



ESXi

SAN Host

NetApp
May 13, 2022

目录

ESXi	1
适用于采用 ONTAP 的 ESXi 7.x 的 NVMe-oF 主机配置	1

ESXi

适用于采用 ONTAP 的 ESXi 7.x 的 NVMe-oF 主机配置

可支持性

ONTAP 9.7 或更高版本支持基于光纤通道的 NVMe（NVMe/FC）。

功能

- ESXi 启动程序主机可以通过相同的适配器端口同时运行 NVMe/FC 和 FCP 流量。请参见 "[Hardware Universe](#)" 有关支持的 FC 适配器和控制器的列表。请参见 "[NetApp 互操作性表](#)" 有关支持的最新配置和版本列表。
- 从 ONTAP 9.9.1 P3 开始，ESXi 7.0 Update 3 支持 NVMe/FC 功能。
- 对于 ESXi 7.0 及更高版本，HPP（高性能插件）是 NVMe 设备的默认插件。

已知限制

不支持以下配置：

- RDM 映射
- VVOL

启用 NVMe/FC

1. 检查 ESXi 主机 NQN 字符串，并验证它是否与 ONTAP 阵列上相应子系统的主机 NQN 字符串匹配：

```
# 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
```

配置 Broadcom/Emulex

1. 请参见，检查所需驱动程序 / 固件是否支持此配置 "[NetApp 互操作性表](#)"。
2. 设置 lpfc 驱动程序参数 lpfc_enable_FC4_type=3 以在 lpfc 驱动程序中启用 NVMe/FC 支持，然后重新启动主机。



从 vSphere 7.0 Update 3 开始，brcmnvme_fc 驱动程序不再可用。因此，lpfc 驱动程序现在包括先前随 brcmnvme_fc 驱动程序提供的基于光纤通道的 NVMe（NVMe/FC）功能。



默认情况下，已为 LPe35000 系列适配器设置 `lpfc_enable_FC4_type=3` 参数。要为 LPe32000 系列和 LPe31000 系列适配器手动设置此命令，必须执行以下命令。

```
# 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
```

配置 Marvell/QLogic

1. 请参见，检查所需驱动程序 / 固件是否支持配置 "[NetApp 互操作性表](#)"。
2. 设置 `qlnativefc driver` 参数 `ql2xnvmesupport=1` 以在 `qlnativefc` 驱动程序中启用 NVMe/FC 支持，然后重新启动主机。

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



默认情况下，QLE 277x 系列适配器会设置 `qlnativefc driver` 参数。要为 QLE 277x 系列适配器手动设置此命令，必须执行以下命令。

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

3. 检查适配器上是否启用了 NVMe：

```

#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

```

验证 NVMe/FC

1. 验证 NVMe/FC 适配器是否列在 ESXi 主机上:

```

# 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. 验证是否已正确创建 NVMe/FC 命名空间:

以下示例中的 UUID 表示 NVMe/FC 命名空间设备。

```
# 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
```



在 ONTAP 9.7 中，NVMe/FC 命名空间的默认块大小为 4k。此默认大小与 ESXi 不兼容。因此，在为 ESXi 创建命名空间时，必须将命名空间块大小设置为 512b。您可以使用 `vserver nvme namespace create` 命令执行此操作。

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

请参见 ["ONTAP 9 命令手册页"](#) 了解更多详细信息。

3. 验证相应 NVMe/FC 命名空间设备的各个 ANA 路径的状态：

```

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}

```

已知问题描述

- 从 ONTAP 9.9.1 P3 开始，可提供 ESXi 7.0 U3（及更高版本）NVMe/FC 支持。这是由于仅在 ONTAP 9.9.1 P3 起提供的关键 NVMe 中止（由 ESXi 7.0 U3 及更高版本发出）修复程序。请参见相应的 burt 公有

报告，网址为 <https://mysupport.netapp.com/site/bugs-online/product/ONTAP/BURT/1420654> 了解详细信息。

相关链接

"TR-4597：采用 ONTAP 的 VMware vSphere"

"NetApp MetroCluster 支持 VMware vSphere 5.x，6.x 和 7.x（2031038）"

"NetApp® SnapMirror® Business Continuity（SM-BC）支持 VMware vSphere 6.x 和 7.x"

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.