

INTEL® OPTANE™ DC PERSISTENT MEMORY (DCPMM) – BASIC SETUP

NOTICES & DISCLAIMERS

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more complete information visit www.intel.com/benchmarks.

Intel processors of the same SKU may vary in frequency or power as a result of natural variability in the production process.

Optimization Notice: Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice. Notice Revision #20110804.

The benchmark results may need to be revised as additional testing is conducted. The results depend on the specific platform configurations and workloads utilized in the testing, and may not be applicable to any particular user's components, computer system or workloads. The results are not necessarily representative of other benchmarks and other benchmark results may show greater or lesser impact from mitigations.

Performance results are based on testing as of 4/22/2019 and may not reflect all publicly available security updates. See configuration disclosure for details. No product can be absolutely secure.

Results have been estimated or simulated using internal Intel analysis or architecture simulation or modeling, and provided to you for informational purposes. Any differences in your system hardware, software or configuration may affect your actual performance.

No computer system can be absolutely secure.

© 2019 Intel Corporation. Intel, the Intel logo, Xeon and Xeon logos are trademarks of Intel Corporation in the U.S. and/or other countries.

OBJECTIVES

Upon completion of this course, participants will be able to:

- Know which tools to use to setup up Intel Optane DC persistent memory (DCPMM)
- Do initial configuration of Intel Optane DC persistent memory in a system
- Run basic commands to validate the memory configuration

AGENDA

CONFIGURATION TOOLS

SETTING UP INTEL® OPTANE™ DC PERSISTENT MEMORY IN SIMPLE EASY STEPS

SETTING UP INTEL® OPTANE™ DC PERSISTENT MEMORY IN LINUX

SETTING UP INTEL® OPTANE™ DC PERSISTENT MEMORY IN WINDOWS SERVER

VALIDATE PERSISTENT MEMORY CONFIGURATION

SUMMARY

QUICK TOOL PRIMER

There are two configuration tools we will be mentioning in this presentation. The IPCMTL tool is a bit more versatile but NDCTL has its place as well

IPMCTL

IPMCTL is an open source utility created and maintained by Intel to manage Intel® Optane™ DC persistent memory modules. IPMCTL, which works in both Linux and Windows, is a vendor specific tool, meaning it is not able to be used for managing NVDIMMs from vendors other than Intel. The full project is open source and can be seen on GitHub, and downloaded from here:

<https://pkgs.org/download/ipmctl>

NDCTL

NDCTL is a utility for managing the Linux LIBNVDIMM Kernel subsystem. It is designed to work with various non-volatile memory devices (NVDIMMs) from different vendors. The LIBNVDIMM subsystem defines a kernel device model and control message interface for platform NVDIMM resources like those defined by the ACPI v6.0 NFIT (NVDIMM Firmware Interface Table). Most versions of Linux already contain the NDCTL package and it's just a matter of installing it. For directions, go to: <https://docs.pmem.io/persistent-memory/getting-started-guide/installing-ndctl>

SETTING UP INTEL® OPTANE™ DC PERSISTENT MEMORY IN SIMPLE EASY STEPS

Configure for Memory Mode (AKA MM, 2LM)

Install

- Install 2nd Generation Intel® Xeon® Scalable processor (Gold or Platinum tier or Silver 4215)
- Install DDR4 & Intel® Optane™ DC persistent memory in server chassis
- Power on the server & boot to OS

Use

- Check available memory (sum of installed DCPMM)
- Run your application(s)


Configure Intel® Optane™ DC persistent memory for App Direct (AKA Application Direct, AD, 1LM, Persistent Memory and PM)

Configure

- Create App Direct Regions (goal)
- Reboot the server (required to enable new configuration)
- Create App Direct Namespace
- Create filesystem on namespace device

Use

- Mount the persistent memory file system
- Configure application to use persistent memory filesystem
- Run your application(s)



Intel® Optane™ DC persistent memory devices boot in Memory Mode by default.


SETTING UP INTEL® OPTANE™ DC PERSISTENT MEMORY IN LINUX

Configure for Memory Mode (AKA MM, 2LM)

- Install**
 - Install 2nd Generation Intel® Xeon® Scalable processor (Gold or Platinum tier or Silver 4215)
 - Install DDR4 & Intel® Optane™ DC persistent memory in server chassis
- Use**
 - Power on the server
 - Check available memory (sum of installed Intel® Optane™ DC persistent memory)
 - Run your application(s)

Configure Intel® Optane™ DC persistent memory for App Direct (AKA Application Direct, AD, 1LM, Persistent Memory and PM)

- Configure**
 - `ipmctl create -goal persistentmemorytype=appdirect`
 - Reboot the server (required to enable new configuration)
 - `ndctl create-namespace`
 - `mkfs -t ext4 /dev/pmem0`
- Use**
 - `mkdir /pmem0; mount -o dax /dev/pmem0 /pmem0`
 - Configure application to use /pmem0 filesystem
 - Run your application(s)



Intel® Optane™ DC persistent memory devices boot in Memory Mode by default.

New goal requests take a few seconds and a reboot for processing.

```
[root@B7R19Z2S1 ~]# ndctl create-namespace
{
  "dev": "namespace1.0",
  "mode": "fsdax",
  "map": "dev",
  "size": "1488.37 GiB (1598.13 GB)",
  "uuid": "270bbb16-0a80-4f56-b6a2-0848c09cbf7d",
  "sector_size": 512,
  "align": 2097152,
  "blockdev": "pmem1"
}
```


SETTING UP INTEL® OPTANE™ DC PERSISTENT MEMORY IN WINDOWS SERVER 2019

Configure for Memory Mode (AKA MM, 2LM)

Install

- Install 2nd Generation Intel® Xeon® Scalable processor (Gold or Platinum tier or Silver 4215)
- Install DDR4 & Intel® Optane™ DC persistent memory in server chassis

Use

- Power on the server
- Check available memory (sum of installed Intel® Optane™ DC persistent memory)
- Run your application(s)


Configure Intel® Optane™ DC persistent memory for App Direct (AKA Application Direct, AD, 1LM, Persistent Memory and PM)

Configure

- `ipmctl create -goal persistentmemorytype=appdirect`
- Reboot the server (required to enable new configuration)
- `Get-PmemUnusedRegion | New-PmemDisk`
- Create filesystem on namespace device & mount as drive letter

Use

- Configure application to use persistent memory filesystem
- Run your application(s)



Intel® Optane™ DC persistent memory devices boot in Memory Mode by default

New goal requests take a few seconds and a reboot for processing.

Powershell Command (copy & paste)

```
Get-Disk |  
Where partitionstyle -eq 'raw' |  
Initialize-Disk -PartitionStyle GPT -PassThru |  
New-Partition -AssignDriveLetter -UseMaximumSize |  
Format-Volume -FileSystem NTFS -IsDAX:$true -  
Confirm:$false
```


VALIDATE PERSISTENT MEMORY CONFIGURATION

Basic Command Flow	Command Examples
Display system total DCPMM Memory Resources	<ul style="list-style-type: none">ipmctl show -memoryresources

```
[root@B7R19Z2S1 ~]# ipmctl show -memoryresources
Capacity=3029.4 GiB
MemoryCapacity=0.0 GiB
AppDirectCapacity=3024.0 GiB
UnconfiguredCapacity=0.0 GiB
InaccessibleCapacity=5.4 GiB
ReservedCapacity=0.0 GiB
```

Basic Command Flow	Command Examples
Show Persistent Memory	<ul style="list-style-type: none">ndctl list -Roripmctl show -region

```
[root@B7R19Z2S1 ~]# ndctl list -R
[
  {
    "dev": "region1",
    "size": 1623497637888,
    "available_size": 1623497637888,
    "max_available_extent": 1623497637888,
    "type": "pmem",
    "iset_id": 7882565198289448140,
    "persistence_domain": "memory_controller"
  },
  {
    "dev": "region0",
    "size": 1623497637888,
    "available_size": 1623497637888,
    "max_available_extent": 1623497637888,
    "type": "pmem",
    "iset_id": 5885781714394492108,
    "persistence_domain": "memory_controller"
  }
]
```

```
[root@B7R19Z2S1 ~]# ipmctl show -region

SocketID | ISetID | PersistentMemoryType | Capacity | FreeCapacity | HealthState
=====
0x0000 | 0x51ae7f48f7f92ccc | AppDirect | 1512.0 GiB | 1512.0 GiB | Healthy
0x0001 | 0x6d647f48c2e62ccc | AppDirect | 1512.0 GiB | 0.0 GiB | Healthy
```

SUMMARY

- Intel Optane DC persistent memory (DCPMM) ships by default in Memory Mode and requires no configuration
- To use DCPMM in Application Direct mode requires configuration. Two popular tools are IPMCTL and NDCTL (for Linux)

