



ESXi

SAN Host

NetApp
May 25, 2022

Table of Contents

ESXi 1

 NVMe-oF Host Configuration for ESXi 7.x with ONTAP 1

ESXi

NVMe-oF Host Configuration for ESXi 7.x with ONTAP

Supportability

NVMe over Fibre Channel (NVMe/FC) is supported on ONTAP 9.7 or later.

Features

- ESXi initiator host can run both NVMe/FC and FCP traffic through the same adapter ports. See the [Hardware Universe](#) for a list of supported FC adapters and controllers. See the [NetApp Interoperability Matrix](#) for the most current list of supported configurations and versions.
- Beginning with ONTAP 9.9.1 P3, NVMe/FC feature is supported for ESXi 7.0 update 3.
- For ESXi 7.0 and later releases, HPP (high performance plugin) is the default plugin for NVMe devices.

Known limitations

The following configurations are not supported:

- RDM mapping
- VVols

Enable NVMe/FC

1. Check the ESXi host NQN string and verify that it matches with the host NQN string for the corresponding subsystem on the ONTAP array:

```
# esxcli nvme info get
Host NQN: nqn.2014-08.com.vmware:nvme:nvme-esx

# vservers nvme subsystem host show -vservers vservers_nvme
Vserver Subsystem          Host NQN
-----
vservers_nvme ss_vservers_nvme nqn.2014-08.com.vmware:nvme:nvme-esx
```

Configure Broadcom/Emulex

1. Check whether the configuration is supported with required driver/firmware by referring to [NetApp Interoperability Matrix](#).
2. Set the `lpfc` driver parameter `lpfc_enable_fc4_type=3` for enabling NVMe/FC support in the `lpfc` driver and reboot the host.



Starting with vSphere 7.0 update 3, the `brcmnvme_fc` driver is no longer available. Therefore, the `lpfc` driver now includes the NVMe over Fibre Channel (NVMe/FC) functionality previously delivered with the `brcmnvme_fc` driver.



The `lpfc_enable_fc4_type=3` parameter is set by default for the LPe35000-series adapters. You must perform the following command to set it manually for LPe32000-series and LPe31000-series adapters.

```
# esxcli system module parameters set -m lpfc -p lpfc_enable_fc4_type=3

#esxcli system module parameters list -m lpfc | grep lpfc_enable_fc4_type
lpfc_enable_fc4_type          int      3          Defines what FC4 types
are supported

#esxcli storage core adapter list
HBA Name  Driver  Link State  UID
Capabilities  Description
-----  -
vmhba1    lpfc    link-up     fc.200000109b95456f:100000109b95456f
Second Level Lun ID (0000:86:00.0) Emulex Corporation Emulex LPe36000
Fibre Channel Adapter  FC HBA
vmhba2    lpfc    link-up     fc.200000109b954570:100000109b954570
Second Level Lun ID (0000:86:00.1) Emulex Corporation Emulex LPe36000
Fibre Channel Adapter  FC HBA
vmhba64   lpfc    link-up     fc.200000109b95456f:100000109b95456f
(0000:86:00.0) Emulex Corporation Emulex LPe36000 Fibre Channel Adapter
NVMe HBA
vmhba65   lpfc    link-up     fc.200000109b954570:100000109b954570
(0000:86:00.1) Emulex Corporation Emulex LPe36000 Fibre Channel Adapter
NVMe HBA
```

Configure Marvell/QLogic

1. Check whether configuration is supported with required driver/firmware by referring to [NetApp Interoperability Matrix](#).
2. Set the `qlnativefc` driver parameter `ql2xnvmesupport=1` for enabling NVMe/FC support in the `qlnativefc` driver and reboot the host.

```
# esxcfg-module -s 'ql2xnvmesupport=1' qlnativefc
```



The `qlnativefc` driver parameter is set by default for the Qle 277x-series adapters. You must perform the following command to set it manually for Qle 277x series adapters.

```
esxcfg-module -l | grep qlnativefc
qlnativefc          4      1912
```

3. Check whether `nvme` is enabled on the adapter:

```

#esxcli storage core adapter list
HBA Name   Driver      Link State  UID
Capabilities      Description
-----
-----
vmhba3     qlnativefc  link-up     fc.20000024ff1817ae:21000024ff1817ae
Second Level Lun ID (0000:5e:00.0) QLogic Corp QLE2742 Dual Port 32Gb
Fibre Channel to PCIe Adapter FC Adapter
vmhba4     qlnativefc  link-up     fc.20000024ff1817af:21000024ff1817af
Second Level Lun ID (0000:5e:00.1) QLogic Corp QLE2742 Dual Port 32Gb
Fibre Channel to PCIe Adapter FC Adapter
vmhba64    qlnativefc  link-up     fc.20000024ff1817ae:21000024ff1817ae
(0000:5e:00.0) QLogic Corp QLE2742 Dual Port 32Gb Fibre Channel to PCIe
Adapter NVMe FC Adapter
vmhba65    qlnativefc  link-up     fc.20000024ff1817af:21000024ff1817af
(0000:5e:00.1) QLogic Corp QLE2742 Dual Port 32Gb Fibre Channel to PCIe
Adapter NVMe FC Adapter

```

Validate NVMe/FC

1. Verify that NVMe/FC adapter is listed on the ESXi host:

```

# esxcli nvme adapter list

Adapter  Adapter Qualified Name      Transport Type  Driver
Associated Devices
-----
-----
vmhba64  aqn:qlnativefc:21000024ff1817ae  FC              qlnativefc
vmhba65  aqn:qlnativefc:21000024ff1817af  FC              qlnativefc
vmhba66  aqn:lpfc:100000109b579d9c        FC              lpfc
vmhba67  aqn:lpfc:100000109b579d9d        FC              lpfc

```

2. Verify that the NVMe/FC namespaces are properly created:

The UUIDs in the following example represent the NVMe/FC namespace devices.

```
# esxcfg-mpath -b
uuid.5084e29a6bb24fbca5ba076eda8ecd7e : NVMe Fibre Channel Disk
(uuid.5084e29a6bb24fbca5ba076eda8ecd7e)
  vmhba65:C0:T0:L1 LUN:1 state:active fc Adapter: WWNN:
20:00:34:80:0d:6d:72:69 WWPN: 21:00:34:80:0d:6d:72:69 Target: WWNN:
20:17:00:a0:98:df:e3:d1 WWPN: 20:2f:00:a0:98:df:e3:d1
  vmhba65:C0:T1:L1 LUN:1 state:active fc Adapter: WWNN:
20:00:34:80:0d:6d:72:69 WWPN: 21:00:34:80:0d:6d:72:69 Target: WWNN:
20:17:00:a0:98:df:e3:d1 WWPN: 20:1a:00:a0:98:df:e3:d1
  vmhba64:C0:T0:L1 LUN:1 state:active fc Adapter: WWNN:
20:00:34:80:0d:6d:72:68 WWPN: 21:00:34:80:0d:6d:72:68 Target: WWNN:
20:17:00:a0:98:df:e3:d1 WWPN: 20:18:00:a0:98:df:e3:d1
  vmhba64:C0:T1:L1 LUN:1 state:active fc Adapter: WWNN:
20:00:34:80:0d:6d:72:68 WWPN: 21:00:34:80:0d:6d:72:68 Target: WWNN:
20:17:00:a0:98:df:e3:d1 WWPN: 20:19:00:a0:98:df:e3:d1
```



In ONTAP 9.7, the default block size for a NVMe/FC namespace is 4K. This default size is not compatible with ESXi. Therefore, when creating namespaces for ESXi, you must set the namespace block size as 512b. You can do this using the `vserver nvme namespace create` command.

Example

```
vserver nvme namespace create -vserver vs_1 -path /vol/nsvol/namespacel -size
100g -ostype vmware -block-size 512B
```

Refer to the [ONTAP 9 Command man pages](#) for additional details.

3. Verify the status of the individual ANA paths of the respective NVMe/FC namespace devices:

```

esxcli storage hpp path list -d uuid.5084e29a6bb24fbca5ba076eda8ecd7e
fc.200034800d6d7268:210034800d6d7268-
fc.201700a098dfe3d1:201800a098dfe3d1-
uuid.5084e29a6bb24fbca5ba076eda8ecd7e
  Runtime Name: vmhba64:C0:T0:L1
  Device: uuid.5084e29a6bb24fbca5ba076eda8ecd7e
  Device Display Name: NVMe Fibre Channel Disk
(uuid.5084e29a6bb24fbca5ba076eda8ecd7e)
  Path State: active
  Path Config: {TPG_id=0,TPG_state=AO,RTP_id=0,health=UP}

fc.200034800d6d7269:210034800d6d7269-
fc.201700a098dfe3d1:201a00a098dfe3d1-
uuid.5084e29a6bb24fbca5ba076eda8ecd7e
  Runtime Name: vmhba65:C0:T1:L1
  Device: uuid.5084e29a6bb24fbca5ba076eda8ecd7e
  Device Display Name: NVMe Fibre Channel Disk
(uuid.5084e29a6bb24fbca5ba076eda8ecd7e)
  Path State: active
  Path Config: {TPG_id=0,TPG_state=AO,RTP_id=0,health=UP}

fc.200034800d6d7269:210034800d6d7269-
fc.201700a098dfe3d1:202f00a098dfe3d1-
uuid.5084e29a6bb24fbca5ba076eda8ecd7e
  Runtime Name: vmhba65:C0:T0:L1
  Device: uuid.5084e29a6bb24fbca5ba076eda8ecd7e
  Device Display Name: NVMe Fibre Channel Disk
(uuid.5084e29a6bb24fbca5ba076eda8ecd7e)
  Path State: active unoptimized
  Path Config: {TPG_id=0,TPG_state=ANO,RTP_id=0,health=UP}

fc.200034800d6d7268:210034800d6d7268-
fc.201700a098dfe3d1:201900a098dfe3d1-
uuid.5084e29a6bb24fbca5ba076eda8ecd7e
  Runtime Name: vmhba64:C0:T1:L1
  Device: uuid.5084e29a6bb24fbca5ba076eda8ecd7e
  Device Display Name: NVMe Fibre Channel Disk
(uuid.5084e29a6bb24fbca5ba076eda8ecd7e)
  Path State: active unoptimized
  Path Config: {TPG_id=0,TPG_state=ANO,RTP_id=0,health=UP}

```

Known issue

- ESXi 7.0 U3 (and later) NVMe/FC support is available starting ONTAP 9.9.1 P3 onwards. This is due to key NVMe abort (issued by ESXi 7.0 U3 and later) fixes that is available starting ONTAP 9.9.1 P3 only. Refer to

the respective burt public report at <https://mysupport.netapp.com/site/bugs-online/product/ONTAP/BURT/1420654> for details.

Related Links

[TR-4597-VMware vSphere with ONTAP](#)

[VMware vSphere 5.x, 6.x and 7.x support with NetApp MetroCluster \(2031038\)](#)

[VMware vSphere 6.x and 7.x support with NetApp® SnapMirror® Business Continuity \(SM-BC\)](#)

Copyright Information

Copyright © 2022 NetApp, Inc. All rights reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means-graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system-without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark Information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.