



Intel® Volume Management Device (Intel® VMD) Enabled NVMe Driver for VMware* ESXi* Version 2.7

Release Notes

May 2021
Version 1.0



Legal Notices and Disclaimers

Intel may make changes to specifications and product descriptions at any time, without notice. Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined". Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them. The information here is subject to change without notice. Do not finalize a design with this information.

Intel disclaims all express and implied warranties, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and non-infringement, as well as any warranty arising from course of performance, course of dealing, or usage in trade.

Benchmark results were obtained prior to implementation of recent software patches and firmware updates intended to address exploits referred to as "Spectre" and "Meltdown". Implementation of these updates may make these results inapplicable to your device or system.

Tests document performance of components on a particular test, in specific systems. Differences in hardware, software, or configuration will affect actual performance. Consult other sources of information to evaluate performance as you consider your purchase. Results have been estimated based on internal Intel analysis and are provided for informational purposes only. Any difference in system hardware or software design or configuration may affect actual performance.

All documented performance test results are obtained in compliance with JESD218 Standards; refer to individual sub-sections within this document for specific methodologies. See www.jedec.org for detailed definitions of JESD218 Standards. Intel does not control or audit the design or implementation of third party benchmark data or Web sites referenced in this document. Intel encourages all of its customers to visit the referenced Web sites or others where similar performance benchmark data are reported and confirm whether the referenced benchmark data are accurate and reflect performance of systems available for purchase.

The products described in this document may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

For copies of this document, documents that are referenced within, or other Intel literature please contact your Intel representative.

Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order.

*Other names and brands may be claimed as the property of others. Copyright © 2021 Intel Corporation. All rights reserved.



Driver Revision History

Date	Driver Version	Description
May 2021	<ul style="list-style-type: none">Intel VMD release 2.7 for VMWare® ESXi® 6.5/6.7/7.0	<ul style="list-style-type: none">Intel VMD-enabled NVMe driver for Generation 1,2,3 Intel® Xeon Scalable Processors



Contents

1	Introduction.....	5
1.1	Overview.....	5
1.2	New in This Release.....	5
1.3	Defect Submission Process.....	6
2	Support.....	7
2.1	Supported Operating Systems.....	7
2.2	Supported Platforms.....	7
2.3	Supported Configurations.....	7
2.4	Supported PCIe NVMe SSD.....	7
3	Release Package Contents.....	8
3.1	Hypervisor Components.....	8
3.2	UEFI Drivers necessary to be loaded into the BIOS with this release**.....	8
3.3	Intel VMD-enabled LED Management.....	8
3.4	Other documents included with this release.....	9
3.5	Limitations.....	9
3.5.1	Upgrading VMD 2.x version.....	9
3.5.2	Workaround When DataStore1 Disappears.....	9
3.5.3	Reboot Required After Migrating RAID 1 From Pass Through.....	12
3.5.4	Writing Core Dump to Disk Not Supported in this release on RAID 1.....	12
3.5.5	Certain NVMe Switches Cause VMD LED Status on Other Slots to be OFF During Hotplug.....	12
3.5.6	vSphere Hotplug Insertion Event Tab Limitation.....	13
3.5.7	Intel® Virtual RAID on CPU (Intel® VROC) HII Menu in Pass Thru mode.....	13
4	Known Issues in this release.....	14
5	Resolved Issues.....	15

1 Introduction

1.1 Overview

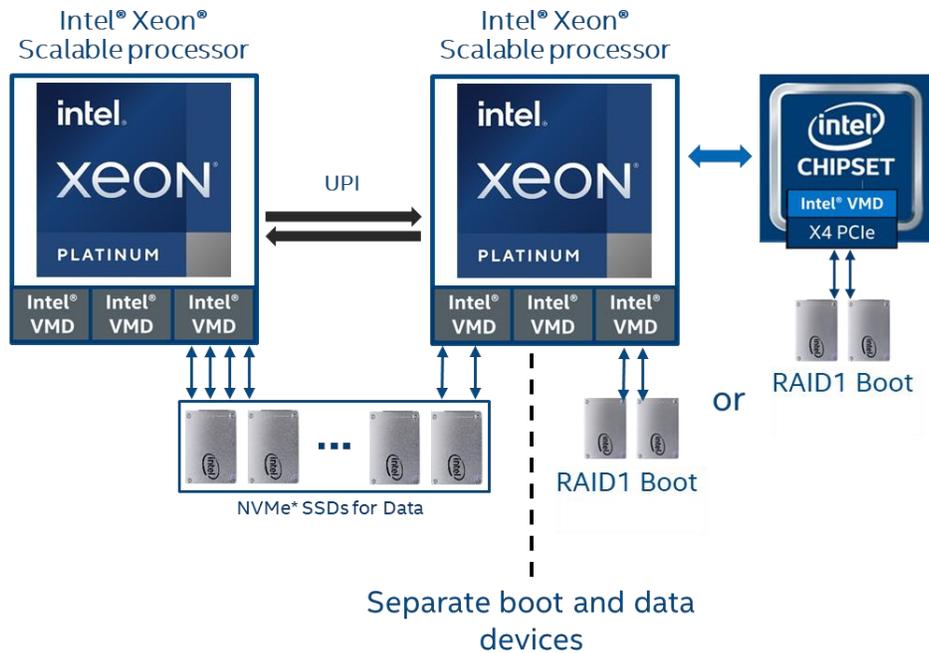
The Intel® Volume Management Device (Intel® VMD) – Enabled NVMe Driver release for VMWare® ESXi® package contains the 2.7 build release version of the Intel VMD-Enabled NVMe Driver to support Generation 1,2, 3 Intel® Xeon Scalable Processors platforms using VMWare® ESXi® Hypervisor. This build is specifically for Generation 3 Intel® Xeon Scalable Processors platforms, and is backward compatible to Intel Generation 1 and 2 Intel® Xeon Scalable Processor platforms (non-PCH RAID 1).

Intel VMD-enabled NVMe driver assists in the management of CPU and PCH attached PCIe NVMe SSDs (Generation 3 Intel® Xeon Scalable Processors). Features include ability for PCIe NVMe Surprise Hot Plug, LED Management, and Error handling optimal in a VMWare® ESXi® VSAN environment.

1.2 New in This Release

Intel VMD NVMe driver version 2.7 supports the Generation 3 Intel® Xeon Scalable Processors Platform for devices behind VMD off of the PCH, and ability to create RAID 1 off of CPU or PCH attached NVMe. Only a single (1) RAID1 volume is supported per platform, and specifically for boot image purposes. In addition, users must be certain that the boot volume (either a single device or 2 Disk RAID1 volume), is on a dedicated Intel VMD Domain/Controller. This means that 1) the datastores cannot be placed on the boot volume and 2) boot device(s) and the data devices cannot be attached to the same Intel VMD Domain. See below image as an example.

*Note: for Intel VMD Direct Assign feature, the NVMe enabled on VMD root ports cannot be direct bassigned to the data store on other VMD enabled NVMe.





For Generation 1 and 2 Intel® Xeon Scalable Processors platforms, RAID 1 can be created off of the CPU attached NVMe.

New in Intel VMD 2.7, Only VMD enabled NVMe devices can be managed with this driver version. New LED Management tools and syntax must also be used with the 2.7 version.

Note: the name of the Intel VMD driver has been updated from “intel-nvme-vmd” to “iavmd”.

Note: When VMD is disabled, the native VMware NVMe driver will load on SSD PCIe NVMe devices.

1.3 Defect Submission Process

With this release, Intel will accept, and process issues reported by customers via the Intel premier Support (IPS) portal.

To submit an issue, please use the Intel Premier Support (IPS) tool. Information, training and details can be found starting at the below website. Your local Intel FAE can also provide you the necessary requirements to enable you to submit an IPS issue (also known as a “case”) including an account setup if you do not already have one.

<http://www.intel.com/content/www/us/en/design/support/ips/training/welcome.html>

2 Support

2.1 Supported Operating Systems

- VMWare® ESXi® version 6.5*/6.7/7.0/7.0U1
- Only VMWare® ESXi® 6.7U3P4 and 7.0U1 version will be supported by VMWare on Intel Generation 3 Intel® Xeon Scalable Processor platforms.
- VMWare® ESXi® 6.5/6.7/7.0/7.0U1 will be supported on Intel Skylake Platforms.

2.2 Supported Platforms

- **Intel® Xeon® Scalable platforms**
 - Intel® C620 series chipset

2.3 Supported Configurations

- Up to 2 level deep switch
- Up to 48 PCIe NVMe SSDs

2.4 Supported PCIe NVMe SSD

Intel VMD-enabled NVMe driver supports most shipping enterprise Intel NVMe SSDs as well as 3rd Party NVMe SSDs. * Please see your Intel AE for the most current list of supported / validated 3rd party PCIe NVMe SSDs.

**Intel NVMe Dual controller devices not supported in this release (DC P3608)*

3 Release Package Contents

3.1 Hypervisor Components

Please note that VMWare will only support ESXi version 6.7U3P4 and 7.0U1 on Intel Generation 3 Intel® Xeon Scalable Processor platforms.

All driver releases of ESXi* below will be acceptable on Purley Skylake Platforms (CPU attached RAID 1).

VMware* ESXi* 6.5 Intel VMD NVMe Driver
iavmd-2.7.0.1157-1OEM.650.0.0.4598673.x86_64.vib
VMD LED Management tool for ESXi* 6.5
intel-vmdr-user-6.5-2.7.2173.vib

VMware* ESXi* 6.7 Intel VMD NVMe Driver
iavmd-2.7.0.1157-1OEM.670.0.0.8169922.x86_64.vib
VMD LED Management tool for ESXi* 6.7
intel-vmdr-user-6.7-2.7.2173.vib

VMware* ESXi* 7.0/7.0U1 Intel VMD NVMe Driver
iavmd-2.7.0.1157-1OEM.700.1.0.15843807.x86_64.vib
VMD LED Management tool for ESXi* 7.0/7.0U1
intel-vmdr-user-7.0-2.7.2173.vib

3.2 UEFI Drivers necessary to be loaded into the BIOS with this release**

Intel® VROC UEFI 7.5 drivers used for validation with this package:

VMDVROC_1.efi

VMDVROC_2.efi

RCmpVROC.efi – Utility for verifying system compliance

**Both UEFI drivers are required to perform enumeration and exposure of Intel VMD-enabled attached devices in the pre-boot environment

3.3 Intel VMD-enabled LED Management

New LED Management command line tool included in this release:

intel-vmdr-user-x.x-2.7.2173.vib

Instructions for using this LED tool are found in the User Guide included in this package.

3.4 Other documents included with this release

Intel_VMD_NVMe_VMWare_User_Guide_2_6.pdf

3.5 Limitations

3.5.1 ESXi 7.0U1 hotplug shows DEAD state in ESXi log

During hotplug insertion and removal, a DEAD state may be seen in ESXi logs when sending the **esxcfg-mpath -L** command after hotplug. This log entry does not affect hotplug or device functionality. The log entrees can be cleared manually as a workaround by sending the following command after hotplug:

```
esxcli storage core adapter rescan -a
```

3.5.2 Upgrading VMD 2.x version

If upgrading an existing VMD 1.x driver to version 2.x, the system module parameters must be disabled for Driver log verbosity level. This is only when the hypervisor is installed on SATA/USB or devices not behind VMD.

Before upgrading the driver, check if the logLevel is set in the current driver module parameters by entering the following command:

```
esxcli system module parameters list -m iavmd
```

If the loglevel is set, you will see it under "Value". Above it is circled in white and is set at logLevel "5".

*If there is no number here, under "value", and is blank, you can proceed to upgrade the driver.

If there is a number here, you must first disable this Driver Log verbosity by sending the following command with null quotations:

```
esxcfg-module -set-options "" iavmd
```

You can now proceed to upgrade the VMD driver to the next 2.6 version

3.5.3 Workaround When DataStore1 Disappears

There are two situations where the datastore1 may disappear and the following workaround will resolve the issue.

1. After Upgrading from VMware NVMe inbox Driver to Intel VMD Driver
2. If RAID 1 is degraded and OS is migrating to a spare or newly inserted device

3.5.3.1 Workaround Steps

Warning: If installing VMware ESXi Hypervisor on an NVMe SSD, datastore1 may be inaccessible after upgrading to Intel VMD NVMe Driver or migration of a RAID 1 boot volume. Please follow one of the workarounds listed in this section to correct this issue.

Note: A reboot is required after installing the iavmd driver.

There are two existing issues that may cause a previously created datastore to disappear when updating from VMware inbox NVMe driver to an Intel NVMe driver or to the Intel VMD driver.

Issue #1:

Data stores save SCSI Inquiry VPD 80h (serial number) into their filesystem metadata. In order to use a data store, the VPD 80h data needs to match what is stored in the metadata.

The SCSI-to-NVMe translation reference provides 3 possible ways to create a VPD 80h response and Intel VMD driver translation is different than VMware's translation.

Intel always uses a v1.0 translation for VPD 80h. This is compatible with data stores created with drivers that pre-date EUI-64 (all Intel NVMe drivers), and compatible with Intel NVMe drives that have EUI-64 collisions.

However, the Intel VMD driver translation is incompatible with data stores created with a driver using a EUI-64 translation (VMware*).

The VMware* translation uses EUI-64 if the device supports it, then v1.0 if it does not. This translation is incompatible with data stores created with drivers that pre-date EUI-64, and incompatible in the presence of drives with EUI-64 collisions.

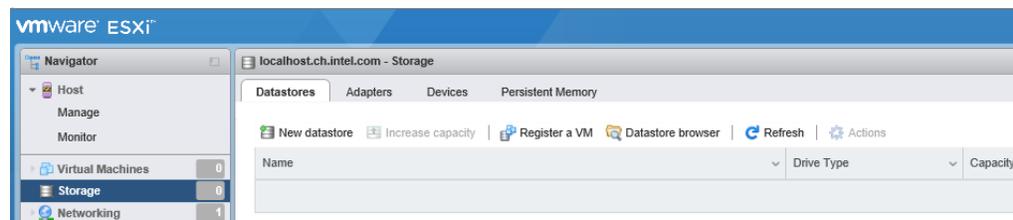
Issue #1 Workaround: Use Re-Signature Process to locate Data Stores created previously by VMware* ESXi* NVMe driver

Please use the following example for Re-signature process. Instructions can be found at the following links from VMware and KB article for Re-signature process:

<https://docs.vmware.com/en/VMware-vSphere/6.5/com.vmware.vsphere.storage.doc/GUID-EEFEB765-A41F-4B6D-917C-BB9ABB80FC80.html>

<https://kb.vmware.com/s/article/1011387>

Example of Data Stores previously created are missing:



b

Steps to re-signature process:

At command prompt in ESXi shell: Type: **esxcli storage vmfs snapshot list**

```
[root@localhost:~] esxcli storage vmfs snapshot list
5aac198a-455a5658-cc98-a0369fa162f0
  Volume Name: DataStore2
  VMFS UUID: 5aac198a-455a5658-cc98-a0369fa162f0
  Can mount: true
  Reason for un-mountability:
  Can resignature: true
  Reason for non-resignaturability:
  Unresolved Extent Count: 1

5aac19e5-acc5004-9605-a0369fa162f0
  Volume Name: DataStore4
  VMFS UUID: 5aac19e5-acc5004-9605-a0369fa162f0
  Can mount: true
  Reason for un-mountability:
  Can resignature: true
  Reason for non-resignaturability:
  Unresolved Extent Count: 1

5aac2abc-fa70525e-1e1a-a4bf013952b3
  Volume Name: datastore1
  VMFS UUID: 5aac2abc-fa70525e-1e1a-a4bf013952b3
```

Two Examples using the either the label or VMFS UUID

Example of Re-signature the snapshot/replica LUN by using uuid:

```
# esxcli storage vmfs snapshot resignature -u 5aac2abc-fa70525e-1e1a-a4bf013952b3
```

Example of Re-signature using the label "datastore1"

```
# esxcli storage vmfs snapshot resignature -l "datastore1"
```

Output should state, "No unresolved VMFS snapshots with volume label / uuid found."

```
5aac2abc-fa70525e-1e1a-a4bf013952b3
  Volume Name: datastore1
  VMFS UUID: 5aac2abc-fa70525e-1e1a-a4bf013952b3
  Can mount: true
  Reason for un-mountability:
  Can resignature: true
  Reason for non-resignaturability:
  Unresolved Extent Count: 1

[root@localhost:~] esxcli storage vmfs snapshot resignature -l "datastore1"
No unresolved VMFS snapshots with volume label 'datastore1' found.
[root@localhost:~]
```

If the output does not appear as above, enter the command again.

Rescan the vSphere Web Client to see that the datastore shows up again.

If this fails to restore the datastore, proceed to Issue #2 Workaround.

Issue #2: SCSI Layer Translation for VMWare* ESXi*

VMware ESXi NVMe SCSI layer translation reports a default logical block size for all NVMe devices, and stores the physical block size within the filesystem metadata of the data store. The

Intel VMD NVMe driver will report the logical block size from the NVMe device. This incompatibility between the two physical block sizes reported will cause datastore1 to disappear and re-signature process outlined above will not work without the following additional workaround.

Issue #2 Workaround: Use Module Parameters SNT_COMPAT

In order to correct this issue, the module parameters can be set right after vib installation. Install the Intel VMD driver and reboot. The driver's scsi-to-nvme translation compatibility mode can be set to **off** by running the following commands:

*First, set the system into maintenance mode any time the module parameters command is changed

```
esxcli system maintenanceMode set -e true
```

```
esxcli system module parameters set -iavmd -p "SNT_COMPAT=0"
```

Reboot the system after setting this parameter. Setting to off will have the NVMe SCSI layer use VMware's translations and keep compatibility between driver updates.

Note, this is a module wide parameter: this cannot be used to select the translation on individual drives.

3.5.4 Immediate Reboot Required After Migrating a System Device From Pass Through to RAID 1 using CLI tool in Hypervisor

A reboot is required if using the CLI tool in the hypervisor environment to migrate a bootable pass through (non-RAID) ESXi hypervisor drive to a second drive, to create a RAID 1 volume. DO NOT perform any IO to the RAID 1 volume before rebooting.

3.5.5 Writing Core Dump to Disk Not Supported in this release on RAID 1

This feature will be supported in a future release.

3.5.6 Certain NVMe Switches Cause VMD LED Status on Other Slots to be OFF During Hotplug

Certain switches may cause the LED Status Locate blinking state to go to OFF when hot plugging other NVMe devices attached to the same switch. This issue cannot be reproduced on other switches Intel VMD has validated, however, it may occur on those not tested.



3.5.7 vSphere Hotplug Insertion Event Tab Limitation

When hot removing a VMD-enabled NVMe SSD, and hot inserting the exact same drive, vSphere events/monitor reports a warning and the device as inaccessible. The Event tab does not report that the drive is re-inserted. This is not specific to NVMe, but, occurs with SATA devices as well. Please refer to the Devices tab that shows when the device is re-inserted, or the VMKernel.log to validate that the hot reinserted NVMe device is correctly logged.

3.5.8 Intel® Virtual RAID on CPU (Intel® VROC) HII Menu in Pass Thru mode

Intel VMD and Intel VROC UEFI drivers are packaged together. Intel VMD UEFI driver enumerates and assigns resources for all NVMe devices under the root port. The Intel VROC UEFI driver exposes those devices to the system.

Due to this packaging, the devices in the UEFI HII Bios menu will be found under the Intel® Virtual RAID on CPU (Intel® VROC) HII menu when Intel VMD is enabled.

Intel VROC in Pass Thru mode is seen so that the user knows that NVMe RAID is not supported when Intel VROC is in pass thru mode.

4 Known Issues in this release

18013603803 - VMD 2.6 – ESXi7.0 U1 – WilsonCity – Datastore is missing after reboot after migration

(see Section 3.5.2 for re-signature workaround)

18013580769 - VMD 2.6 - ESXi7.0 U1 – WilsonCity – Datastore is not visible on degraded OS RAID volume after reboot

(see Section 3.5.2 for Workaround)

18013500569 - VMD 2.6 - ESXi 6.7/7.0 - WilsonCity - Migration does not start when disk without partitions is used as migration source. Partitions are assumed to be GPT-based.

VMDRCLI returns an error and migration does not start

"ERROR: VOLUME CREATE STATUS ERROR (INVALID DATA DISK SIZE)"

Workaround: Do not use the "s" parameter when starting the RAID migration

i.e.

1. Make sure that source migration drive is "clean" - that it does not have partitions

"partedUtil get /dev/disks/t10.NVMe____Intel_SSD....."

2. Try to start RAID migration

/opt/intel/bin/intel-vmdr-user createvol vmhba4 -d 0,1 -l 1 -b 64 -c 352000 -n Test ~~-s 1~~



5 Resolved Issues

Issues Fixed in this release of Intel VMD-enabled NVMe Driver for VMWare* ESXi*

IPS # 00531193 - Whitely - ESXi7.0 NVMe removal doesn't work

18013580789 - VMD 2.6 - ESXi7.0 U1 - WilsonCity - RAID volume disks are in INCOMPATIBLE state after marking SPARE disk in HII with degraded RAID

18013579720 - VMD 2.6 - ESXi 6.7 - can't boot to OS after mark as PT on volume member where OS is installed (under rebuild to spare)

14014167028 - VMD 2.6 - ESXi7.0 U1 - Queue depth set at default value causes performance requirements failure during vSAN certification.