



NVMe/FC Host Configuration for ESXi 7.0 with ONTAP

ONTAP SAN Host

Amanda Stroman
July 30, 2020

Table of Contents

- NVMe/FC Host Configuration for ESXi 7.0 with ONTAP 1
 - Supportability 1
 - Known limitations 1
 - Enabling NVMe/FC with ANA 1
 - Configuring the Broadcom FC adapter for NVMe/FC 2
 - Validating NVMe/FC 4

NVMe/FC Host Configuration for ESXi 7.0 with ONTAP

Supportability

NVMe/FC is supported on ONTAP 9.7 or later for ESXi 7.0.

ESXi initiator host can run both NVMe/FC & FCP traffic through the same adapter ports. See the [Hardware Universe](#) for a list of supported FC adapters and controllers. For the most current list of supported configurations & versions, see the [NetApp Interoperability Matrix](#).

Known limitations

The following are not supported:

- RDM mapping
- VVols

Enabling NVMe/FC with ANA

1. Disable the HppManageDegradedPaths parameter for improved interoperability with ONTAP:

```
# esxcfg-advcfg -s 0 /Misc/HppManageDegradedPaths
```

2. Reboot the host.
3. After reboot, verify that the HppManageDegradedPaths parameter is now disabled:

```
# esxcfg-advcfg -g /Misc/HppManageDegradedPaths  
Value of HppManageDegradedPaths is 0
```

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

Example

```
# esxcli nvme info get
Host NQN: nqn.2014-08.com.vmware:nvme:chat-54-113

*> vserver nvme subsystem host show -vserver co_nv_fc_esx
Vserver Subsystem Host NQN
-----
-----
co_nv_fc_esx
      subsys_chat_54_113_nvme
              nqn.2014-08.com.vmware:nvme:chat-54-113
```

Configuring the Broadcom FC adapter for NVMe/FC

1. Install the recommended lpfc driver by copying it to a temporary folder and then executing the following command:

```
# esxcli software vib install -d /tmp/t/Emulex-FCoE-FC-lpfc-12.4.224.0-
offline-bundle-13621872.zip --no-sig-check
Installation Result
  Message: The update completed successfully, but the system needs to
be rebooted for the changes to be effective.
  Reboot Required: true
  VIBs Installed: EMU_bootbank_lpfc_12.4.224.0-1OEM.688.0.0.13621872
  VIBs Removed: EMU_bootbank_lpfc_12.4.211.6-1OEM.688.0.0.13621872
  VIBs Skipped:
```

2. If necessary, set the lpfc driver parameter `lpfc_enable_fc4_type=3` for enabling NVMe/FC support in the lpfc driver:



This parameter is set by default for the LPe35000-series adapters. You must perform the following step to set it manually for LPe32000-series & LPe31000-series adapters.

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

3. Use the `elxmgmt` utility to upgrade the Broadcom FC adapter firmware to the recommended version:

```
# esxcli software vib install -d /tmp/t/Emulex-elxmgmt-6.8.7-12.4.211.7.zip --no-sig-check
Installation Result
  Message: The update completed successfully, but the system needs to be rebooted for the changes to be effective.
  Reboot Required: true
  VIBs Installed: EMU_bootbank_emu-esx-elxmgmt_12.4.211.7-01
  VIBs Removed:
  VIBs Skipped:
...
```

4. Reboot the host.
5. After reboot, verify that the recommended lpfc driver and adapter firmware versions have applied and the initiator ports are online:

```
# esxcli storage san fc list
Adapter: vmhba3
  Port ID: 010600
  Node Name: 20:00:00:90:fa:e0:ec:8e
  Port Name: 10:00:00:90:fa:e0:ec:8e
  Speed: 32 Gbps
  Port Type: NPort
  Port State: ONLINE
  Model Description: Emulex LightPulse LPe32002-M2 2-Port 32Gb Fibre
Channel Adapter
  Hardware Version: 0000000c
  OptionROM Version: 12.4.217.2
  Firmware Version: 12.4.217.2
  Driver Name: lpfc
  DriverVersion: 12.4.224.0

Adapter: vmhba4
  Port ID: 010F00
  Node Name: 20:00:00:90:fa:e0:ec:8f
  Port Name: 10:00:00:90:fa:e0:ec:8f
  Speed: 32 Gbps
  Port Type: NPort
  Port State: ONLINE
  Model Description: Emulex LightPulse LPe32002-M2 2-Port 32Gb Fibre
Channel Adapter
  Hardware Version: 0000000c
  OptionROM Version: 12.4.217.2
  Firmware Version: 12.4.217.2
  Driver Name: lpfc
  DriverVersion: 12.4.224.0
```

Validating NVMe/FC

1. Verify that the ONTAP target NVMe/FC controllers are properly discovered on the ESXi host:

```
# esxcli nvme controller list
```

```
Name
```

```
Controller Number  Adapter  Transport Type  Is Online
```

```
-----  
-----  
-----  
nqn.1992-  
08.com.netapp:sn.e7f89c2c245d11e9975300a098dfce55:subsystem.interop_57_v  
m_01#vmhba32#204900a098dfe3d1:204a00a098dfe3d1                259  
vmhba32  FC                false  
nqn.1992-  
08.com.netapp:sn.e7f89c2c245d11e9975300a098dfce55:subsystem.interop_57_v  
m_09#vmhba32#204900a098dfe3d1:204a00a098dfe3d1                263  
vmhba32  FC                false  
nqn.1992-  
08.com.netapp:sn.e7f89c2c245d11e9975300a098dfce55:subsystem.interop_57_v  
m_11#vmhba32#204900a098dfe3d1:204a00a098dfe3d1                267  
vmhba32  FC                false  
nqn.1992-  
08.com.netapp:sn.e7f89c2c245d11e9975300a098dfce55:subsystem.interop_57_v  
m_10#vmhba32#204900a098dfe3d1:204a00a098dfe3d1                265  
vmhba32  FC                false  
nqn.1992-  
08.com.netapp:sn.e7f89c2c245d11e9975300a098dfce55:subsystem.interop_57_v  
m_02#vmhba32#204900a098dfe3d1:204a00a098dfe3d1                261  
vmhba32  FC                false
```

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.0d12b7cd97344be8a53b7913f8f72f04 : NVMe Fibre Channel Disk
(uuid.0d12b7cd97344be8a53b7913f8f72f04)
  vmhba65:C0:T9:L30 LUN:30 state:active fc Adapter: WWNN:
20:00:00:90:fa:e0:ec:8f WWPN: 10:00:00:90:fa:e0:ec:8f Target: WWNN:
20:49:00:a0:98:df:e3:d1 WWPN: 20:4d:00:a0:98:df:e3:d1
  vmhba64:C0:T9:L30 LUN:30 state:active fc Adapter: WWNN:
20:00:00:90:fa:e0:ec:8e WWPN: 10:00:00:90:fa:e0:ec:8e Target: WWNN:
20:49:00:a0:98:df:e3:d1 WWPN: 20:4c:00:a0:98:df:e3:d1
  vmhba64:C0:T5:L30 LUN:30 state:standby fc Adapter: WWNN:
20:00:00:90:fa:e0:ec:8e WWPN: 10:00:00:90:fa:e0:ec:8e Target: WWNN:
20:49:00:a0:98:df:e3:d1 WWPN: 20:4a:00:a0:98:df:e3:d1
  vmhba65:C0:T0:L30 LUN:30 state:standby fc Adapter: WWNN:
20:00:00:90:fa:e0:ec:8f WWPN: 10:00:00:90:fa:e0:ec:8f Target: WWNN:
20:49:00:a0:98:df:e3:d1 WWPN: 20:4b:00:a0:98:df:e3:d1

uuid.49de7683950d47c9898f51443d893910 : NVMe Fibre Channel Disk
(uuid.49de7683950d47c9898f51443d893910)
  vmhba65:C0:T12:L39 LUN:39 state:active fc Adapter: WWNN:
20:00:00:90:fa:e0:ec:8f WWPN: 10:00:00:90:fa:e0:ec:8f Target: WWNN:
20:3a:00:a0:98:df:e3:d1 WWPN: 20:27:00:a0:98:df:e3:d1
  vmhba65:C0:T13:L39 LUN:39 state:standby fc Adapter: WWNN:
20:00:00:90:fa:e0:ec:8f WWPN: 10:00:00:90:fa:e0:ec:8f Target: WWNN:
20:3a:00:a0:98:df:e3:d1 WWPN: 20:29:00:a0:98:df:e3:d1
  vmhba64:C0:T12:L39 LUN:39 state:active fc Adapter: WWNN:
20:00:00:90:fa:e0:ec:8e WWPN: 10:00:00:90:fa:e0:ec:8e Target: WWNN:
20:3a:00:a0:98:df:e3:d1 WWPN: 20:3b:00:a0:98:df:e3:d1
  vmhba64:C0:T13:L39 LUN:39 state:standby fc Adapter: WWNN:
20:00:00:90:fa:e0:ec:8e WWPN: 10:00:00:90:fa:e0:ec:8e Target: WWNN:
20:3a:00:a0:98:df:e3:d1 WWPN: 20:28:00: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 512b. You can do this using the `vserver nvme namespace create` command.

Example

```
vserver nvme namespace create -vserver vs_1 -path /vol/nsvol/namespace1 -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
```



```
fc.20000090fae0ec8f:10000090fae0ec8f-
fc.204900a098dfe3d1:204d00a098dfe3d1-
uuid.1aa669c5376240a28ae47d8d549586ea
  Runtime Name: vmhba65:C0:T9:L33
  Device: uuid.1aa669c5376240a28ae47d8d549586ea
  Device Display Name: NVMe Fibre Channel Disk
(uuid.1aa669c5376240a28ae47d8d549586ea)
  Path State: active
```

```
fc.20000090fae0ec8e:10000090fae0ec8e-
fc.204900a098dfe3d1:204a00a098dfe3d1-
uuid.1aa669c5376240a28ae47d8d549586ea
  Runtime Name: vmhba64:C0:T5:L33
  Device: uuid.1aa669c5376240a28ae47d8d549586ea
  Device Display Name: NVMe Fibre Channel Disk
(uuid.1aa669c5376240a28ae47d8d549586ea)
  Path State: standby
```

```
:leveloffset: -1
```

```
<<<
```

```
*Copyright Information*
```

Copyright © 2021 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.