



NVMesh Interoperability Matrix

2.5.2 — Last update: 15 December 2022

NVIDIA - Mellanox

Table of Contents

1. Copyright and Trademark Information	3
2. Preface	4
3. Hardware	5
3.1. Servers	6
3.2. Network Interface Cards (NICs).....	7
3.2.1. Deprecated	9
3.3. Physical Storage Devices and Media	10
3.3.1. NVMe Devices	11
3.3.1.1. Caveats and Notes	14
3.3.2. Non-NVMe Devices.....	15
3.4. Metadata Region Support.....	16
3.5. Network Switches.....	17
4. Software	18
4.1. Operating Systems.....	19
4.1.1. Red Hat Enterprise Linux and OpenShift.....	20
4.1.2. Ubuntu	23
4.2. Browsers.....	25
4.3. Software Dependencies	26
4.4. File Systems	27

1. Copyright and Trademark Information

© 2015-2022 Excelero, Inc. All rights reserved. Specifications are subject to change without notice. Excelero, the Excelero logo, MeshProtect, and Remote-Direct-Drive-Access (RDDA) are trademarks of Excelero, Inc. in the United States and/or other countries. NVMesh® is a registered trademark of Excelero, Inc. in the United States.

Mellanox and ConnectX are registered trademarks of NVIDIA.

Intel is a registered trademark of Intel Corporation. Xeon and Core are trademarks of Intel Corporation. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

Red Hat, Red Hat Enterprise Linux, the Red Hat logo and OpenShift are trademarks of Red Hat, Inc., registered in the United States and other countries.

Linux® is the registered trademark of Linus Torvalds in the United States and other countries.

XFS® is a trademark of Silicon Graphics International Corp. or its subsidiaries in the United States and/or other countries.

Ubuntu® is a trademark of Canonical or its subsidiaries in the United States and/or other countries.

Node.js® is an official trademark of Joyent. NVMesh is not formally related to or endorsed by the official Joyent Node.js open source or commercial project

OpenStack® is an official trademark of the OpenStack Foundation.

MONGO and MONGODB are trademarks of MongoDB Inc., registered in the United States and other countries.

Microsoft® and Azure® are trademarks of Microsoft, Inc., registered in the United States and other countries.

2. Preface

Excelero™ creates innovative, high performance storage solutions that accelerate business applications and deliver outstanding return on investment with the lowest cost of ownership. The NVMesh® software defined block storage product offers Elastic NVMe: the performance of local NVMe server flash with the convenience, efficiency and redundancy of an all-flash-array. For details, go to: www.excelero.com.

 This document details hardware and software interoperability with **NVMesh 2.5.2**.

AUDIENCE

The primary audience for this document is intended to be storage and/or application administration personnel responsible for installing and deploying the Excelero **NVMesh 2.5.2** product.

NON-DISCLOSURE REQUIREMENTS

© Copyright 2015-2022 Excelero, Inc. All rights reserved. This document contains the confidential and proprietary information of Excelero, Inc. Do not reproduce or distribute without the prior written consent of Excelero.

FEEDBACK

We continually try to improve the quality and usefulness of Excelero documentation. If you have any corrections, feedback, or requests for additional documentation, send an e-mail message to support@excelero.com

3. Hardware

This section details the supported and validated hardware components for **NVMesh 2.5.2**.

- [Servers](#)
- [Network Interface Cards](#)
- [NVMe Devices](#)
- [Metadata Region Support](#)
- [Network Switches](#)

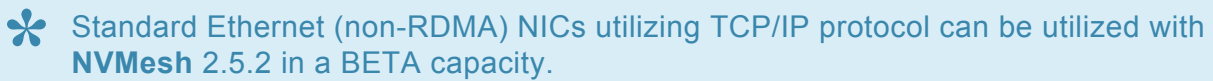
3.1. Servers

NVMesh 2.5.2 runs on standard Intel® and AMD® x86 servers with NVMe drives and RDMA NICs.

Note: AMD EPYC (Naples and Rome) based platforms are currently certified. There are caveats in terms of using RDDA. Consult with [Excelero Technical Support](#) for more information.

3.2. Network Interface Cards (NICs)

Supported Standard Ethernet NICs



In general, there is no certification for TCP/IP NICs. It is suggested to use 10Gb or faster NICs. Best results are achieved with 25Gb or faster NICs.

Supported RDMA NICs

The following RDMA NICs have been tested and certified as supported with **NVMesh**. NICs are listed in alphabetical order by manufacturer.

Supported Protocols:

- RoCEv2 = RDMA over Converged Ethernet, version 2
- InfiniBand

Vendor	Model or Series	Specific Part Number(s)	Supported Protocols	Firmware Level	Notes
Mellanox	ConnectX-4 VPI	MCX455A-ECAT, MCX456A-ECAT, MCX455A-FCAT, MCX456A-FCAT, MCX453A-FCAT, MCX454A-FCAT	RoCEv2, InfiniBand	12.20.1010	
Mellanox	ConnectX-4 EN	MCX415A-CCAT, MCX416A-CCAT, MCX415A-BCAT, MCX416A-BCAT, MCX415A-GCAT, MCX416A-GCAT, MCX413A-BCAT, MCX414A-BCAT, MCX413A-GCAT, MCX414A-GCAT	RoCEv2	12.20.1010	
Mellanox	ConnectX-4 Lx EN	MCX4111A-XCAT, MCX4121A-XCAT, MCX4111A-ACAT, MCX4121A-ACAT, MCX4131A-BCAT, MCX4131A-GCAT	RoCEv2	14.20.1010	
Mellanox	ConnectX-5 VPI	MCX555A-ECAT, MCX556A-ECAT, MCX556M-ECAT	RoCEv2, InfiniBand	16.20.1010	
Mellanox	ConnectX-5 EN	MCX515A-CCAT, MCX516A-CCAT	RoCEv2	16.20.1010	
Mellanox	ConnectX-6	Various	RoCEv2, InfiniBand	20.26.1040	See important notes below.
Mellanox	ConnectX-6 DX	Various	RoCEv2,	22.27.1016	See

			InfiniBand	important notes below.
--	--	--	------------	------------------------

Notes:

- The Mellanox Connect-IB adapter line is *not* supported.
- If you have interest in support for the Broadcom NetXtreme® E-Series 10Gb/25Gb RDMA Ethernet Controller, BCM57414, please contact Excelero.

! ConnectX-6 Family: The firmware level listed above is the minimum supported version.
RDDA protocol: If you enable **RDDA** in `nvmesh.conf` you **MUST** enable the `ONE_QP_PER_RECOVERY=1` firmware setting in the network adapter.
With earlier ConnectX cards, if you did not set this, it resulted in poor performance.
With the ConnectX-6, not enabling this can result in data corruption.

3.2.1. Deprecated

The following Network Interface Cards (NICs) were supported in previous versions of **NVMesh** but are no longer supported for new installations.

Customers with a valid support contract and installations with these devices will continue to be supported for the duration of that support contract.

Support for these devices may be completely removed in future versions of **NVMesh**.

Vendor	Model or Series	Specific Part Number(s)	Supported Protocols
ATTO	FastFrame™ 3 Series	N311, N312, N351, N325, N332	RoCEv2
Mellanox®	ConnectX-3 VPI Pro	MCX353A-FCCT, MCX354A-FCCT	RoCEv2, InfiniBand
Mellanox	ConnectX-3 EN Pro	MCX311A-XCCT, MCX312B-XCCT, MCX312C-XCCT, MCX313A-BCCT, MCX314A-BCCT	RoCEv2
Mellanox	ConnectX-3 VPI	MCX353A-QCBT, MCX354A-QCBT, MCX353A-FCBT, MCX354A-FCBT	InfiniBand

3.3. Physical Storage Devices and Media

3.3.1. NVMe Devices

Erasure Coded (EC) Volumes Support

The current implementation of Erasure Coded Volumes requires NVMe drives to support a 4K + 8B sector size. This is sometimes referred to as “metadata support”, “long blocks”, “VSS” or as “end-to-end data protection”. In the table below, the “EC Certified” column denotes whether a given drive can be used with such volumes. All certified drives can be used for other volume types.

To identify if a drive **not** listed below could potentially run the current Erasure Coding functionality, refer to [Metadata Region Support](#).

NVMe Drives Certified For Use With NVMesh

Drives are listed in alphabetical order by manufacturer.

Manufacturer	Model/ Series	EC Certified	Form Factor	Notes
Dell	Express Flash PM1725a	Yes	U.2	
Intel	DC P4800X Optane™	No	U.2, AIC	See Caveats and Notes .
	DC P5800X Optane™	No	U.2	See Caveats and Notes .
Kioxia	CD5	No	U.2	
	CM5	Yes	U.2, AIC	Initial formatting for drives of size 12 TB or higher requires additional configuration. Contact Exceero support for details.
	CD6	No	U.2	
	CM6	Yes	U.2	Initial formatting for drives of size 12 TB or higher requires additional configuration. Contact Exceero support for details.
Micron	7100 Series	No	U.2, AIC	
	7300 Series	No	U.2, M.2	
	7400	No	U.3, M.2, E1.S	

	Series			
	9100 Series	No	U.2, AIC	
	9200 Series	No	U.2	
	9300 Series	No	U.2	See Caveats and Notes .
Netlist	N1951 Series	Yes	U.2, AIC	
	N1962 Series	Yes	U.2	
Samsung	PM953	No	U.2	
	PM963	No	U.2, M.2	
	PM983	No	U.2, NF1	
	PM9A3	No	E1.S, U.2, M.2	
	PM1725a	Yes	U.2, AIC	
	PM1725b	Yes	U.2, AIC	
	PM1733	Yes	U.2	It is recommended to initially precondition such drives by writing twice to the entire drive. This improves read performance.
	SM963	No	U.2	
	SZ985	No	AIC	
	XS1715	No	U.2, AIC	
Solidigm	DC P432X Series	No	U.2, E1.L	
	DC P4420 Series	No	U.2	
	DC P4500 Series	No	U.2, AIC	
	DC P4600 Series	No	U.2, AIC	
	DC P4510	No	U.2	

	Series			
	DC P4610 Series	No	U.2	
	D5-P5316 Series	Yes	U.2, E1.L	While the drive model is EC-certified, this is an unrecommended usage of such drives.
	D7-P5500 Series	Yes	U.2	
	D7-P5510 Series	Yes	U.2	
	D7-P5600 Series	Yes	U.2	
WD	Ultrastar SN200 Series	Yes	U.2, AIC	
	Ultrastar SN620	No	U.2, AIC	
	Ultrastar SN630	No	U.2, AIC	
	Ultrastar SN640	No	U.2, AIC	
	Ultrastar SN840	Yes	U.2	

Form Factor Definitions:

- U.2 = SFF 2.5" Drive
- AIC = Add-in-card (PCIe edge connector card)
- M.2 = PCIe Mini Card
- NF1 = Samsung NF1 sometimes referred to as M.3
- E1.L = EDSFF 1U Long
- E1.S = EDSFF 1U Short

3.3.1.1. Caveats and Notes

Drive Specific Notes

Intel Optane P4800X & P5800X

To use such drives with **NVMesh 2.5.2**, perform these steps prior to formatting the drives, as they do not support TRIM operations.

On all **Targets** with such drives:

1. Use the `nvme list` command to obtain the model name of the drives. This may require installing the `nvme-cli` package.
2. Run `/opt/NVMesh/common-repo/tools/toma_rpc disk-models set <modelname> i s_using_nvme_trim_before_zero off` with the model name obtained from the previous command. If there are multiple models of Optane drives, repeat this command for each model.

Micron 9300

Moving from 512b format to 4k format for these drives requires resetting them. To avoid issues with drives not reverting from reset after format, it is recommended to format the drives to 4k first using the `nvme` tool from the `nvme-cli` package with the reset flag set (-r). If they do not revert, it will be necessary to power-cycle the machines.

Deprecated Drives

The following NVMe drives were supported in previous versions of **NVMesh**, but are no longer supported for new installations.

Customers with a valid support contract and installations with these devices will continue to be supported for the duration of that support contract.

Support for these devices may be completely removed in future versions of **NVMesh**.

Manufacturer	Model/Series	Form Factor	Notes
Intel®	DC P3600/3700 Series	U.2, AIC	
WD (HGST)	Ultrastar® SN100 Series	U.2, AIC	

3.3.2. Non-NVMe Devices

Non-NVMe devices can be used as storage media.

See [Add non-NVMe Drives](#) for more detail and instructions.

! It is not possible to use virtual block devices connected with NVMe-over-Fabrics, as they are incorrectly identified as NVMe drives.

3.4. Metadata Region Support

MeshProtect Distributed Parity volumes can only be created on drives formatted with 4K + 8B sectors. This format requires drives that support an NVMe feature named *Metadata Region*. This is sometimes referred to as “metadata support”, “long blocks”, or as “End-to-end data protection”.

To identify support for metadata region, use the Linux nvme command:

```
nvme id-ns
```

In the output, search for the `mc` value:

```
mc          : 0x3
```

The possible settings are:

<pre>m c : 0</pre>	Metadata Region not supported.
<pre>m c : 0x1</pre>	Metadata supported. Metadata as part of extended data LBA is supported on the drive. This mode is not support for NVMesh 2.5.2 .
<pre>m c : 0x2</pre>	Metadata supported. Metadata referred to with a metadata pointer is supported on the drive. The metadata pointer is the preferred mode for NVMesh 2.5.2 parity-based volumes.
<pre>m c : 0x3</pre>	Metadata supported. Both metadata options are supported on the drive. This mode is support for NVMesh 2.5.2

It is also important to validate that the drive supports 4K blocks + 8 bytes metadata. In the output, search for an `lbaf` entry with `ms:8` (8 bytes metadata) and `lbads:12` ($2^{12} = 4096$ bytes data).

```
lbaf 3 : ms:8 lbads:12 rp:0x2
```


3.5. Network Switches

The following network switches are known to work with **NVMesh 2.5.2**:

Vendor/ Model	Switch OS	Version	Notes
Arista	EOS		
Cisco NX-OS		7.0(3)I4(1)	
Cumulus Networks	Cumulus Linux	3.6.x / 4.2.x	Can only be configured with ECN. Global pause or PFC mode are not supported, except on Mellanox switches
Dell Force10 (z9100)	FTOS		
Mellanox	MLNX-OS/ Onyx	3.6.6107	
Supermicro C3632	Cumulus Linux	3.5.x	Can only be configured with ECN. Global pause or PFC mode are not supported

4. Software

This section details the supported and validated software components for **NVMesh 2.5.2**.

4.1. Operating Systems

The following Linux distributions, with distinct combinations of kernel versions and versions of OpenFabrics Enterprise Distribution (OFED), have been tested and certified as supported with **NVMesh 2.5.2**.

- [Red Hat Linux / CentOS](#)
- [Ubuntu Server](#)

TCP Support is available only for some select kernels, which are specifically highlighted in the following sections.

4.1.1. Red Hat Enterprise Linux and OpenShift

The following Red Hat® Enterprise Linux® and CentOS distributions, with distinct combinations of kernel versions and versions of OpenFabrics Enterprise Distribution (OFED), have been tested and certified as supported with **NVMesh 2.5.2**. For other similar distributions, contact [Excelero Technical Support](#).

Red Hat Linux 8.5

✿ TCP-only is not supported yet (i.e. **Targets** and **Clients** must have an RDMA NIC to enable TCP)

Mellanox OFEDs	Kernels
Inbox driver (no OFED) 5.4-3.1.0.0	4.18.0-348.el8.x86_64 4.18.0-348.2.1.el8_5.x86_64 4.18.0-348.7.1.el8_5.x86_64
5.4-3.1.0.0	4.18.0-348.20.1.el8_5.x86_64

Red Hat Linux 8.4 and OpenShift

✿ TCP-only is not supported yet (i.e. **Targets** and **Clients** must have an RDMA NIC to enable TCP)

Mellanox OFEDs	Kernels
Inbox driver (no OFED) 5.4-1.0.3.0	4.18.0-305.el8.x86_64 4.18.0-305.7.1.el8.x86_64 4.18.0-305.10.2.el8_4.x86_64 4.18.0-305.12.1.el8_4.x86_64 4.18.0-305.17.1.el8_4.x86_64


Red Hat Linux 8.3

Mellanox OFEDs	Kernels
Inbox driver (no OFED) 5.4-1.0.3.0 5.2-2.2.0.0	4.18.0-240.el8.x86_64 4.18.0-240.1.1.el8.x86_64 4.18.0-240.10.1.el8.x86_64 4.18.0-240.15.1.el8.x86_64 4.18.0-240.22.1.el8.x86_64

Red Hat Linux 8.2

Mellanox OFEDs	Kernels
	4.18.0-193.el8.x86_64
Inbox driver (no OFED)	4.18.0-193.6.3.el8_2.x86_64
5.4-1.0.3.0	4.18.0-193.13.2.el8_2.x86_64
5.2-2.2.0.0	4.18.0-193.14.3.el8_2.x86_64
5.0-2.1.8.0	4.18.0-193.19.1.el8_2.x86_64
	4.18.0-193.28.1.el8_2.x86_64

Red Hat Linux / CentOS 7.9

 TCP-only is supported only with 5.4.x kernels (i.e. **Targets** and **Clients** using 3.10.0-1160.x kernels must have an RDMA NIC to enable TCP)

Mellanox OFEDs	Kernels
	3.10.0-1160.el7.x86_64
	3.10.0-1160.6.1.el7.x86_64
	3.10.0-1160.11.1.el7.x86_64
Inbox driver (no OFED)	3.10.0-1160.15.2.el7.x86_64
5.4-1.0.3.0	3.10.0-1160.45.1.el7.x86_64
5.2-2.2.0.0	3.10.0-1160.49.1.el7.x86_64
5.1-2.3.7.1	3.10.0-1160.53.1.el7.x86_64
	3.10.0-1160.62.1.el7.x86_64
	5.4.80
	5.4.111
	5.10.41

Red Hat Linux / CentOS 7.8

Mellanox OFEDs	Kernels
Inbox driver (no OFED)	3.10.0-1127.el7.x86_64
5.4-1.0.3.0	3.10.0-1127.8.2.el7.x86_64
5.2-2.2.0.0	3.10.0-1127.10.1.el7.x86_64
5.0-2.1.8.0	3.10.0-1127.13.1.el7.x86_64
	3.10.0-1127.18.2.el7.x86_64

Red Hat Linux / CentOS 7.7

Mellanox OFEDs	Kernels
Inbox driver (no OFED)	3.10.0-1062.el7.x86_64
5.4-1.0.3.0	3.10.0-1062.9.1.el7.x86_64
5.2-2.2.0.0	3.10.0-1062.12.1.el7.x86_64

5.0-2.1.8.0	3.10.0-1062.18.1.el7.x86_64
-------------	-----------------------------

4.1.2. Ubuntu

The following Ubuntu Server distributions, with distinct combinations of kernel versions and versions of OpenFabrics Enterprise Distribution (OFED), have been tested and certified as supported with **NVMesh 2.5.2**.

NVMf access functionality requires manual installation of python 3.6. Contact [Excelero Technical Support](#) for instructions.

Ubuntu 20.04

Mellanox OFEDs	Kernels
5.4-1.0.3.0	5.4.0-66 5.4.0-80 5.4.0-81 5.4.0-99 5.4.0-100 (1)
5.1-2.5.8.0	5.4.0-66
5.1-2.3.7.1	5.4.0-33 5.4.0-40 5.4.0-52 5.4.0-60 5.4.0-62 5.4.0-64 5.4.0-65 5.4.0-66
Inbox driver (no OFED) (1)	5.4.0-100
TCP-only (1)	5.4.0-100

Ubuntu 18.04

Mellanox OFEDs	Kernels
5.1-2.3.7.1	4.15.0-122 4.15.0-123 4.15.0-124 4.15.0-128 4.15.0-129 4.15.0-130 4.15.0-132 4.15.0-134 4.15.0-135 4.15.0-136
5.4-1.0.3.0 (1)	5.4.0-100

Inbox driver (no OFED) (1)	5.4.0-100 4.15.0-136
TCP-only (1)	5.4.0-100 4.15.0-136

Note Id	Note Description
1	Added with build 2.5.1-5

4.2. Browsers

The **NVMesh 2.5.2 Management** GUI is tested for quality assurance with Google Chrome across Linux, MacOS and Windows. Other browsers may work fine, but if you experience unexpected behavior in the web GUI, you may be asked to try the same operation with Chrome.

4.3. Software Dependencies

The **NVMesh 2.5.2 Management** modules requires certain versions of packages. The following versions have been tested and certified as supported.

Supporting Software	Information	Required Version
MongoDB	https://docs.mongodb.com/v4.2/	4.2
NodeJS	https://nodejs.org/en/about/releases/	12.x

4.4. File Systems

The following standard Linux file systems are known to work with **NVMesh 2.5.2**.

- ext3
- ext4
- xfs

The following clustered or distributed Linux file systems are known to work with **NVMesh 2.5.2**.

- ThinkParq BeeGFS
- IBM SpectrumScale (GPFS)
- HPE CXFS
- DDN Lustre
- Quantum StorNext
- [OCFS2](#) with limitations. Contact [Excelero Technical Support](#) for more details.