

Enabling Optane Drives as IMDT

Jem Camba - 2018-12-09 - in Storage, Networking, & Databases

This article describes how to convert Intel Optane drives from block device mode to memory drive mode, functioning as memory for the operating system and its processes.

There are two methods described in this article:

- Automated install - for use with servers that are pre-configured and licensed by Packet
- Low-level install - for use on any server manually

Automated Install is much simpler, and should be used where possible.

Supported Operating System Distributions

- Ubuntu 16.04+
- CentOS 7+

Automated Install

To install:

1. Deploy and log onto the server using a supported operating system distribution.
2. Log onto the server
3. Run: `curl -v -L -O https://raw.githubusercontent.com/packethost/imdt-installer/master/imdt-deployer.sh`
4. Run `chmod +x imdt-deployer.sh`
5. Run `./imdt-deployer.sh`
6. Reboot

Errors

Possible errors are:

- The installer reports that the CPU does not support Intel Virtualization. If so, you cannot use IMDT on this server.
- The installer reports that no Optane drives are found. If so, you cannot use IMDT on this server.
- The installer reports that Intel Virtualization is not enabled. If so, you must boot into the BIOS, enable them, and reboot.

Low-Level Install

Prerequisites

1. In order to use Optane Drives as IMDT, you need a license from Intel for each Optane drive that you wish to convert. Contact your Intel representative for the necessary licenses. License keys normally are in the format **XXXX-XXXXXXX**
2. You must have a server with IMDT-capable Optane drives, like the ones in the Packet-Intel [AccelerateWithOptane lab](#).
3. As of this writing, IMDT works *only* with the list of supported operating system distributions
4. Each Optane drive used as memory must be associated with a different NUMA. Instructions on how to determine which NUMA each drive is associated with, and how to select drives, are in the "Installation" section.

Determining NUMA-NVMe Association

To determine which NVMe is associated with which NUMA, run the following:

```
# for i in /sys/module/nvme/drivers/*nvme*/nvme/nvm*; do dev=$(basename $i);  
start=${i#/sys/module/nvme/drivers/*nvme*/}; id=${start%/nvme/*}; numa=$(cat  
/sys/module/nvme/drivers/*nvme*/$id/numa*); echo $dev $numa; done
```

The following example is taken from one of our AccelerateWithOptane servers:

```
# for i in /sys/module/nvme/drivers/*nvme*/nvme/nvm*; do dev=$(basename $i);  
start=${i#/sys/module/nvme/drivers/*nvme*/}; id=${start%/nvme/*}; numa=$(cat  
/sys/module/nvme/drivers/*nvme*/$id/numa*); echo $dev $numa; done
```

```
nvme0 0  
nvme1 0  
nvme2 0  
nvme3 1
```

The nvme drives `nvme0` , `nvme1` and `nvme2` all are associated with NUMA `0` , while `nvme3` is associated with NUMA `1` . The first nvme drive actually is NAND, not Optane, and isn't a valid candidate, leaving you with the 3 drives `nvme1` , `nvme2` , and `nvme3` . If you want just one drive to be used as memory, use any one of them. If you want 2 drives, pick one of `nvme1` or `nvme2`, associated with NUMA `0` , and `nvme3` , associated with NUMA `1`

Installation

1. Deploy a server with a supported operating system, per the normal Packet procedure.
2. Log onto the server
3. Download the IMDT installer from <http://memorydrv.com/downloads/> As of this writing, the latest release is `8.6.2535.19` , so you would download it as `curl -O -L http://memorydrv.com/downloads/8.6.2535.19/imdt_installer-8.6.2535.19.sh`
4. Make the installer executable as `chmod +x imdt_installer*`
5. Determine which Optane drives you want to convert to memory by discovering which drives are associated with which NUMA; see above.
6. Run the installer you downloaded previously to generate an `imdt.list` and request a license file: `./imdt_installer-8.6.2535.19.sh in -n` . When requested, enter the appropriate license keys retrieved earlier, and your email address.
7. Wait to receive the license keys by email. The license keys are sent as an attachment to the email and must be downloaded *in toto*, not copied or modified in any way.
8. Upload the license file to the server.
9. Run the installer again, pointing it to the license file uploaded in the previous step: `./imdt_installer-8.6.2535.19.sh in -n <license_file>` . Accept all defaults.
10. Boot the system into BIOS, and modify the BIOS settings per the next section.
11. Reboot.
12. When the system comes up, look for the startup message `Loading Intel Memory Drive Technology` , followed by an IMDT splash screen, and eventually the normal kernel boot.
13. Check that you are running IMDT by running `imdtversion` .

Required BIOS Settings

The required BIOS settings are as follows. Paths to settings tend to be similar between various firmware, but may not be identical on your actual server.

- `Intel Virtualization Technology` must be **enabled**. It normally can be found at `BIOS | Advanced | Process Configuration | Intel Virtualization`

Technology

- **Intel Virtualization Technology for Directed I/O** must be **enabled**. It normally can be found at **BIOS | Advanced | Integrated I/O Configuration | Intel VT for Directed I/O**
- The first boot device must be one of the Optane drives. They normally show up as **UEFI Misc Device** and **UEFI Misc Device 2** , but may vary based on your system. Move these to the top of the list.

When changing the boot order, it is **critical** that you do not change the ordering of any of the UEFI IPv4 boot options relative to each other. These can cause problems for deprovisioning.

Tags

Optane