



System: NVMe settings

SANtricity 11.6

NetApp
October 22, 2024

Table of Contents

- System: NVMe settings 1
 - Concepts 1
 - How tos 3
 - FAQs 5

System: NVMe settings

Concepts

NVMe overview

Some controllers include a port for implementing NVMe (Non-Volatile Memory Express) over fabrics. NVMe allows for high-performance communication between hosts and the storage array.

What is NVMe?

NVM stands for "Non-Volatile Memory" and is persistent memory used in many types of storage devices. NVMe (NVM Express) is a standardized interface or protocol designed specifically for high-performance multi-queue communication with NVM devices.

What is NVMe over Fabrics?

NVMe over Fabrics (NVMe-oF) is a technology specification that enables NVMe message-based commands and data to transfer between a host computer and storage over a network. An NVMe storage array (called a *subsystem*) can be accessed by a host using a fabric. NVMe commands are enabled and encapsulated in transport abstraction layers on both the host side and the subsystem side. This extends the high performance NVMe interface end-to-end from the host to the storage and standardizes and simplifies the command set.

NVMe-oF storage is presented to a host as a local block storage device. The volume (called a *namespace*) can be mounted to a file system as with any other block storage device. You can use the REST API, the SMcli, or SANtricity System Manager to provision your storage as needed.

What is an NVMe Qualified Name (NQN)?

The NVMe Qualified Name (NQN) is used to identify the remote storage target. The NVMe qualified name for the storage array is always assigned by the subsystem and may not be modified. There is only one NVMe Qualified Name for the entire array. The NVMe Qualified Name is limited to 223 characters in length. You can compare it to an iSCSI Qualified Name.

What is a namespace and a namespace ID?

A namespace is the equivalent of a logical unit in SCSI, which relates to a volume in the array. The namespace ID (NSID) is equivalent to a logical unit number (LUN) in SCSI. You create the NSID at namespace creation time, and can set it to a value between 1 and 255.

What is an NVMe controller?

Similar to a SCSI I_T nexus, which represents the path from the host's initiator to the storage system's target, an NVMe controller created during the host connection process provides an access path between a host and the namespaces in the storage array. An NQN for the host plus a host port identifier uniquely identify an NVMe controller. While an NVMe controller can only be associated with a single host, it can access multiple namespaces.

You configure which hosts can access which namespaces and set the namespace ID for the host using SANtricity System Manager. Then, when the NVMe controller is created, the list of namespace IDs accessible by the NVMe controller is created and used to configure the permissible connections.

NVMe terminology

Learn how the NVMe terms apply to your storage array.

Term	Description
InfiniBand	InfiniBand (IB) is a communications standard for data transmission between high-performance servers and storage systems.
Namespace	A namespace is NVM storage that is formatted for block access. It is analogous to a logical unit in SCSI, which relates to a volume in the storage array.
Namespace ID	The namespace ID is the NVMe controller's unique identifier for the namespace, and can be set to a value between 1 and 255. It is analogous to a logical unit number (LUN) in SCSI.
NQN	NVMe Qualified Name (NQN) is used to identify the remote storage target (the storage array).
NVM	Non-Volatile Memory (NVM) is persistent memory used in many types of storage devices.
NVMe	Non-Volatile Memory Express (NVMe) is an interface designed for flash-based storage devices, such as SSD drives. NVMe reduces I/O overhead and includes performance improvements, as compared to previous logical-device interfaces.
NVMe-oF	Non-Volatile Memory Express over Fabrics (NVMe-oF) is a specification that enables NVMe commands and data to transfer over a network between a host and storage.
NVMe controller	An NVMe controller is created during the host connection process. It provides an access path between a host and the namespaces in the storage array.
NVMe queue	A queue is used for passing commands and messages over the NVMe interface.
NVMe subsystem	The storage array with an NVMe host connection.
RDMA	Remote direct memory access (RDMA) enables more direct data movement in and out of a server by implementing a transport protocol in the network interface card (NIC) hardware.
RoCE	RDMA over Converged Ethernet (RoCE) is a network protocol that allows remote direct memory access (RDMA) over an Ethernet network.
SSD	Solid-state disks (SSDs) are data storage devices that use solid state memory (flash) to store data persistently. SSDs emulate conventional hard drives, and are available with the same interfaces that hard drives use.

How tos

Configure NVMe over InfiniBand ports

If your controller includes an NVMe over InfiniBand connection, you can configure the NVMe port settings from the System page.

Before you begin

- Your controller must include an NVMe over InfiniBand host port; otherwise, the NVMe over InfiniBand settings are not available in System Manager.
- You must know the IP address of the host connection.



The NVMe over InfiniBand settings and functions appear only if your storage array's controller includes an NVMe over InfiniBand port.

Steps

1. Select **Settings > System**.
2. Under **NVMe over InfiniBand settings**, select **Configure NVMe over InfiniBand ports**.
3. Select the controller with the NVMe over InfiniBand port you want to configure. Click **Next**.
4. Select the HIC port you want to configure from the drop-down list, and then enter the IP address.

If you are configuring an EF600 storage array with a 200Gb-capable HIC, this dialog box displays two IP Address fields, one for a physical port (external) and one for a virtual port (internal). You should assign a unique IP address for both ports. These settings allow the host to establish a path between each port, and for the HIC to achieve maximum performance. If you do not assign an IP address to the virtual port, the HIC will run at approximately half its capable speed.

5. Click **Finish**.
6. Reset the NVMe over InfiniBand port by clicking **Yes**.

Configure NVMe over RoCE ports

If your controller includes a connection for NVMe over RoCE (RDMA over Converged Ethernet), you can configure the NVMe port settings from the System page.

Before you begin


- Your controller must include an NVMe over RoCE host port; otherwise, the NVMe over RoCE settings are not available in System Manager.
- You must know the IP address of the host connection.

Steps

1. Select **Settings > System**.
2. Under **NVMe over ROCE settings**, select **Configure NVMe over ROCE ports**.
3. Select the controller with the NVMe over RoCE port you want to configure. Click **Next**.
4. Select the HIC port you want to configure from the drop-down list. Click **Next**.
5. Configure the port settings.

To see all port settings, click the **Show more port settings** link on the right of the dialog box.

Field Details

Port Setting	Description
Configured ethernet port speed	Select the speed that matches the speed capability of the SFP on the port.
Enable IPv4 / Enable IPv6	<div>Select one or both options to enable support for IPv4 and IPv6 networks.</div> <div> If you want to disable port access, deselect both check boxes.</div>
MTU size (Available by clicking Show more port settings .)	<div>If necessary, enter a new size in bytes for the Maximum Transmission Unit (MTU).</div> <div>The default Maximum Transmission Unit (MTU) size is 1500 bytes per frame. You must enter a value between 1500 and 9000.</div>

If you selected **Enable IPv4**, a dialog box opens for selecting IPv4 settings after you click **Next**. If you selected **Enable IPv6**, a dialog box opens for selecting IPv6 settings after you click **Next**. If you selected both options, the dialog box for IPv4 settings opens first, and then after you click **Next**, the dialog box for IPv6 settings opens.

6. Configure the IPv4 and/or IPv6 settings, either automatically or manually.

Field Details

Port setting	Description
Automatically obtain configuration	Select this option to obtain the configuration automatically.
Manually specify static configuration	Select this option, and then enter a static address in the fields. (If desired, you can cut and paste addresses into the fields.) For IPv4, include the network subnet mask and gateway. For IPv6, include the routable IP address and router IP address. If you are configuring an EF600 storage array with a 200Gb-capable HIC, this dialog box displays two sets of fields for network parameters, one for a physical port (external) and one for a virtual port (internal). You should assign unique parameters for both ports. These settings allow the host to establish a path between each port, and for the HIC to achieve maximum performance. If you do not assign an IP address to the virtual port, the HIC will run at approximately half its capable speed.

7. Click **Finish**.

View NVMe over Fabrics statistics

You can view data about the NVMe over Fabrics connections to your storage array.

About this task

System Manager shows these types of NVMe over Fabrics statistics. All statistics are read-only and cannot be set.

- **NVMe Subsystem statistics** — Shows statistics for the NVMe controller and its queue. The NVMe controller provides an access path between a host and the namespaces in the storage array. You can review the NVMe subsystem statistics for such items as connection failures, resets, and shutdowns.
- **RDMA Interface statistics** — Provides statistics for all NVMe over Fabrics ports on the RDMA interface, which includes performance statistics and link error information associated with each switch port. This tab only appears when NVMe over Fabrics ports are available.

You can view each of these statistics as raw statistics or as baseline statistics. Raw statistics are all of the statistics that have been gathered since the controllers were started. Baseline statistics are point-in-time statistics that have been gathered since you set the baseline time.

Steps

1. Select **Settings > System**.
2. Select **View NVMe over Fabrics Statistics**.
3. **Optional:** To set the baseline, click **Set new baseline**.

Setting the baseline sets a new starting point for the collection of the statistics. The same baseline is used for all NVMe statistics.

FAQs

How do I interpret NVMe over Fabrics statistics?

The View NVMe over Fabrics Statistics dialog box displays statistics for the NVMe subsystem and the RDMA interface. All statistics are read-only, and cannot be set.

- **NVMe Subsystem statistics** — Shows statistics for the NVMe controller and its queue. The NVMe controller provides an access path between a host and the namespaces in the storage array. You can review the NVMe subsystem statistics for such items as connection failures, resets, and shutdowns. For more information about these statistics, click **View legend for table headings**.
- **RDMA Interface statistics** — Provides statistics for all NVMe over Fabrics ports on the RDMA interface, which includes performance statistics and link error information associated with each switch port. This tab only appears when NVMe over Fabrics ports are available. For more information about the statistics, click **View legend for table headings**.

You can view each of these statistics as raw statistics or as baseline statistics. Raw statistics are all of the statistics that have been gathered since the controllers were started. Baseline statistics are point-in-time statistics that have been gathered since you set the baseline time.

What else do I need to do to configure or diagnose NVMe over InfiniBand?

The following table lists the System Manager functions that you can use to configure and

manage NVMe over InfiniBand sessions.



The NVMe over InfiniBand settings are available only if your storage array's controller includes an NVMe over InfiniBand port.

Configure and diagnose NVMe over InfiniBand

Action	Location
Configure NVMe over InfiniBand ports	<ol style="list-style-type: none">1. Select Hardware.2. Select Show back of shelf.3. Select a controller.4. Select Configure NVMe over InfiniBand ports. <p>or</p> <ol style="list-style-type: none">1. Select Settings > System.2. Scroll down to NVMe over InfiniBand settings, and then select Configure NVMe over InfiniBand Ports.
View NVMe over InfiniBand statistics	<ol style="list-style-type: none">1. Select Settings > System.2. Scroll down to NVMe over InfiniBand settings, and then select View NVMe over Fabrics Statistics.

What else do I need to do to configure or diagnose NVMe over RoCE?

You can configure and manage NVMe over RoCE from the Hardware and Settings pages.



The NVMe over RoCE settings are available only if your storage array's controller includes an NVMe over RoCE port.

Configure and diagnose NVMe over RoCE

Action	Location
Configure NVMe over RoCE ports	<ol style="list-style-type: none">1. Select Hardware.2. Select Show back of shelf.3. Select a controller.4. Select Configure NVMe over RoCE ports. <p>or</p> <ol style="list-style-type: none">1. Select Settings > System.2. Scroll down to NVMe over RoCE settings, and then select Configure NVMe over RoCE Ports.

Action	Location
View NVMe over Fabrics statistics	<ol style="list-style-type: none"> 1. Select Settings › System. 2. Scroll down to NVMe over RoCE settings, and then select View NVMe over Fabrics Statistics.

Why are there two IP addresses for one physical port?

The EF600 storage array can include two HICs — one external and one internal.

In this configuration, the external HIC is connected to an internal, auxiliary HIC. Each physical port that you can access from the external HIC has an associated virtual port from the internal HIC.

To achieve maximum 200Gb performance, you must assign a unique IP address for both the physical and virtual ports so the host can establish connections to each. If you do not assign an IP address to the virtual port, the HIC will run at approximately half its capable speed.

Why are there two sets of parameters for one physical port?

The EF600 storage array can include two HICs — one external and one internal.

In this configuration, the external HIC is connected to an internal, auxiliary HIC. Each physical port that you can access from the external HIC has an associated virtual port from the internal HIC.

To achieve maximum 200Gb performance, you must assign parameters for both the physical and virtual ports so the host can establish connections to each. If you do not assign parameters to the virtual port, the HIC will run at approximately half its capable speed.

Copyright information

Copyright © 2024 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.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

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.