any kind, and SAP SE or its affiliated companies shall not be liable for errors or omissions with respect to the materials. The only warranties for SAP SE or SAP affiliate company products and services are those that are set forth in the express warranty statements accompanying such products and services, if any. Nothing herein should be construed as constituting an additional warranty.

SAP and other SAP products and services mentioned herein as well as their respective logos are trademarks or registered trademarks of SAP SE (or an SAP affiliate company) in Germany and other countries. Please see <http://www.sap.com/trademark> for additional trademark information and notices.

© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.