



SANtricity® software

Express Configuration

For Linux®

May 2021 | 215-13046_2021-05_en-us
doccomments@netapp.com

Contents

Decide whether to use this Express Guide.....6

Fibre Channel Express Setup.....9

 Verify the Linux configuration is supported.....9

 Configure IP addresses using DHCP.....9

 Install and configure Linux Unified Host Utilities..... 10

 Install SANtricity Storage Manager for SMcli (SANtricity software version 11.53 or earlier)..... 10

 Access SANtricity System Manager and use the Setup wizard..... 11

 Configure the multipath software..... 12

 Setting up the multipath.conf file 14

 Configure the FC switches..... 14

 Determine host WWPNs and make the recommended settings..... 14

 Create partitions and filesystems..... 15

 Verify storage access on the host..... 16

 Record FC-specific information for Linux..... 16

 FC worksheet for Linux..... 16

SAS Express Setup..... 18

 Verify the Linux configuration is supported..... 18

 Configure IP addresses using DHCP..... 18

 Install and configure Host Utilities..... 19

 Install SANtricity Storage Manager for SMcli (SANtricity software version 11.53 or earlier)..... 19

 Access SANtricity System Manager and use the Setup wizard..... 20

 Configure the multipath software..... 21

 Setting up the multipath.conf file 22

 Determine SAS host identifiers - Linux 23

 Create partitions and filesystems..... 23

 Verify storage access on the host..... 24

 Record SAS-specific information for Linux..... 24

 SAS worksheet - Linux..... 25

iSCSI Express Setup..... 26

 Verify the Linux configuration is supported..... 26

 Configure IP addresses using DHCP..... 26

 Install and configure Host Utilities..... 27

 Install SANtricity Storage Manager for SMcli (SANtricity software version 11.53 or earlier)..... 27

 Access SANtricity System Manager and use the Setup wizard..... 28

 Configure the multipath software..... 29

 Setting up the multipath.conf file..... 30

 Configure the switches..... 31

 Configure networking..... 31

 Configure array-side networking..... 31

 Configure host-side networking..... 33

 Verify IP network connections..... 35

Create partitions and filesystems.....	36
Verify storage access on the host.....	37
Record iSCSI-specific information for Linux.....	37
iSCSI worksheet - Linux.....	37
iSER over InfiniBand Express Setup.....	39
Verify the Linux configuration is supported.....	39
Configure IP addresses using DHCP.....	39
Configure subnet manager.....	40
Install and configure Host Utilities.....	41
Install SANtricity Storage Manager for SMcli (SANtricity software version 11.53 or earlier).....	41
Access SANtricity System Manager and use the Setup wizard.....	42
Configure the multipath software.....	43
Setting up the multipath.conf file.....	44
Configure network connections—iSER over Infiniband.....	44
Configure networking for storage attached hosts—iSER over Infiniband.....	45
Create partitions and filesystems.....	47
Verify storage access on the host.....	48
Record iSER over InfiniBand-specific information for Linux.....	48
iSER over InfiniBand worksheet.....	48
SRP over InfiniBand Express Setup.....	51
Verify the Linux configuration is supported.....	51
Configure IP addresses using DHCP.....	51
Configure subnet manager.....	52
Install and configure Host Utilities.....	53
Install SANtricity Storage Manager for SMcli (SANtricity software version 11.53 or earlier).....	53
Access SANtricity System Manager and use the Setup wizard.....	54
Configure the multipath software.....	55
Setting up the multipath.conf file.....	56
Determining host port GUIDs and making the recommended settings—SRP over Infiniband.....	56
Configure network connections—SRP over Infiniband.....	57
Create partitions and filesystems.....	58
Verify storage access on the host.....	59
Record SRP over InfiniBand-specific information for Linux.....	59
SRP over Infiniband worksheet.....	60
NVMe over InfiniBand Express Setup.....	61
Verify the Linux configuration is supported.....	61
NVMe over InfiniBand restrictions.....	61
Configure IP addresses using DHCP.....	62
Install SANtricity Storage Manager for SMcli (SANtricity software version 11.53 or earlier).....	63
Access SANtricity System Manager and use the Setup wizard.....	64
Configure subnet manager.....	65
Set up NVMe over InfiniBand on the host side.....	66
Configure storage array NVMe over InfiniBand connections.....	69
Discover and connect to the storage from the host.....	69
Define a host.....	72
Assign a volume.....	73
Display the volumes visible to the host.....	74
Set up failover.....	74

Configuring the host to run failover.....	74
Accessing NVMe volumes.....	75
Accessing NVMe volumes for virtual device targets.....	75
Accessing NVMe volumes for physical NVMe device targets.....	77
Create filesystems	78
Create filesystems (RHEL 7 and SLES 12).....	78
Create filesystems (SLES 15).....	79
Verify storage access on the host.....	80
Record NVMe over InfiniBand-specific information for Linux.....	80
NVMe over InfiniBand worksheet.....	80

NVMe over RoCE Express Setup..... 83

Verify the Linux configuration is supported.....	83
NVMe over RoCE restrictions.....	83
Configuring IP addresses using DHCP.....	84
Install SANtricity Storage Manager for SMcli (SANtricity software version 11.53 or earlier).....	85
Access SANtricity System Manager and use Setup wizard.....	86
Configure the switch.....	87
Set up NVMe over RoCE on the host side.....	87
Configure storage array NVMe over RoCE connections.....	89
Discover and connect to the storage from the host.....	91
Define a host.....	92
Assign a volume.....	93
Display the volumes visible to the host.....	94
Set up failover on the host.....	94
Configuring the host to run failover.....	95
Accessing NVMe Volumes.....	95
Accessing NVMe volumes for virtual device targets.....	95
Accessing NVMe volumes for physical NVMe device targets.....	97
Create filesystems	98
Create filesystems (RHEL 7 and SLES 12).....	98
Create filesystems (SLES 15)	99
Verify storage access on the host.....	100
Record NVMe over RoCE-specific information for Linux.....	100
NVMe over RoCE worksheet for Linux.....	100

NVMe over Fibre Channel Express Setup..... 103

Verify the Linux configuration is supported.....	103
NVMe over Fibre Channel restrictions.....	103
Configuring IP addresses using DHCP.....	104
Install SANtricity Storage Manager for SMcli (SANtricity software version 11.53 or earlier).....	105
Access SANtricity System Manager and use Setup wizard.....	106
Configure the FC switches.....	107
Set up NVMe over Fibre Channel on the host side	107
Define a host.....	108
Assign a volume	109
Display the volumes visible to the host.....	110
Set up failover on the host	110
Configuring the host to run failover.....	111
Accessing NVMe volumes.....	111
Accessing NVMe volumes for virtual device targets.....	113
Accessing NVMe volumes for physical NVMe device targets.....	114

Create partitions and filesystems.....	116
Verify storage access on the host	117
Record NVMe over Fibre Channel-specific information for Linux.....	117
NVMe over Fibre Channel worksheet for Linux.....	117
 Where to find additional information - Linux	119
 Copyright and trademark.....	120
Copyright.....	120
Trademark.....	120

Decide whether to use this Express Guide

The express method for installing your storage array and accessing SANtricity System Manager is appropriate for setting up a standalone Linux host to an E-Series storage system. It is designed to get the storage system up and running as quickly as possible with minimal decision points.

Note: The configuration that the express method provides might not meet the needs of your production environment. For additional options for installing and configuring the storage system, see the *SANtricity Power Guide* for your operating system on the [E-Series and SANtricity Document Resources](#) page.

The express method includes the following steps:

1. Setting up one of the following communication environments:

- Fibre Channel (FC)
- iSCSI
- SAS
- iSER over Infiniband
- SRP over Infiniband
- NVMe over Infiniband
- NVMe over RoCE
- NVMe over Fibre Channel

2. Creating logical volumes on the storage array.

3. Making the volumes available to the data host.

This guide is based on the following assumptions:

Component	Assumptions
Hardware	<ul style="list-style-type: none">• You have used the Installation and Setup Instructions included with the controller shelves to install the hardware.• You have connected cables between the optional drive shelves and the controllers.• You have applied power to the storage system.• You have installed all other hardware (for example, management station, switches) and made the necessary connections.• If you are using NVMe over Infiniband, NVMe over RoCE, or NVMe over Fibre Channel, each EF300, EF600, EF570, or E5700 controller contains at least 32 GB of RAM.
Host	<ul style="list-style-type: none">• You have made a connection between the storage system and the data host.• You have installed the host operating system.• You are not using Linux as a virtualized guest.• You are not configuring the data (I/O attached) host to boot from SAN.• You have installed any OS updates as listed under the NetApp Interoperability Matrix Tool.

Component	Assumptions
Storage management station	<ul style="list-style-type: none"> You are using a 1 Gbps or faster management network. You are using a separate station for management rather than the data (I/O attached) host. You are using out-of-band management, in which a storage management station sends commands to the storage system through the Ethernet connections to the controller. You have attached the management station to the same subnet as the storage management ports.
IP addressing	<ul style="list-style-type: none"> You have installed and configured a DHCP server. You have not yet made an Ethernet connection between the management station and the storage system.
Storage provisioning	<ul style="list-style-type: none"> You will not use shared volumes. You will create pools rather than volume groups.
Protocol: FC	<ul style="list-style-type: none"> You have made all host-side FC connections and activated switch zoning. You are using NetApp-supported FC HBAs and switches. You are using FC HBA driver and firmware versions as listed in the NetApp Interoperability Matrix Tool.
Protocol: iSCSI	<ul style="list-style-type: none"> You are using Ethernet switches capable of transporting iSCSI traffic. You have configured the Ethernet switches according to the vendor's recommendation for iSCSI.
Protocol: SAS	<ul style="list-style-type: none"> You are using NetApp-supported SAS HBAs. You are using SAS HBA driver and firmware versions as listed in the NetApp Interoperability Matrix Tool.
Protocol: iSER over InfiniBand	<ul style="list-style-type: none"> You are using an InfiniBand fabric. You are using IB-iSER HBA driver and firmware versions as listed in the NetApp Interoperability Matrix Tool.
Protocol: SRP over InfiniBand	<ul style="list-style-type: none"> You are using an InfiniBand fabric. You are using IB-SRP driver and firmware versions as listed in the NetApp Interoperability Matrix Tool.
Protocol: NVMe over InfiniBand	<ul style="list-style-type: none"> You have received the 100G or 200G host interface cards in an EF300, EF600, EF570, or E5700 storage system pre-configured with the NVMe over InfiniBand protocol or the controllers were ordered with standard IB ports and need to be converted to NVMe-oF ports. You are using an InfiniBand fabric. You are using NVMe/IB driver and firmware versions as listed in the NetApp Interoperability Matrix Tool.
Protocol: NVMe over RoCE	<ul style="list-style-type: none"> You have received the 100G or 200G host interface cards in an EF300, EF600, EF570, or E5700 storage system pre-configured with the NVMe over RoCE protocol or the controllers were ordered with standard IB ports and need to be converted to NVMe-oF ports. You are using NVMe/RoCE driver and firmware versions as listed in the NetApp Interoperability Matrix Tool.

Component	Assumptions
Protocol: NVMe over Fibre Channel	<ul style="list-style-type: none">• You have received the 32G host interface cards in an EF300, EF600, EF570, or E5700 storage system pre-configured with the NVMe over Fibre Channel protocol or the controllers were ordered with standard FC ports and need to be converted to NVMe-oF ports.• You are using NVMe/FC driver and firmware versions as listed in the NetApp Interoperability Matrix Tool.

Note: The instructions in this guide include examples for SUSE Linux Enterprise Server (SLES) and for Red Hat Enterprise Linux (RHEL). Examples for RHEL are specific to RHEL7.

If these assumptions are not correct for your installation, or if you want more conceptual background information, see the SANtricity Power Guide for your operating system on the [E-Series and SANtricity Document Resources page](#).

Related information

[NetApp Interoperability Matrix Tool](#)
[Windows express configuration](#)

Fibre Channel Express Setup

Steps

1. [Verify the Linux configuration is supported](#) on page 9
2. [Configure IP addresses using DHCP](#) on page 9
3. [Install and configure Linux Unified Host Utilities](#) on page 10
4. [Install SANtricity Storage Manager for SMcli \(SANtricity software version 11.53 or earlier\)](#) on page 10
5. [Access SANtricity System Manager and use the Setup wizard](#) on page 11
6. [Configure the multipath software](#) on page 12
7. [Setting up the multipath.conf file](#) on page 14
8. [Configure the FC switches](#) on page 14
9. [Determine host WWPNS and make the recommended settings](#) on page 14
10. [Create partitions and filesystems](#) on page 15
11. [Verify storage access on the host](#) on page 16
12. [Record FC-specific information for Linux](#) on page 16

Verify the Linux configuration is supported

To ensure reliable operation, you create an implementation plan and then use the NetApp Interoperability Matrix Tool (IMT) to verify that the entire configuration is supported.

Steps

1. Go to the [NetApp Interoperability Matrix Tool](#).
2. Click on the **Solution Search** tile.
3. In the **Protocols > SAN Host** area, click the **Add** button next to **E-Series SAN Host**.
4. Click **View Refine Search Criteria**.
The **Refine Search Criteria** section is displayed. In this section you may select the protocol that applies, as well as other criteria for the configuration such as Operating System, NetApp OS, and Host Multipath driver. Select the criteria you know you want for your configuration, and then see what compatible configuration elements apply. As necessary, make the updates for your operating system and protocol that are prescribed in the tool. Detailed information for your chosen configuration is accessible on the View Supported Configurations page by clicking the **right page arrow**.

Configure IP addresses using DHCP

In this express method for configuring communications between the management station and the storage array, you use Dynamic Host Configuration Protocol (DHCP) to provide IP addresses. Each storage array has either one controller (simplex) or two controllers (duplex), and each controller has two storage management ports. Each management port will be assigned an IP address.

Before you begin

You have installed and configured a DHCP server on the same subnet as the storage management ports.

About this task

The following instructions refer to a storage array with two controllers (a duplex configuration).

1. If you have not already done so, connect an Ethernet cable to the management station and to management port 1 on each controller (A and B).

The DHCP server assigns an IP address to port 1 of each controller.

Note: Do not use management port 2 on either controller. Port 2 is reserved for use by NetApp technical personnel.

Important: If you disconnect and reconnect the Ethernet cable, or if the storage array is power-cycled, DHCP assigns IP addresses again. This process occurs until static IP addresses are configured. It is recommended that you avoid disconnecting the cable or power-cycling the array.

If the storage array cannot get DHCP-assigned IP addresses within 30 seconds, the following default IP addresses are set:

- Controller A, port 1: 169.254.128.101
 - Controller B, port 1: 169.254.128.102
 - Subnet mask: 255.255.0.0
2. Locate the MAC address label on the back of each controller, and then provide your network administrator with the MAC address for port 1 of each controller.
Your network administrator needs the MAC addresses to determine the IP address for each controller. You will need the IP addresses to connect to your storage system through your browser.

Install and configure Linux Unified Host Utilities

Linux Unified Host Utilities 7.1 includes tools to manage NetApp storage, including failover policies and physical paths.

Steps

1. Use the [NetApp Interoperability Matrix Tool](#) to determine the appropriate version of Unified Host Utilities 7.1 to install.

The versions are listed in a column within each supported configuration.

2. Download the Unified Host Utilities 7.1 from [NetApp Support](#).

Note: Alternatively, you can use the SANtricity `SMdevices` utility to perform the same functions as the Unified Host Utility tool. The `SMdevices` utility is included as part of the `SMutils` package. The `SMutils` package is a collection of utilities to verify what the host sees from the storage array. It is included as part of the SANtricity software installation.

Related information

[Linux Unified Host Utilities 7.1 Installation Guide](#)

Install SANtricity Storage Manager for SMcli (SANtricity software version 11.53 or earlier)

When you install the SANtricity Storage Manager software on your management station, the command line interface (CLI) is installed to help you manage your array. By also installing the software on the host, the Host Context Agent is installed that helps the host push configuration information to the storage array controllers through the I/O path.

Before you begin

Important: For SANtricity software 11.60 and newer, the SANtricity Secure CLI (SMcli) is included in the SANtricity OS and downloadable through the SANtricity System Manager. For more information on how to download the SMcli through the SANtricity System Manager, refer

to the *Download command line interface (CLI)* topic under the SANtricity System Manager Online Help.

- You are using SANtricity software 11.53 or earlier.
- You must have the correct administrator or superuser privileges.
- You must have ensured that the system that will contain the SANtricity Storage Manager client has the following minimum requirements:
 - **RAM:** 2 GB for Java Runtime Engine
 - **Disk space:** 5 GB
 - **OS/Architecture:** Refer to [NetApp Support Downloads > Software > E-Series/EF-Series SANtricity Storage Manager](#) for guidance on determining the supported operating system versions and architectures.

About this task

This section describes how to install SANtricity Storage Manager on both the Windows and Linux OS platforms, because both Windows and Linux are common management station platforms when Linux is used for the data host.

Steps

1. Download the SANtricity software release from [NetApp Support Downloads > Software > E-Series/EF-Series SANtricity Storage Manager](#).
2. Run the SANtricity installer.

Windows	Linux
Double-click the SMIA*.exe installation package to start the installation.	<ol style="list-style-type: none"> a. Go to the directory where the SMIA*.bin installation package is located. b. If the temp mount point does not have execute permissions, set the <code>IATEMPDIR</code> variable. Example: <code>IATEMPDIR=/root ./SMIA-LINUX64-11.25.0A00.0002.bin</code> c. Run the <code>chmod +x SMIA*.bin</code> command to grant execute permission to the file. d. Run the <code>./SMIA*.bin</code> command to start the installer.

3. Use the installation wizard to install the software on the management station.

Access SANtricity System Manager and use the Setup wizard

You use the Setup wizard in SANtricity System Manager to configure your storage array.

Before you begin

- You have ensured that the device from which you will access SANtricity System Manager contains one of the following browsers:

Browser	Minimum version
Google Chrome	47
Microsoft Internet Explorer	11
Microsoft Edge	EdgeHTML 12
Mozilla Firefox	31
Safari	9

- You are using out-of-band management.

About this task

The wizard automatically relaunches when you open System Manager or refresh your browser and *at least one* of the following conditions is met:

- No pools and volume groups are detected.
- No workloads are detected.
- No notifications are configured.

Steps

1. From your browser, enter the following URL:

https://<DomainNameOrIPAddress>

`IPAddress` is the address for one of the storage array controllers.

The first time SANtricity System Manager is opened on an array that has not been configured, the `Set Administrator Password` prompt appears. Role-based access management configures four local roles: admin, support, security, and monitor. The latter three roles have random passwords that cannot be guessed. After you set a password for the admin role you can change all of the passwords using the admin credentials. See *SANtricity System Manager online help* for more information on the four local user roles.

2. Enter the System Manager password for the admin role in the `Set Administrator Password` and `Confirm Password` fields, and then select the **Set Password** button. When you open System Manager and no pools, volumes groups, workloads, or notifications have been configured, the Setup wizard launches.
3. Use the Setup wizard to perform the following tasks:
 - **Verify hardware (controllers and drives)** – Verify the number of controllers and drives in the storage array. Assign a name to the array.
 - **Verify hosts and operating systems** – Verify the host and operating system types that the storage array can access.
 - **Accept pools** – Accept the recommended pool configuration for the express installation method. A pool is a logical group of drives.
 - **Configure alerts** – Allow System Manager to receive automatic notifications when a problem occurs with the storage array.
 - **Enable AutoSupport** – Automatically monitor the health of your storage array and have dispatches sent to technical support.
4. If you have not already created a volume, create one by going to **Storage > Volumes > Create > Volume**.

For more information, see the online help for SANtricity System Manager.

Configure the multipath software

Multipath software provides a redundant path to the storage array in case one of the physical paths is disrupted. The multipath software presents the operating system with a single virtual device that represents the active physical paths to the storage. The multipath software also manages the failover process that updates the virtual device. You use the device mapper multipath (DM-MP) tool for Linux installations.

Before you begin

You have installed the required packages on your system.

- For Red Hat (RHEL) hosts, verify the packages are installed by running `rpm -q device-mapper-multipath`.
- For SLES hosts, verify the packages are installed by running `rpm -q multipath-tools`.

About this task

By default, DM-MP is disabled in RHEL and SLES. Complete the following steps to enable DM-MP components on the host.

If you have not already installed the operating system, use the media supplied by your operating system vendor.

Steps

1. If a `multipath.conf` file is not already created, run the `# touch /etc/multipath.conf` command.
2. Use the default multipath settings by leaving the `multipath.conf` file blank.
3. Start the multipath service.

```
# systemctl start multipathd
```

4. Save your kernel version by running the `uname -r` command.

```
# uname -r
3.10.0-327.el7.x86_64
```

You will use this information when you assign volumes to the host.

5. Do one of the following to enable the `multipathd` daemon on boot.

If you are using....	Do this...
RHEL 7.x and 8.x systems:	<code>systemctl enable multipathd</code>
SLES 12.x and 15.x systems:	<code>systemctl enable multipathd</code>

6. Rebuild the `initramfs` image or the `initrd` image under `/boot` directory:

If you are using....	Do this...
RHEL 7.x and 8.x systems:	<code>dracut --force --add multipath</code>
SLES 12.x and 15.x systems:	<code>dracut --force --add multipath</code>

7. Make sure that the newly created `/boot/initrams-*` image or `/boot/initrd-*` image is selected in the boot configuration file.
For example, for grub it is `/boot/grub/menu.lst` and for grub2 it is `/boot/grub2/menu.cfg`.
8. Use the "Create host manually" procedure in the online help to check whether the hosts are defined. Verify that each host type is either **Linux DM-MP (Kernel 3.10 or later)** if you enable the Automatic Load Balancing feature, or **Linux DM-MP (Kernel 3.9 or earlier)** if you disable the Automatic Load Balancing feature. If necessary, change the selected host type to the appropriate setting.
9. Reboot the host.

Setting up the multipath.conf file

The `multipath.conf` file is the configuration file for the multipath daemon, `multipathd`. The `multipath.conf` file overrides the built-in configuration table for `multipathd`. Any line in the file whose first non-white-space character is `#` is considered a comment line. Empty lines are ignored.

Note: For SANtricity operating system 8.30 and newer, NetApp recommends using the default settings as provided.

The `multipath.conf` are available in the following locations:

- For SLES, `/usr/share/doc/packages/multipath-tools/multipath.conf.synthetic`
- For RHEL, `/usr/share/doc/device-mapper-multipath-0.4.9/multipath.conf`

Configure the FC switches

Configuring (zoning) the Fibre Channel (FC) switches enables the hosts to connect to the storage array and limits the number of paths. You zone the switches using the management interface for the switches.

Before you begin

- You must have administrator credentials for the switches.
- You must have used your HBA utility to discover the WWPN of each host initiator port and of each controller target port connected to the switch.

About this task

For details about zoning your switches, see the switch vendor's documentation.

Each initiator port must be in a separate zone with all of its corresponding target ports.

Steps

1. Log in to the FC switch administration program, and then select the zoning configuration option.
2. Create a new zone that includes the first host initiator port and that also includes all of the target ports that connect to the same FC switch as the initiator.
3. Create additional zones for each FC host initiator port in the switch.
4. Save the zones, and then activate the new zoning configuration.

Determine host WWPNs and make the recommended settings

You install an FC HBA utility so you can view the worldwide port name (WWPN) of each host port. Additionally, you can use the HBA utility to change any settings recommended in the Notes column of the [NetApp Interoperability Matrix Tool](#) for the supported configuration.

About this task

Guidelines for HBA utilities:

- Most HBA vendors offer an HBA utility. You will need the correct version of HBA for your host operating system and CPU. Examples of FC HBA utilities include:
 - Emulex OneCommand Manager for Emulex HBAs
 - QLogic QConverge Console for QLogic HBAs
- Host I/O ports might automatically register if the host context agent is installed.

Steps

- 1. Download the appropriate utility from your HBA vendor's web site.
 - 2. Install the utility.
 - 3. Select the appropriate settings in the HBA utility.
- Appropriate settings for your configuration are listed in the Notes column of the IMT.

Create partitions and filesystems

A new LUN has no partition or file system when the Linux host first discovers it. You must format the LUN before it can be used. Optionally, you can create a file system on the LUN.

Before you begin

The host must have discovered the LUN.

In the /dev/mapper folder, you have run the ls command to see the available disks.

About this task

You can initialize the disk as a basic disk with a GUID partition table (GPT) or Master boot record (MBR).

Format the LUN with a file system such as ext4. Some applications do not require this step.

Steps

- 1. Retrieve the SCSI ID of the mapped disk by issuing the sanlun lun show -p command.
- The SCSI ID is a 33-character string of hexadecimal digits, beginning with the number 3. If user-friendly names are enabled, Device Mapper reports disks as mpath instead of by a SCSI ID.

```
# sanlun lun show -p

E-Series Array: ictm1619s01c01-SRP(60080e50002908b40000000054efb9d2)
Volume Name:
Preferred Owner: Controller in Slot B
Current Owner: Controller in Slot B
Mode: RDAC (Active/Active)
UTM LUN: None
LUN: 116
LUN Size:
Product: E-Series
Host Device: mpathr(360080e50004300ac000007575568851d)
Multipath Policy: round-robin 0
Multipath Provider: Native

-----
-----
host      controller
path      path      /dev/      host      controller
state     type      node      adapter   target
-----
-----
up        secondary sdcx      host14    A1
up        secondary sdat      host10    A2
up        secondary sdbv      host13    B1
```

- 2. Create a new partition according to the method appropriate for your Linux OS release.

Typically, characters identifying the partition of a disk are appended to the SCSI ID (the number 1 or p3 for instance).

```
# parted -a optimal -s -- /dev/mapper/360080e5000321bb8000092b1535f887a mklabel  
gpt mkpart primary ext4 0% 100%
```

3. Create a file system on the partition.

The method for creating a file system varies depending on the file system chosen.

```
# mkfs.ext4 /dev/mapper/360080e5000321bb8000092b1535f887a1
```

4. Create a folder to mount the new partition.

```
# mkdir /mnt/ext4
```

5. Mount the partition.

```
# mount /dev/mapper/360080e5000321bb8000092b1535f887a1 /mnt/ext4
```

Verify storage access on the host

Before using the volume, you verify that the host can write data to the volume and read it back.

Before you begin

You must have initialized the volume and formatted it with a file system.

Steps

1. On the host, copy one or more files to the mount point of the disk.
2. Copy the files back to a different folder on the original disk.
3. Run the `diff` command to compare the copied files to the originals.

After you finish

Remove the file and folder that you copied.


Record FC-specific information for Linux

Select the FC worksheet to record your protocol-specific storage configuration information. You need this information to perform provisioning tasks.

FC worksheet for Linux

You can use this worksheet to record FC storage configuration information. You need this information to perform provisioning tasks.

The illustration shows a host connected to an E-Series storage array in two zones. One zone is indicated by the blue line; the other zone is indicated by the red line. Any single port has two paths to the storage (one to each controller).



Host identifiers

Callout No.	Host (initiator) port connections	WWPN
1	Host	<i>not applicable</i>
2	Host port 0 to FC switch zone 0	
7	Host port 1 to FC switch zone 1	

Target identifiers

Callout No.	Array controller (target) port connections	WWPN
3	Switch	<i>not applicable</i>
6	Array controller (target)	<i>not applicable</i>
5	Controller A, port 1 to FC switch 1	
9	Controller A, port 2 to FC switch 2	
4	Controller B, port 1 to FC switch 1	
8	Controller B, port 2 to FC switch 2	

Mapping host

Mapping host name	
Host OS type	

SAS Express Setup

Steps

1. [Verify the Linux configuration is supported](#) on page 18
2. [Configure IP addresses using DHCP](#) on page 18
3. [Install and configure Host Utilities](#) on page 19
4. [Install SANtricity Storage Manager for SMcli \(SANtricity software version 11.53 or earlier\)](#) on page 19
5. [Access SANtricity System Manager and use the Setup wizard](#) on page 20
6. [Configure the multipath software](#) on page 21
7. [Setting up the multipath.conf file](#) on page 22
8. [Determine SAS host identifiers - Linux](#) on page 23
9. [Create partitions and filesystems](#) on page 23
10. [Verify storage access on the host](#) on page 24
11. [Record SAS-specific information for Linux](#) on page 24

Verify the Linux configuration is supported

To ensure reliable operation, you create an implementation plan and then use the NetApp Interoperability Matrix Tool (IMT) to verify that the entire configuration is supported.

Steps

1. Go to the [NetApp Interoperability Matrix Tool](#).
2. Click on the **Solution Search** tile.
3. In the **Protocols > SAN Host** area, click the **Add** button next to **E-Series SAN Host**.
4. Click **View Refine Search Criteria**.
The **Refine Search Criteria** section is displayed. In this section you may select the protocol that applies, as well as other criteria for the configuration such as Operating System, NetApp OS, and Host Multipath driver. Select the criteria you know you want for your configuration, and then see what compatible configuration elements apply. As necessary, make the updates for your operating system and protocol that are prescribed in the tool. Detailed information for your chosen configuration is accessible on the View Supported Configurations page by clicking the **right page arrow**.

Configure IP addresses using DHCP

In this express method for configuring communications between the management station and the storage array, you use Dynamic Host Configuration Protocol (DHCP) to provide IP addresses. Each storage array has either one controller (simplex) or two controllers (duplex), and each controller has two storage management ports. Each management port will be assigned an IP address.

Before you begin

You have installed and configured a DHCP server on the same subnet as the storage management ports.

About this task

The following instructions refer to a storage array with two controllers (a duplex configuration).

1. If you have not already done so, connect an Ethernet cable to the management station and to management port 1 on each controller (A and B).

The DHCP server assigns an IP address to port 1 of each controller.

Note: Do not use management port 2 on either controller. Port 2 is reserved for use by NetApp technical personnel.

Important: If you disconnect and reconnect the Ethernet cable, or if the storage array is power-cycled, DHCP assigns IP addresses again. This process occurs until static IP addresses are configured. It is recommended that you avoid disconnecting the cable or power-cycling the array.

If the storage array cannot get DHCP-assigned IP addresses within 30 seconds, the following default IP addresses are set:

- Controller A, port 1: 169.254.128.101
 - Controller B, port 1: 169.254.128.102
 - Subnet mask: 255.255.0.0
2. Locate the MAC address label on the back of each controller, and then provide your network administrator with the MAC address for port 1 of each controller.
Your network administrator needs the MAC addresses to determine the IP address for each controller. You will need the IP addresses to connect to your storage system through your browser.

Install and configure Host Utilities

Linux Unified Host Utilities 7.1 includes tools to manage NetApp storage, including failover policies and physical paths.

Steps

1. Use the [NetApp Interoperability Matrix Tool](#) to determine the appropriate version of Unified Host Utilities 7.1 to install.

The versions are listed in a column within each supported configuration.

2. Download the Unified Host Utilities 7.1 from [NetApp Support](#).

Note: Alternatively, you can use the SANtricity `SMdevices` utility to perform the same functions as the Unified Host Utility tool. The `SMdevices` utility is included as part of the `SMutils` package. The `SMutils` package is a collection of utilities to verify what the host sees from the storage array. It is included as part of the SANtricity software installation.

Install SANtricity Storage Manager for SMcli (SANtricity software version 11.53 or earlier)

When you install the SANtricity Storage Manager software on your management station, the command line interface (CLI) is installed to help you manage your array. By also installing the software on the host, the Host Context Agent is installed that helps the host push configuration information to the storage array controllers through the I/O path.

Before you begin

Important: For SANtricity software 11.60 and newer, the SANtricity Secure CLI (SMcli) is included in the SANtricity OS and downloadable through the SANtricity System Manager. For more information on how to download the SMcli through the SANtricity System Manager, refer to the *Download the command line interface (CLI)* topic under the SANtricity System Manager Online Help.

- You are using SANtricity software 11.53 or earlier.
- You must have the correct administrator or superuser privileges.
- You must have ensured that the system that will contain the SANtricity Storage Manager client has the following minimum requirements:
 - **RAM:** 2 GB for Java Runtime Engine
 - **Disk space:** 5 GB
 - **OS/Architecture:** Refer to [NetApp Support Downloads > Software > E-Series/EF-Series SANtricity Storage Manager](#) for guidance on determining the supported operating system versions and architectures.

About this task

This section describes how to install SANtricity Storage Manager on both the Windows and Linux OS platforms, because both Windows and Linux are common management station platforms when Linux is used for the data host.

Steps

1. Download the SANtricity software release from [NetApp Support Downloads > Software > E-Series/EF-Series SANtricity Storage Manager](#).
2. Run the SANtricity installer.

Windows	Linux
Double-click the SMIA*.exe installation package to start the installation.	<ol style="list-style-type: none"> a. Go to the directory where the SMIA*.bin installation package is located. b. If the temp mount point does not have execute permissions, set the <code>IATEMPDIR</code> variable. Example: <code>IATEMPDIR=/root ./SMIA-LINUX64-11.25.0A00.0002.bin</code> c. Run the <code>chmod +x SMIA*.bin</code> command to grant execute permission to the file. d. Run the <code>./SMIA*.bin</code> command to start the installer.

3. Use the installation wizard to install the software on the management station.

Access SANtricity System Manager and use the Setup wizard

You use the Setup wizard in SANtricity System Manager to configure your storage array.

Before you begin

- You have ensured that the device from which you will access SANtricity System Manager contains one of the following browsers:

Browser	Minimum version
Google Chrome	47
Microsoft Internet Explorer	11
Microsoft Edge	EdgeHTML 12
Mozilla Firefox	31
Safari	9

- You are using out-of-band management.

About this task

The wizard automatically relaunches when you open System Manager or refresh your browser and *at least one* of the following conditions is met:

- No pools and volume groups are detected.
- No workloads are detected.
- No notifications are configured.

Steps

1. From your browser, enter the following URL:

https://<DomainNameOrIPAddress>

IPAddress is the address for one of the storage array controllers.

The first time SANtricity System Manager is opened on an array that has not been configured, the `Set Administrator Password` prompt appears. Role-based access management configures four local roles: admin, support, security, and monitor. The latter three roles have random passwords that cannot be guessed. After you set a password for the admin role you can change all of the passwords using the admin credentials. See *SANtricity System Manager online help* for more information on the four local user roles.

2. Enter the System Manager password for the admin role in the `Set Administrator Password` and `Confirm Password` fields, and then select the **Set Password** button. When you open System Manager and no pools, volumes groups, workloads, or notifications have been configured, the Setup wizard launches.
3. Use the Setup wizard to perform the following tasks:
 - **Verify hardware (controllers and drives)** – Verify the number of controllers and drives in the storage array. Assign a name to the array.
 - **Verify hosts and operating systems** – Verify the host and operating system types that the storage array can access.
 - **Accept pools** – Accept the recommended pool configuration for the express installation method. A pool is a logical group of drives.
 - **Configure alerts** – Allow System Manager to receive automatic notifications when a problem occurs with the storage array.
 - **Enable AutoSupport** – Automatically monitor the health of your storage array and have dispatches sent to technical support.
4. If you have not already created a volume, create one by going to **Storage > Volumes > Create > Volume**.

For more information, see the online help for SANtricity System Manager.

Configure the multipath software

Multipath software provides a redundant path to the storage array in case one of the physical paths is disrupted. The multipath software presents the operating system with a single virtual device that represents the active physical paths to the storage. The multipath software also manages the failover process that updates the virtual device. You use the device mapper multipath (DM-MP) tool for Linux installations.

Before you begin

You have installed the required packages on your system.

- For Red Hat (RHEL) hosts, verify the packages are installed by running `rpm -q device-mapper-multipath`.
- For SLES hosts, verify the packages are installed by running `rpm -q multipath-tools`.

About this task

By default, DM-MP is disabled in RHEL and SLES. Complete the following steps to enable DM-MP components on the host.

If you have not already installed the operating system, use the media supplied by your operating system vendor.

Steps

- 1. If a `multipath.conf` file is not already created, run the `# touch /etc/multipath.conf` command.
- 2. Use the default multipath settings by leaving the `multipath.conf` file blank.
- 3. Start the multipath service.

```
# systemctl start multipathd
```

- 4. Save your kernel version by running the `uname -r` command.

```
# uname -r
3.10.0-327.el7.x86_64
```

You will use this information when you assign volumes to the host.

- 5. Do one of the following to enable the `multipathd` daemon on boot.

If you are using....	Do this...
RHEL 7.x and 8.x systems:	<code>systemctl enable multipathd</code>
SLES 12.x and 15.x systems:	<code>systemctl enable multipathd</code>

- 6. Rebuild the `initramfs` image or the `initrd` image under `/boot` directory:

If you are using....	Do this...
RHEL 7.x and 8.x systems:	<code>dracut --force --add multipath</code>
SLES 12.x and 15.x systems:	<code>dracut --force --add multipath</code>

- 7. Make sure that the newly created `/boot/initramfs-*` image or `/boot/initrd-*` image is selected in the boot configuration file.
For example, for grub it is `/boot/grub/menu.lst` and for grub2 it is `/boot/grub2/menu.cfg`.
- 8. Use the "Create host manually" procedure in the online help to check whether the hosts are defined. Verify that each host type is either **Linux DM-MP (Kernel 3.10 or later)** if you enable the Automatic Load Balancing feature, or **Linux DM-MP (Kernel 3.9 or earlier)** if you disable the Automatic Load Balancing feature. If necessary, change the selected host type to the appropriate setting.
- 9. Reboot the host.

Setting up the `multipath.conf` file

The `multipath.conf` file is the configuration file for the multipath daemon, `multipathd`. The `multipath.conf` file overrides the built-in configuration table for `multipathd`. Any line in the file whose first non-white-space character is `#` is considered a comment line. Empty lines are ignored.

Note: For SANtricity operating system 8.30 and newer, NetApp recommends using the default settings as provided.

The `multipath.conf` are available in the following locations:

- For SLES, `/usr/share/doc/packages/multipath-tools/multipath.conf.synthetic`
- For RHEL, `/usr/share/doc/device-mapper-multipath-0.4.9/multipath.conf`

Determine SAS host identifiers - Linux

For the SAS protocol, you find the SAS addresses using the HBA utility, then use the HBA BIOS to make the appropriate configuration settings.

About this task

Guidelines for HBA utilities:

- Most HBA vendors offer an HBA utility. Depending on your host operating system and CPU, use either the LSI-sas2flash(6G) or sas3flash(12G) utility.
- Host I/O ports might automatically register if the host context agent is installed.

Steps

1. Download the HBA utility from your HBA vendor's web site.
2. Install the utility.
3. Use the HBA BIOS to select the appropriate settings for your configuration.

See the Notes column of the [NetApp Interoperability Matrix Tool](#) for recommendations.

Create partitions and filesystems

A new LUN has no partition or file system when the Linux host first discovers it. You must format the LUN before it can be used. Optionally, you can create a file system on the LUN.

Before you begin

The host must have discovered the LUN.

In the `/dev/mapper` folder, you have run the `ls` command to see the available disks.

About this task

You can initialize the disk as a basic disk with a GUID partition table (GPT) or Master boot record (MBR).

Format the LUN with a file system such as ext4. Some applications do not require this step.

Steps

1. Retrieve the SCSI ID of the mapped disk by issuing the `sanlun lun show -p` command.

The SCSI ID is a 33-character string of hexadecimal digits, beginning with the number 3. If user-friendly names are enabled, Device Mapper reports disks as `mpath` instead of by a SCSI ID.

```
# sanlun lun show -p
```

```
E-Series Array: ictml619s01c01-SRP(60080e50002908b40000000054efb9d2)
Volume Name:
Preferred Owner: Controller in Slot B
Current Owner: Controller in Slot B
Mode: RDAC (Active/Active)
UTM LUN: None
LUN: 116
LUN Size:
```

```
Product: E-Series
Host Device: mpathr(360080e50004300ac000007575568851d)
Multipath Policy: round-robin 0
Multipath Provider: Native
-----
-----
host      controller
path      path      /dev/      host      controller
state     type       node      adapter   target
-----
-----
up        secondary  sdcx      host14     A1
up        secondary  sdat      host10     A2
up        secondary  sdbv      host13     B1
```

2. Create a new partition according to the method appropriate for your Linux OS release.
Typically, characters identifying the partition of a disk are appended to the SCSI ID (the number 1 or p3 for instance).

```
# parted -a optimal -s -- /dev/mapper/360080e5000321bb8000092b1535f887a mklabel
gpt mkpart primary ext4 0% 100%
```

3. Create a file system on the partition.
The method for creating a file system varies depending on the file system chosen.

```
# mkfs.ext4 /dev/mapper/360080e5000321bb8000092b1535f887a1
```

4. Create a folder to mount the new partition.

```
# mkdir /mnt/ext4
```

5. Mount the partition.

```
# mount /dev/mapper/360080e5000321bb8000092b1535f887a1 /mnt/ext4
```

Verify storage access on the host

Before using the volume, you verify that the host can write data to the volume and read it back.

Before you begin

You must have initialized the volume and formatted it with a file system.

Steps

- 1. On the host, copy one or more files to the mount point of the disk.
- 2. Copy the files back to a different folder on the original disk.
- 3. Run the diff command to compare the copied files to the originals.

After you finish


Remove the file and folder that you copied.

Record SAS-specific information for Linux

Record your protocol-specific storage configuration information on the SAS worksheet. You need this information to perform provisioning tasks.

SAS worksheet - Linux

You can use this worksheet to record SAS storage configuration information. You need this information to perform provisioning tasks.



Host Identifiers

Callout No.	Host (initiator) port connections	SAS address
1	Host	<i>not applicable</i>
2	Host (initiator) port 1 connected to Controller A, port 1	
3	Host (initiator) port 1 connected to Controller B, port 1	
4	Host (initiator) port 2 connected to Controller A, port 1	
5	Host (initiator) port 2 connected to Controller B, port 1	

Target Identifiers

Recommended configurations consist of two target ports.

Mappings Host

Mappings Host Name	
Host OS Type	

iSCSI Express Setup

Steps

1. [Verify the Linux configuration is supported](#) on page 26
2. [Configure IP addresses using DHCP](#) on page 26
3. [Install and configure Host Utilities](#) on page 27
4. [Install SANtricity Storage Manager for SMcli \(SANtricity software version 11.53 or earlier\)](#) on page 27
5. [Access SANtricity System Manager and use the Setup wizard](#) on page 28
6. [Configure the multipath software](#) on page 29
7. [Setting up the multipath.conf file](#) on page 30
8. [Configure the switches](#) on page 31
9. [Configure networking](#) on page 31
10. [Configure array-side networking](#) on page 31
11. [Configure host-side networking](#) on page 33
12. [Verify IP network connections](#) on page 35
13. [Create partitions and filesystems](#) on page 36
14. [Verify storage access on the host](#) on page 37
15. [Record iSCSI-specific information for Linux](#) on page 37

Verify the Linux configuration is supported

To ensure reliable operation, you create an implementation plan and then use the NetApp Interoperability Matrix Tool (IMT) to verify that the entire configuration is supported.

Steps

1. Go to the [NetApp Interoperability Matrix Tool](#).
2. Click on the **Solution Search** tile.
3. In the **Protocols > SAN Host** area, click the **Add** button next to **E-Series SAN Host**.
4. Click **View Refine Search Criteria**.
The **Refine Search Criteria** section is displayed. In this section you may select the protocol that applies, as well as other criteria for the configuration such as Operating System, NetApp OS, and Host Multipath driver. Select the criteria you know you want for your configuration, and then see what compatible configuration elements apply. As necessary, make the updates for your operating system and protocol that are prescribed in the tool. Detailed information for your chosen configuration is accessible on the View Supported Configurations page by clicking the **right page arrow**.

Configure IP addresses using DHCP

In this express method for configuring communications between the management station and the storage array, you use Dynamic Host Configuration Protocol (DHCP) to provide IP addresses. Each storage array has either one controller (simplex) or two controllers (duplex), and each controller has two storage management ports. Each management port will be assigned an IP address.

Before you begin

You have installed and configured a DHCP server on the same subnet as the storage management ports.

About this task

The following instructions refer to a storage array with two controllers (a duplex configuration).

1. If you have not already done so, connect an Ethernet cable to the management station and to management port 1 on each controller (A and B).

The DHCP server assigns an IP address to port 1 of each controller.

Note: Do not use management port 2 on either controller. Port 2 is reserved for use by NetApp technical personnel.

Important: If you disconnect and reconnect the Ethernet cable, or if the storage array is power-cycled, DHCP assigns IP addresses again. This process occurs until static IP addresses are configured. It is recommended that you avoid disconnecting the cable or power-cycling the array.

If the storage array cannot get DHCP-assigned IP addresses within 30 seconds, the following default IP addresses are set:

- Controller A, port 1: 169.254.128.101
 - Controller B, port 1: 169.254.128.102
 - Subnet mask: 255.255.0.0
2. Locate the MAC address label on the back of each controller, and then provide your network administrator with the MAC address for port 1 of each controller.
Your network administrator needs the MAC addresses to determine the IP address for each controller. You will need the IP addresses to connect to your storage system through your browser.

Install and configure Host Utilities

Linux Unified Host Utilities 7.1 includes tools to manage NetApp storage, including failover policies and physical paths.

Steps

1. Use the [NetApp Interoperability Matrix Tool](#) to determine the appropriate version of Unified Host Utilities 7.1 to install.

The versions are listed in a column within each supported configuration.

2. Download the Unified Host Utilities 7.1 from [NetApp Support](#).

Note: Alternatively, you can use the SANtricity `SMdevices` utility to perform the same functions as the Unified Host Utility tool. The `SMdevices` utility is included as part of the `SMutils` package. The `SMutils` package is a collection of utilities to verify what the host sees from the storage array. It is included as part of the SANtricity software installation.

Install SANtricity Storage Manager for SMcli (SANtricity software version 11.53 or earlier)

When you install the SANtricity Storage Manager software on your management station, the command line interface (CLI) is installed to help you manage your array. By also installing the software on the host, the Host Context Agent is installed that helps the host push configuration information to the storage array controllers through the I/O path.

Before you begin

Important: For SANtricity software 11.60 and newer, the SANtricity Secure CLI (SMcli) is included in the SANtricity OS and downloadable through the SANtricity System Manager. For more information on how to download the SMcli through the SANtricity System Manager, refer

to the *Download the command line interface (CLI)* topic under the SANtricity System Manager Online Help.

- You are using SANtricity software 11.53 or earlier.
- You must have the correct administrator or superuser privileges.
- You must have ensured that the system that will contain the SANtricity Storage Manager client has the following minimum requirements:
 - **RAM:** 2 GB for Java Runtime Engine
 - **Disk space:** 5 GB
 - **OS/Architecture:** Refer to [NetApp Support Downloads > Software > E-Series/EF-Series SANtricity Storage Manager](#) for guidance on determining the supported operating system versions and architectures.

About this task

This section describes how to install SANtricity Storage Manager on both the Windows and Linux OS platforms, because both Windows and Linux are common management station platforms when Linux is used for the data host.

Steps

1. Download the SANtricity software release from [NetApp Support Downloads > Software > E-Series/EF-Series SANtricity Storage Manager](#).
2. Run the SANtricity installer.

Windows	Linux
Double-click the SMIA*.exe installation package to start the installation.	<ol style="list-style-type: none">a. Go to the directory where the SMIA*.bin installation package is located.b. If the temp mount point does not have execute permissions, set the <code>IATEMPDIR</code> variable. Example: <code>IATEMPDIR=/root ./SMIA-LINUX64-11.25.0A00.0002.bin</code>c. Run the <code>chmod +x SMIA*.bin</code> command to grant execute permission to the file.d. Run the <code>./SMIA*.bin</code> command to start the installer.

3. Use the installation wizard to install the software on the management station.

Access SANtricity System Manager and use the Setup wizard

You use the Setup wizard in SANtricity System Manager to configure your storage array.

Before you begin

- You have ensured that the device from which you will access SANtricity System Manager contains one of the following browsers:

Browser	Minimum version
Google Chrome	47
Microsoft Internet Explorer	11
Microsoft Edge	EdgeHTML 12
Mozilla Firefox	31
Safari	9

- You are using out-of-band management.

About this task

If you are an iSCSI user, you closed the Setup wizard while configuring iSCSI.

The wizard automatically relaunches when you open System Manager or refresh your browser and *at least one* of the following conditions is met:

- No pools and volume groups are detected.
- No workloads are detected.
- No notifications are configured.

Steps

1. From your browser, enter the following URL:

`https://<DomainNameOrIPAddress>`

`IPAddress` is the address for one of the storage array controllers.

The first time SANtricity System Manager is opened on an array that has not been configured, the `Set Administrator Password` prompt appears. Role-based access management configures four local roles: admin, support, security, and monitor. The latter three roles have random passwords that cannot be guessed. After you set a password for the admin role you can change all of the passwords using the admin credentials. See *SANtricity System Manager online help* for more information on the four local user roles.

2. Enter the System Manager password for the admin role in the `Set Administrator Password` and `Confirm Password` fields, and then select the **Set Password** button. When you open System Manager and no pools, volumes groups, workloads, or notifications have been configured, the Setup wizard launches.
3. Use the Setup wizard to perform the following tasks:
 - **Verify hardware (controllers and drives)** – Verify the number of controllers and drives in the storage array. Assign a name to the array.
 - **Verify hosts and operating systems** – Verify the host and operating system types that the storage array can access.
 - **Accept pools** – Accept the recommended pool configuration for the express installation method. A pool is a logical group of drives.
 - **Configure alerts** – Allow System Manager to receive automatic notifications when a problem occurs with the storage array.
 - **Enable AutoSupport** – Automatically monitor the health of your storage array and have dispatches sent to technical support.
4. If you have not already created a volume, create one by going to **Storage > Volumes > Create > Volume**.

For more information, see the online help for SANtricity System Manager.

Configure the multipath software

Multipath software provides a redundant path to the storage array in case one of the physical paths is disrupted. The multipath software presents the operating system with a single virtual device that represents the active physical paths to the storage. The multipath software also manages the failover process that updates the virtual device. You use the device mapper multipath (DM-MP) tool for Linux installations.

Before you begin

You have installed the required packages on your system.

- For Red Hat (RHEL) hosts, verify the packages are installed by running `rpm -q device-mapper-multipath`.
- For SLES hosts, verify the packages are installed by running `rpm -q multipath-tools`.

About this task

By default, DM-MP is disabled in RHEL and SLES. Complete the following steps to enable DM-MP components on the host.

If you have not already installed the operating system, use the media supplied by your operating system vendor.

Steps

1. If a `multipath.conf` file is not already created, run the `# touch /etc/multipath.conf` command.
2. Use the default multipath settings by leaving the `multipath.conf` file blank.
3. Start the multipath service.

```
# systemctl start multipathd
```

4. Save your kernel version by running the `uname -r` command.

```
# uname -r
3.10.0-327.el7.x86_64
```

You will use this information when you assign volumes to the host.

5. Do one of the following to enable the `multipathd` daemon on boot.

If you are using....	Do this...
RHEL 7.x and 8.x systems:	<code>systemctl enable multipathd</code>
SLES 12.x and 15.x systems:	<code>systemctl enable multipathd</code>

6. Rebuild the `initramfs` image or the `initrd` image under `/boot` directory:

If you are using....	Do this...
RHEL 7.x and 8.x systems:	<code>dracut --force --add multipath</code>
SLES 12.x and 15.x systems:	<code>dracut --force --add multipath</code>

7. Use the "Create host manually" procedure in the online help to check whether the hosts are defined. Verify that each host type is either **Linux DM-MP (Kernel 3.10 or later)** if you enable the Automatic Load Balancing feature, or **Linux DM-MP (Kernel 3.9 or earlier)** if you disable the Automatic Load Balancing feature. If necessary, change the selected host type to the appropriate setting.
8. Reboot the host.

Setting up the multipath.conf file

The `multipath.conf` file is the configuration file for the multipath daemon, `multipathd`. The `multipath.conf` file overrides the built-in configuration table for `multipathd`. Any line in the file whose first non-white-space character is `#` is considered a comment line. Empty lines are ignored.

Note: For SANtricity operating system 8.30 and newer, NetApp recommends using the default settings as provided.

The `multipath.conf` are available in the following locations:

- For SLES, `/usr/share/doc/packages/multipath-tools/multipath.conf.synthetic`
- For RHEL, `/usr/share/doc/device-mapper-multipath-0.4.9/multipath.conf`

Configure the switches

You configure the switches according to the vendor's recommendations for iSCSI. These recommendations might include both configuration directives as well as code updates.

You must ensure the following:

- You have two separate networks for high availability. Make sure that you isolate your iSCSI traffic to separate network segments.
- You must enable flow control **end to end**.
- If appropriate, you have enabled jumbo frames.

Note: Port channels/LACP is not supported on the controller's switch ports. Host-side LACP is not recommended; multipathing provides the same, and in some cases better, benefits.

Configure networking

You can set up your iSCSI network in many ways, depending on your data storage requirements.

Consult your network administrator for tips on selecting the best configuration for your environment.

To configure an iSCSI network with basic redundancy, connect each host port and one port from each controller to separate switches, and partition each set of host ports and controller ports on separate network segments or VLANs.

You must enable send and receive hardware flow control **end to end**. You must disable priority flow control.

If you are using jumbo frames within the IP SAN for performance reasons, make sure to configure the array, switches, and hosts to use jumbo frames. Consult your operating system and switch documentation for information on how to enable jumbo frames on the hosts and on the switches. To enable jumbo frames on the array, complete the steps in *Configuring array-side networking—iSCSI*.

Note: Many network switches have to be configured above 9,000 bytes for IP overhead. Consult your switch documentation for more information.

Configure array-side networking

You use the SANtricity System Manager GUI to configure iSCSI networking on the array side.

Before you begin

- You must know the IP address or domain name for one of the storage array controllers.
- You or your system administrator must have set up a password for the System Manager GUI, or you must configured Role-Based Access Control (RBAC) or LDAP and a directory service for the appropriate security access to the storage array. See the *SANtricity System Manager online help* for more information about Access Management.

About this task

This task describes how to access the iSCSI port configuration from the Hardware page. You can also access the configuration from **System > Settings > Configure iSCSI ports**.

Steps

1. From your browser, enter the following URL:

https://<DomainNameOrIPAddress>

IPAddress is the address for one of the storage array controllers.

The first time SANtricity System Manager is opened on an array that has not been configured, the **Set Administrator Password** prompt appears. Role-based access management configures four local roles: admin, support, security, and monitor. The latter three roles have random passwords that cannot be guessed. After you set a password for the admin role you can change all of the passwords using the admin credentials. See *SANtricity System Manager online help* for more information on the four local user roles.

2. Enter the System Manager password for the admin role in the **Set Administrator Password** and **Confirm Password** fields, and then select the **Set Password** button. When you open System Manager and no pools, volumes groups, workloads, or notifications have been configured, the Setup wizard launches.
3. Close the Setup wizard.
You will use the wizard later to complete additional setup tasks.
4. Select **Hardware**.
5. If the graphic shows the drives, click **Show back of shelf**.
The graphic changes to show the controllers instead of the drives.
6. Click the controller with the iSCSI ports you want to configure.
The controller's context menu appears.
7. Select **Configure iSCSI ports**.
The Configure iSCSI Ports dialog box opens.
8. In the drop-down list, select the port you want to configure, and then click **Next**.
9. Select the configuration port settings, and then click **Next**.

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

Port Setting	Description
Configured ethernet port speed	<p>Select the desired speed.</p> <p>The options that appear in the drop-down list depend on the maximum speed that your network can support (for example, 10 Gbps).</p> <p>Note: The optional 25Gb iSCSI host interface cards available on the controllers do not auto-negotiate speeds. You must set the speed for each port to either 10 Gb or 25 Gb. All ports must be set to the same speed.</p>
Enable IPv4 / Enable IPv6	Select one or both options to enable support for IPv4 and IPv6 networks.
TCP listening port (Available by clicking Show more port settings .)	<p>If necessary, enter a new port number.</p> <p>The listening port is the TCP port number that the controller uses to listen for iSCSI logins from host iSCSI initiators. The default listening port is 3260. You must enter 3260 or a value between 49152 and 65535.</p>
MTU size (Available by clicking Show more port settings .)	<p>If necessary, enter a new size in bytes for the Maximum Transmission Unit (MTU).</p> <p>The default Maximum Transmission Unit (MTU) size is 1500 bytes per frame. You must enter a value between 1500 and 9000.</p>

Port Setting	Description
Enable ICMP PING responses	Select this option to enable the Internet Control Message Protocol (ICMP). The operating systems of networked computers use this protocol to send messages. These ICMP messages determine whether a host is reachable and how long it takes to get packets to and from that host.

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.

10. Configure the IPv4 and/or IPv6 settings, either automatically or manually. To see all port settings, click the **Show more settings** link on the right of the dialog box.

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. For IPv4, include the network subnet mask and gateway. For IPv6, include the routable IP address and router IP address.

11. Click **Finish**.
12. Close System Manager.

Configure host-side networking

You configure iSCSI networking on the host side by setting the number of node sessions per physical path, turning on the appropriate iSCSI services, configuring the network for the iSCSI ports, creating iSCSI face bindings, and establishing the iSCSI sessions between initiators and targets.

About this task

In most cases, you can use the inbox software-initiator for iSCSI CNA/NIC. You do not need to download the latest driver, firmware, and BIOS. Refer to the [NetApp Interoperability Matrix Tool](#) to determine code requirements.

Steps

1. Check the `node.session.nr_sessions` variable in the `/etc/iscsi/iscsid.conf` file to see the default number of sessions per physical path. If necessary, change the default number of sessions to one session.

```
node.session.nr_sessions = 1
```

2. Change the `node.session.timeo.replacement_timeout` variable in the `/etc/iscsi/iscsid.conf` file to 20, from a default value of 120.

```
node.session.timeo.replacement_timeout=20
```

3. Make sure `iscsid` and `(open-)iscsi` services are on and enabled for boot.

Red Hat Enterprise Linux 7 and 8 (RHEL 7 and RHEL 8)

```
# systemctl start iscsi
# systemctl start iscsid
# systemctl enable iscsi
# systemctl enable iscsid
```

SUSE Linux Enterprise Server 12 and 15 (SLES 12 and SLES 15)

```
# systemctl start iscsid.service
# systemctl enable iscsid.service
```

Optionally, you set `node.startup = automatic` in `/etc/iscsi/iscsid.conf` before running any `iscsiadm` commands to have sessions persist after reboot:

4. Get the host IQN initiator name, which will be used to configure the host to an array.

```
# cat /etc/iscsi/initiatorname.iscsi
```

5. Configure the network for iSCSI ports:

Note: In addition to the public network port, iSCSI initiators should use two NICs or more on separate private segments or VLANs.

- a. Determine the iSCSI port names using the `# ifconfig -a` command.
- b. Set the IP address for the iSCSI initiator ports. The initiator ports should be present on the same subnet as the iSCSI target ports.

```
# vim /etc/sysconfig/network-scripts/ifcfg-<NIC port>
Edit:
BOOTPROTO=none
ONBOOT=yes
NM_CONTROLLED=no
Add:
IPADDR=192.168.xxx.xxx
NETMASK=255.255.255.0
```

Note: Be sure to set the address for both iSCSI initiator ports.

- c. Restart network services.

```
# systemctl restart network
```

- d. Make sure the Linux server can ping *all* of the iSCSI target ports.

6. Configure the iSCSI interfaces by creating two iSCSI iface bindings.

```
iscsiadm -m iface -I iface0 -o new
iscsiadm -m iface -I iface0 -o update -n iface.net_ifacename -v <NIC port1>
```

```
iscsiadm -m iface -I iface1 -o new
iscsiadm -m iface -I iface1 -o update -n iface.net_ifacename -v <NIC port2>
```

Note: To list the interfaces, use `iscsiadm -m iface`.

7. Establish the iSCSI sessions between initiators and targets (four total).

- a. Discover iSCSI targets. Save the IQN (it will be the same with each discovery) in the worksheet for the next step.

```
iscsiadm -m discovery -t sendtargets -p 192.168.0.1:3260 -I iface0 -P 1
```

Note: The IQN looks like the following:

```
iqn.1992-01.com.netapp:2365.60080e50001bf1600000000531d7be3
```

- b. Create the connection between the iSCSI initiators and iSCSI targets, using ifaces.

```
iscsiadm -m node -T ign.1992-01.com.netapp:2365.60080e50001bf1600000000531d7be3  
-p 192.168.0.1:3260 -I iface0 -l
```

- c. List the iSCSI sessions established on the host.

```
# iscsiadm -m session
```

Verify IP network connections

You verify Internet Protocol (IP) network connections by using ping tests to ensure the host and array are able to communicate.

Steps

1. On the host, run one of the following commands, depending on whether jumbo frames are enabled:

- If jumbo frames are not enabled, run this command:

```
ping -I <hostIP> <targetIP>
```

- If jumbo frames are enabled, run the `ping` command with a payload size of 8,972 bytes. The IP and ICMP combined headers are 28 bytes, which when added to the payload, equals 9,000 bytes. The `-s` switch sets the `packet size` bit. The `-d` switch sets the debug option. These options allow jumbo frames of 9,000 bytes to be successfully transmitted between the iSCSI initiator and the target.

```
ping -I <hostIP> -s 8972 -d <targetIP>
```

In this example, the iSCSI target IP address is 192.0.2.8.

```
#ping -I 192.0.2.100 -s 8972 -d 192.0.2.8  
Pinging 192.0.2.8 with 8972 bytes of data:  
Reply from 192.0.2.8: bytes=8972 time=2ms TTL=64  
Reply from 192.0.2.8: bytes=8972 time=2ms TTL=64  
Reply from 192.0.2.8: bytes=8972 time=2ms TTL=64  
Reply from 192.0.2.8: bytes=8972 time=2ms TTL=64  
Ping statistics for 192.0.2.8:  
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),  
    Approximate round trip times in milli-seconds:  
        Minimum = 2ms, Maximum = 2ms, Average = 2ms
```

2. Issue a `ping` command from each host's initiator address (the IP address of the host Ethernet port used for iSCSI) to each controller iSCSI port. Perform this action from each host server in the configuration, changing the IP addresses as necessary.

Note: If the command fails (for example, returns `Packet needs to be fragmented but DF set`), verify the MTU size (jumbo frame support) for the Ethernet interfaces on the host server, storage controller, and switch ports.

Create partitions and filesystems

A new LUN has no partition or file system when the Linux host first discovers it. You must format the LUN before it can be used. Optionally, you can create a file system on the LUN.

Before you begin

The host must have discovered the LUN.

In the `/dev/mapper` folder, you have run the `ls` command to see the available disks.

About this task

You can initialize the disk as a basic disk with a GUID partition table (GPT) or Master boot record (MBR).

Format the LUN with a file system such as ext4. Some applications do not require this step.

Steps

1. Retrieve the SCSI ID of the mapped disk by issuing the `sanlun lun show -p` command.

The SCSI ID is a 33-character string of hexadecimal digits, beginning with the number 3. If user-friendly names are enabled, Device Mapper reports disks as `mpath` instead of by a SCSI ID.

```
# sanlun lun show -p

E-Series Array: ictml619s01c01-SRP(60080e50002908b40000000054efb9d2)
Volume Name:
Preferred Owner: Controller in Slot B
Current Owner: Controller in Slot B
Mode: RDAC (Active/Active)
UTM LUN: None
LUN: 116
LUN Size:
Product: E-Series
Host Device: mpathr(360080e50004300ac000007575568851d)
Multipath Policy: round-robin 0
Multipath Provider: Native
-----
host      controller
path      path      /dev/      host      controller
state     type      node      adapter   target
-----
up        secondary sdcx      host14     A1
up        secondary sdat      host10     A2
up        secondary sdbv      host13     B1
```

2. Create a new partition according to the method appropriate for your Linux OS release.

Typically, characters identifying the partition of a disk are appended to the SCSI ID (the number 1 or p3 for instance).

```
# parted -a optimal -s -- /dev/mapper/360080e5000321bb8000092b1535f887a mklabel
gpt mkpart primary ext4 0% 100%
```

3. Create a file system on the partition.

The method for creating a file system varies depending on the file system chosen.

```
# mkfs.ext4 /dev/mapper/360080e5000321bb8000092b1535f887a1
```

4. Create a folder to mount the new partition.

```
# mkdir /mnt/ext4
```

5. Mount the partition.

```
# mount /dev/mapper/360080e5000321bb8000092b1535f887a1 /mnt/ext4
```

Verify storage access on the host

Before using the volume, you verify that the host can write data to the volume and read it back.

Before you begin

You must have initialized the volume and formatted it with a file system.

Steps

1. On the host, copy one or more files to the mount point of the disk.
2. Copy the files back to a different folder on the original disk.
3. Run the `diff` command to compare the copied files to the originals.

After you finish

Remove the file and folder that you copied.

Record iSCSI-specific information for Linux


Select the iSCSI worksheet to record your protocol-specific storage configuration information. You need this information to perform provisioning tasks.

iSCSI worksheet - Linux

You can use this worksheet to record iSCSI storage configuration information. You need this information to perform provisioning tasks.

Recommended configuration

Recommended configurations consist of two initiator ports and four target ports with one or more VLANs.



Target IQN


Callout No.	Target port connection	IQN
2	Target port	

Mappings host name

Callout No.	Host information	Name and type
1	Mappings host name	
	Host OS type	

Recommended configuration

Recommended configurations consist of two initiator ports and four target ports with one or more VLANs.



iSER over InfiniBand Express Setup

Steps

1. [Verify the Linux configuration is supported](#) on page 39
2. [Configure IP addresses using DHCP](#) on page 39
3. [Configure subnet manager](#) on page 40
4. [Install and configure Host Utilities](#) on page 41
5. [Install SANtricity Storage Manager for SMcli \(SANtricity software version 11.53 or earlier\)](#) on page 41
6. [Access SANtricity System Manager and use the Setup wizard](#) on page 42
7. [Configure the multipath software](#) on page 43
8. [Setting up the multipath.conf file](#) on page 44
9. [Configure network connections—iSER over Infiniband](#) on page 44
10. [Configure networking for storage attached hosts—iSER over Infiniband](#) on page 45
11. [Create partitions and filesystems](#) on page 47
12. [Verify storage access on the host](#) on page 48
13. [Record iSER over InfiniBand-specific information for Linux](#) on page 48

Verify the Linux configuration is supported

To ensure reliable operation, you create an implementation plan and then use the NetApp Interoperability Matrix Tool (IMT) to verify that the entire configuration is supported.

Steps

1. Go to the [NetApp Interoperability Matrix Tool](#).
2. Click on the **Solution Search** tile.
3. In the **Protocols > SAN Host** area, click the **Add** button next to **E-Series SAN Host**.
4. Click **View Refine Search Criteria**.
The **Refine Search Criteria** section is displayed. In this section you may select the protocol that applies, as well as other criteria for the configuration such as Operating System, NetApp OS, and Host Multipath driver. Select the criteria you know you want for your configuration, and then see what compatible configuration elements apply. As necessary, make the updates for your operating system and protocol that are prescribed in the tool. Detailed information for your chosen configuration is accessible on the View Supported Configurations page by clicking the **right page arrow**.

Configure IP addresses using DHCP

In this express method for configuring communications between the management station and the storage array, you use Dynamic Host Configuration Protocol (DHCP) to provide IP addresses. Each storage array has either one controller (simplex) or two controllers (duplex), and each controller has two storage management ports. Each management port will be assigned an IP address.

Before you begin

You have installed and configured a DHCP server on the same subnet as the storage management ports.

About this task

The following instructions refer to a storage array with two controllers (a duplex configuration).

1. If you have not already done so, connect an Ethernet cable to the management station and to management port 1 on each controller (A and B).

The DHCP server assigns an IP address to port 1 of each controller.

Note: Do not use management port 2 on either controller. Port 2 is reserved for use by NetApp technical personnel.

Important: If you disconnect and reconnect the Ethernet cable, or if the storage array is power-cycled, DHCP assigns IP addresses again. This process occurs until static IP addresses are configured. It is recommended that you avoid disconnecting the cable or power-cycling the array.

If the storage array cannot get DHCP-assigned IP addresses within 30 seconds, the following default IP addresses are set:

- Controller A, port 1: 169.254.128.101
 - Controller B, port 1: 169.254.128.102
 - Subnet mask: 255.255.0.0
2. Locate the MAC address label on the back of each controller, and then provide your network administrator with the MAC address for port 1 of each controller.
Your network administrator needs the MAC addresses to determine the IP address for each controller. You will need the IP addresses to connect to your storage system through your browser.

Configure subnet manager

Using an InfiniBand switch to run subnet manager might cause unexpected path loss during high loads. To avoid path loss, configure the subnet manager on one or more of your hosts using opensm.

Steps

1. Install the `opensm` package on any hosts that will be running the subnet manager.
2. Start and enable the `opensm` service.
3. Use the `ibstat -p` command to find GUID0 and GUID1 of the HBA ports. For example:

```
# ibstat -p
0x248a070300a80a80
0x248a070300a80a81
```

4. Start two instances of the subnet manager, one for each subnet, by adding the following commands to `/etc/rc.d/after.local` for SUSE or `/etc/rc.d/rc.local` for Redhat. Substitute the values you found in the last step for GUID0 and GUID1. For P0 and P1, use the subnet manager priorities, with 1 being the lowest and 15 the highest:

```
opensm -B -g GUID0 -p P0 -f /var/log/opensm-ib0.log
opensm -B -g GUID1 -p P1 -f /var/log/opensm-ib1.log
```


Install and configure Host Utilities

Linux Unified Host Utilities 7.1 includes tools to manage NetApp storage, including failover policies and physical paths.

Steps

1. Use the [NetApp Interoperability Matrix Tool](#) to determine the appropriate version of Unified Host Utilities 7.1 to install.
The versions are listed in a column within each supported configuration.
2. Download the Unified Host Utilities 7.1 from [NetApp Support](#).

Note: Alternatively, you can use the SANtricity `SMdevices` utility to perform the same functions as the Unified Host Utility tool. The `SMdevices` utility is included as part of the `SMutils` package. The `SMutils` package is a collection of utilities to verify what the host sees from the storage array. It is included as part of the SANtricity software installation.

Install SANtricity Storage Manager for SMcli (SANtricity software version 11.53 or earlier)

When you install the SANtricity Storage Manager software on your management station, a graphical user interface (GUI) and a command line interface (CLI) are installed by default. These instructions assume that you will install the SANtricity Storage Manager GUI on a management station and *not* on the I/O host.

Before you begin

Important: For SANtricity software 11.60 and newer, the SANtricity Secure CLI (SMcli) is included in the SANtricity OS and downloadable through the SANtricity System Manager. For more information on how to download the SMcli through the SANtricity System Manager, refer to the *Download the command line interface (CLI)* topic under the SANtricity System Manager Online Help.

- You are using SANtricity software 11.53 or earlier.
- You must have the correct administrator or superuser privileges.
- You must have ensured that the system that will contain the SANtricity Storage Manager client has the following minimum requirements:
 - **RAM:** 2 GB for Java Runtime Engine
 - **Disk space:** 5 GB
 - **OS/Architecture:** Refer to [NetApp Support Downloads > Software > E-Series/EF-Series SANtricity Storage Manager](#) for guidance on determining the supported operating system versions and architectures.

About this task

This section describes how to install SANtricity Storage Manager on both the Windows and Linux OS platforms, because both Windows and Linux are common management station platforms when Linux is used for the data host.

Steps

1. Download the SANtricity software release from [NetApp Support Downloads > Software > E-Series/EF-Series SANtricity Storage Manager](#).
2. Run the SANtricity installer.

Windows	Linux
Double-click the <code>SMIA*.exe</code> installation package to start the installation.	<ol style="list-style-type: none"> Go to the directory where the <code>SMIA*.bin</code> installation package is located. If the temp mount point does not have execute permissions, set the <code>IATEMPDIR</code> variable. Example: <code>IATEMPDIR=/root ./SMIA-LINUX64-11.25.0A00.0002.bin</code> Run the <code>chmod +x SMIA*.bin</code> command to grant execute permission to the file. Run the <code>./SMIA*.bin</code> command to start the installer.

- Use the installation wizard to install the software on the management station.

Access SANtricity System Manager and use the Setup wizard

You use the Setup wizard in SANtricity System Manager to configure your storage array.

Before you begin

- You have ensured that the device from which you will access SANtricity System Manager contains one of the following browsers:

Browser	Minimum version
Google Chrome	47
Microsoft Internet Explorer	11
Microsoft Edge	EdgeHTML 12
Mozilla Firefox	31
Safari	9

- You are using out-of-band management.

About this task

The wizard automatically relaunches when you open System Manager or refresh your browser and *at least one* of the following conditions is met:

- No pools and volume groups are detected.
- No workloads are detected.
- No notifications are configured.

Steps

- From your browser, enter the following URL:

`https://<DomainNameOrIPAddress>`

`IPAddress` is the address for one of the storage array controllers.

The first time SANtricity System Manager is opened on an array that has not been configured, the `Set Administrator Password` prompt appears. Role-based access management configures four local roles: admin, support, security, and monitor. The latter three roles have random passwords that cannot be guessed. After you set a password for the admin role you can change all of the passwords using the admin credentials. See *SANtricity System Manager online help* for more information on the four local user roles.

- Enter the System Manager password for the admin role in the `Set Administrator Password` and `Confirm Password` fields, and then select the **Set Password** button.

When you open System Manager and no pools, volumes groups, workloads, or notifications have been configured, the Setup wizard launches.

3. Use the Setup wizard to perform the following tasks:

- **Verify hardware (controllers and drives)** – Verify the number of controllers and drives in the storage array. Assign a name to the array.
- **Verify hosts and operating systems** – Verify the host and operating system types that the storage array can access.
- **Accept pools** – Accept the recommended pool configuration for the express installation method. A pool is a logical group of drives.
- **Configure alerts** – Allow System Manager to receive automatic notifications when a problem occurs with the storage array.
- **Enable AutoSupport** – Automatically monitor the health of your storage array and have dispatches sent to technical support.

4. If you have not already created a volume, create one by going to **Storage > Volumes > Create > Volume**.

For more information, see the online help for SANtricity System Manager.

Configure the multipath software

Multipath software provides a redundant path to the storage array in case one of the physical paths is disrupted. The multipath software presents the operating system with a single virtual device that represents the active physical paths to the storage. The multipath software also manages the failover process that updates the virtual device. You use the device mapper multipath (DM-MP) tool for Linux installations.

Before you begin

You have installed the required packages on your system.

- For Red Hat (RHEL) hosts, verify the packages are installed by running `rpm -q device-mapper-multipath`.
- For SLES hosts, verify the packages are installed by running `rpm -q multipath-tools`.

About this task

By default, DM-MP is disabled in RHEL and SLES. Complete the following steps to enable DM-MP components on the host.

If you have not already installed the operating system, use the media supplied by your operating system vendor.

Steps

1. If a `multipath.conf` file is not already created, run the `# touch /etc/multipath.conf` command.
2. Use the default multipath settings by leaving the `multipath.conf` file blank.
3. Start the multipath service.

```
# systemctl start multipathd
```

4. Save your kernel version by running the `uname -r` command.

```
# uname -r
3.10.0-327.el7.x86_64
```

You will use this information when you assign volumes to the host.

5. Do one of the following to enable the multipathd daemon on boot.

If you are using....	Do this...
RHEL 7.x and 8.x systems:	<code>systemctl enable multipathd</code>
SLES 12.x and 15.x systems:	<code>systemctl enable multipathd</code>

6. Rebuild the `initramfs` image or the `initrd` image under `/boot` directory:

If you are using....	Do this...
RHEL 7.x and 8.x systems:	<code>dracut --force --add multipath</code>
SLES 12.x and 15.x systems:	<code>dracut --force --add multipath</code>

7. Make sure that the newly created `/boot/initramfs-*` image or `/boot/initrd-*` image is selected in the boot configuration file.

For example, for `grub` it is `/boot/grub/menu.lst` and for `grub2` it is `/boot/grub2/menu.cfg`.

8. Use the "Create host manually" procedure in the online help to check whether the hosts are defined. Verify that each host type is either **Linux DM-MP (Kernel 3.10 or later)** if you enable the Automatic Load Balancing feature, or **Linux DM-MP (Kernel 3.9 or earlier)** if you disable the Automatic Load Balancing feature. If necessary, change the selected host type to the appropriate setting.
9. Reboot the host.

Setting up the multipath.conf file

The `multipath.conf` file is the configuration file for the multipath daemon, `multipathd`. The `multipath.conf` file overrides the built-in configuration table for `multipathd`. Any line in the file whose first non-white-space character is `#` is considered a comment line. Empty lines are ignored.

Note: For SANtricity operating system 8.30 and newer, NetApp recommends using the default settings as provided.

The `multipath.conf` are available in the following locations:

- For SLES, `/usr/share/doc/packages/multipath-tools/multipath.conf.synthetic`
- For RHEL, `/usr/share/doc/device-mapper-multipath-0.4.9/multipath.conf`

Configure network connections—iSER over Infiniband

If your configuration uses the iSER over Infiniband protocol, perform the steps in this section.

About this task

When you are using a 56-Gbps HIC with the iSER over Infiniband protocol, additional array configuration is required.

Steps

1. From the **Setup** tab, select **Configure iSCSI Host Ports** to set the storage array iSCSI addresses.
Put the array iSCSI addresses on the same subnet as the host port(s) you will use to create iSCSI sessions. For addresses, see [iSER worksheet](#).
2. From the **Devices** tab, select the storage array and go to **iSER > Manage Settings** to find the IQN.

This information might be necessary when you create iSER sessions from operating systems that do not support send targets discovery. Enter this information in the worksheet, in [iSER worksheet](#).

Configure networking for storage attached hosts—iSER over Infiniband

The InfiniBand OFED driver stack supports running both iSER and SRP simultaneously on the same ports, so no additional hardware is required.

Before you begin

A NetApp recommended OFED is installed on the system. For more information, see the [NetApp Interoperability Matrix Tool](#).

Steps

1. Enable and start iSCSI services on the host(s):

Red Hat Enterprise Linux 7 and 8 (RHEL 7 and RHEL 8)

```
# systemctl start iscsi
# systemctl start iscsid
# systemctl enable iscsi
# systemctl enable iscsid
```

SUSE Linux Enterprise Server 12 and 15 (SLES 12 and SLES 15)

```
# systemctl start iscsid.service
# systemctl enable iscsid.service
```

2. Configure iPoB network interfaces:

- a. Identify the InfiniBand ports that will be used. Document the HW Address (MAC address) of each port.
- b. Configure persistent names for the InfiniBand network interface devices.
- c. Configure the IP address and network information for the iPoB interfaces identified.

The specific interface configuration required might vary depending on the operating system used. Consult your vendor's operating system documentation for specific information on implementation.

- d. Start the IB network interfaces by restarting the networking service or by manually restarting each interface. For example:

```
systemctl restart network
```

- e. Verify connectivity to the target ports. From the host, ping the IP addresses you configured when you configured network connections.

3. Restart services to load the iSER module.

4. Edit the iSCSI settings in `/etc/iscsi/iscsid.conf`.

```
node.startup = automatic replacement_timeout = 20
```

5. Create iSCSI session configurations:

- a. Create iface configuration files for each InfiniBand interface.

Note: The directory location for the iSCSI iface files is operating system dependent. This example is for using Red Hat Enterprise Linux:

```
iscsiadm -m iface -I iser > /var/lib/iscsi/iface-ib0
iscsiadm -m iface -I iser > /var/lib/iscsi/iface-ib1
```

- b. Edit each iface file to set the interface name and initiator IQN. Set the following parameters appropriately for each iface file:

Option	Value
iface.net_ifacename	The interface device name (ex. ib0).
iface.initiatorname	The host initiator IQN documented in the worksheet.

- c. Create iSCSI sessions to the target.

The preferred method to create the sessions is to use the SendTargets discovery method. However, this method does not work on some operating system releases.

Note: Use **Method 2** for RHEL 6.x or SLES 11.3 or later.

- **Method 1 - SendTargets discovery:** Use the SendTargets discovery mechanism to one of the target portal IP addresses. This will create sessions for each of the target portals.

```
iscsiadm -m discovery -t st -p 192.168.130.101 -I iser
```

- **Method 2 - Manual creation:** For each target portal IP address, create a session using the appropriate host interface iface configuration. In this example, interface ib0 is on subnet A and interface ib1 is on subnet B. For these variables, substitute the appropriate value from the worksheet:

- <Target IQN> = storage array Target IQN
- <Target Port IP> = IP address configured on the specified target port

```
# Controller A Port 1
iscsiadm -m node -target <Target IQN> -I iface-ib0 -p <Target Port IP> -l -o new
# Controller B Port 1
iscsiadm -m node -target <Target IQN> -I iface-ib0 -p <Target Port IP> -l -o new
# Controller A Port 2
iscsiadm -m node -target <Target IQN> -I iface-ib1 -p <Target Port IP> -l -o new
# Controller B Port 2
iscsiadm -m node -target <Target IQN> -I iface-ib1 -p <Target Port IP> -l -o new
```

6. Log in to iSCSI sessions.

For each session, run the iscsiadm command to log in to the session.

```
# Controller A Port 1
iscsiadm -m node -target <Target IQN> -I iface -ib0 -p <Target Port IP> -l
# Controller B Port 1
iscsiadm -m node -target <Target IQN> -I iface -ib0 -p <Target Port IP> -l
# Controller A Port 2
iscsiadm -m node -target <Target IQN> -I iface -ib1 -p <Target Port IP> -l
# Controller B Port 2
iscsiadm -m node -target <Target IQN> -I iface -ib1 -p <Target Port IP> -l
```

7. Verify the iSER/iSCSI sessions.

- a. Check the iscsi session status from the host:

```
iscsiadm -m session
```

- b. Check the iscsi session status from the array. From SANtricity System Manager, navigate to the **Storage Array > iSER > View/End Sessions**.

Result

When the OFED/RDMA service starts, the iSER kernel module(s) loads by default when the iSCSI services are running. To complete the iSER connection setup, the iSER module(s) should be loaded. Currently this requires a host reboot.

Create partitions and filesystems

A new LUN has no partition or file system when the Linux host first discovers it. You must format the LUN before it can be used. Optionally, you can create a file system on the LUN.

Before you begin

The host must have discovered the LUN.

In the /dev/mapper folder, you have run the ls command to see the available disks.

About this task

You can initialize the disk as a basic disk with a GUID partition table (GPT) or Master boot record (MBR).

Format the LUN with a file system such as ext4. Some applications do not require this step.

Steps

- 1. Retrieve the SCSI ID of the mapped disk by issuing the `sanlun lun show -p` command.

The SCSI ID is a 33-character string of hexadecimal digits, beginning with the number 3. If user-friendly names are enabled, Device Mapper reports disks as mpath instead of by a SCSI ID.

```
# sanlun lun show -p

E-Series Array: ictml619s01c01-SRP(60080e50002908b40000000054efb9d2)
Volume Name:
Preferred Owner: Controller in Slot B
Current Owner: Controller in Slot B
Mode: RDAC (Active/Active)
UTM LUN: None
LUN: 116
LUN Size:
Product: E-Series
Host Device: mpathr(360080e50004300ac000007575568851d)
Multipath Policy: round-robin 0
Multipath Provider: Native

-----
host      controller
path      path      /dev/      host      controller
state     type      node      adapter   target
-----
up        secondary sdcx      host14     A1
up        secondary sdat      host10     A2
up        secondary sdbv      host13     B1
```

- 2. Create a new partition according to the method appropriate for your Linux OS release.

Typically, characters identifying the partition of a disk are appended to the SCSI ID (the number 1 or p3 for instance).

```
# parted -a optimal -s -- /dev/mapper/360080e5000321bb8000092b1535f887a mklabel
gpt mkpart primary ext4 0% 100%
```

3. Create a file system on the partition.

The method for creating a file system varies depending on the file system chosen.

```
# mkfs.ext4 /dev/mapper/360080e5000321bb8000092b1535f887a1
```

4. Create a folder to mount the new partition.

```
# mkdir /mnt/ext4
```

5. Mount the partition.

```
# mount /dev/mapper/360080e5000321bb8000092b1535f887a1 /mnt/ext4
```

Verify storage access on the host

Before using the volume, you verify that the host can write data to the volume and read it back.

Before you begin

You must have initialized the volume and formatted it with a file system.

Steps

- 1. On the host, copy one or more files to the mount point of the disk.
- 2. Copy the files back to a different folder on the original disk.
- 3. Run the diff command to compare the copied files to the originals.

After you finish

Remove the file and folder that you copied.

Record iSER over InfiniBand-specific information for Linux

Select the iSER over InfiniBand worksheet to record your protocol-specific storage configuration information. You need this information to perform provisioning tasks.

iSER over InfiniBand worksheet

You can use this worksheet to record iSER over Infiniband storage configuration information. You need this information to perform provisioning tasks.

iSER over InfiniBand: Host identifiers

Note: The software initiator IQN is determined during the task, [Configuring storage attached hosts with iSER networking](#).


Locate and document the initiator IQN from each host. For software initiators, the IQN is typically found in the /etc/iscsi/initiatorname.iscsi file.

Callout No.	Host port connections	Software initiator IQN
1	Host (initiator) 1	
n/a		

Callout No.	Host port connections	Software initiator IQN
n/a		
n/a		
n/a		

iSER over InfiniBand: Recommended configuration

Recommended configurations consist of two host (initiator) ports and four target ports.



iSER over InfiniBand: Target IQN

Document the target IQN for the storage array. You will use this information in [Configuring storage attached hosts with iSER networking](#).

Find the Storage Array IQN name using SANtricity: **Storage Array > iSER > Manage Settings**. This information might be necessary when you create iSER sessions from operating systems that do not support send targets discovery.

Callout No.	Array name	Target IQN
6	Array controller (target)	

iSER over InfiniBand: Network configuration

Document the network configuration that will be used for the hosts and storage on the InfiniBand fabric. These instructions assume that two subnets will be used for full redundancy.

Your network administrator can provide the following information. You use this information in the topic, [Configuring storage attached hosts with iSER networking](#).

Subnet A

Define the subnet to be used.

Network Address	Netmask

Document the IQNs to be used by the array ports and each host port.

Callout No.	Array controller (target) port connections	IQN
3	Switch	<i>not applicable</i>
5	Controller A, port 1	
4	Controller B, port 1	
2	Host 1, port 1	
	(Optional) Host 2, port 1	

Subnet B

Define the subnet to be used.

Network Address	Netmask

Document the IQNs to be used by the array ports and each host port.

Callout No.	Array controller (target) port connections	IQN
8	Switch	<i>not applicable</i>
10	Controller A, port 2	
9	Controller B, port 2	
7	Host 1, port 2	
	(Optional) Host 2, port 2	

iSER over InfiniBand: Mapping host name

Note: The mapping host name is created during the workflow.

Mapping host name	
Host OS type	

SRP over InfiniBand Express Setup

Steps

1. [Verify the Linux configuration is supported](#) on page 51
2. [Configure IP addresses using DHCP](#) on page 51
3. [Configure subnet manager](#) on page 52
4. [Install and configure Host Utilities](#) on page 53
5. [Install SANtricity Storage Manager for SMcli \(SANtricity software version 11.53 or earlier\)](#) on page 53
6. [Access SANtricity System Manager and use the Setup wizard](#) on page 54
7. [Configure the multipath software](#) on page 55
8. [Setting up the multipath.conf file](#) on page 56
9. [Determining host port GUIDs and making the recommended settings—SRP over Infiniband](#) on page 56
10. [Configure network connections—SRP over Infiniband](#) on page 57
11. [Create partitions and filesystems](#) on page 58
12. [Verify storage access on the host](#) on page 59
13. [Record SRP over InfiniBand-specific information for Linux](#) on page 59

Verify the Linux configuration is supported

To ensure reliable operation, you create an implementation plan and then use the NetApp Interoperability Matrix Tool (IMT) to verify that the entire configuration is supported.

Steps

1. Go to the [NetApp Interoperability Matrix Tool](#).
2. Click on the **Solution Search** tile.
3. In the **Protocols > SAN Host** area, click the **Add** button next to **E-Series SAN Host**.
4. Click **View Refine Search Criteria**.
The **Refine Search Criteria** section is displayed. In this section you may select the protocol that applies, as well as other criteria for the configuration such as Operating System, NetApp OS, and Host Multipath driver. Select the criteria you know you want for your configuration, and then see what compatible configuration elements apply. As necessary, make the updates for your operating system and protocol that are prescribed in the tool. Detailed information for your chosen configuration is accessible on the View Supported Configurations page by clicking the **right page arrow**.

Configure IP addresses using DHCP

In this express method for configuring communications between the management station and the storage array, you use Dynamic Host Configuration Protocol (DHCP) to provide IP addresses. Each storage array has either one controller (simplex) or two controllers (duplex), and each controller has two storage management ports. Each management port will be assigned an IP address.

Before you begin

You have installed and configured a DHCP server on the same subnet as the storage management ports.

About this task

The following instructions refer to a storage array with two controllers (a duplex configuration).

1. If you have not already done so, connect an Ethernet cable to the management station and to management port 1 on each controller (A and B).

The DHCP server assigns an IP address to port 1 of each controller.

Note: Do not use management port 2 on either controller. Port 2 is reserved for use by NetApp technical personnel.

Important: If you disconnect and reconnect the Ethernet cable, or if the storage array is power-cycled, DHCP assigns IP addresses again. This process occurs until static IP addresses are configured. It is recommended that you avoid disconnecting the cable or power-cycling the array.

If the storage array cannot get DHCP-assigned IP addresses within 30 seconds, the following default IP addresses are set:

- Controller A, port 1: 169.254.128.101
 - Controller B, port 1: 169.254.128.102
 - Subnet mask: 255.255.0.0
2. Locate the MAC address label on the back of each controller, and then provide your network administrator with the MAC address for port 1 of each controller.
Your network administrator needs the MAC addresses to determine the IP address for each controller. You will need the IP addresses to connect to your storage system through your browser.

Configure subnet manager

Using an InfiniBand switch to run subnet manager might cause unexpected path loss during high loads. To avoid path loss, configure the subnet manager on one or more of your hosts using opensm.

Steps

1. Install the opensm package on any hosts that will be running the subnet manager.
2. Start and enable the opensm service.
3. Use the `ibstat -p` command to find GUID0 and GUID1 of the HBA ports. For example:

```
# ibstat -p
0x248a070300a80a80
0x248a070300a80a81
```

4. Start two instances of the subnet manager, one for each subnet, by adding the following commands to `/etc/rc.d/after.local` for SUSE or `etc/rc.d/rc.local` for Redhat. Substitute the values you found in the last step for GUID0 and GUID1. For P0 and P1, use the subnet manager priorities, with 1 being the lowest and 15 the highest:

```
opensm -B -g GUID0 -p P0 -f /var/log/opensm-ib0.log
opensm -B -g GUID1 -p P1 -f /var/log/opensm-ib1.log
```

Install and configure Host Utilities

Linux Unified Host Utilities 7.1 includes tools to manage NetApp storage, including failover policies and physical paths.

Steps

1. Use the [NetApp Interoperability Matrix Tool](#) to determine the appropriate version of Unified Host Utilities 7.1 to install.
The versions are listed in a column within each supported configuration.
2. Download the Unified Host Utilities 7.1 from [NetApp Support](#).

Note: Alternatively, you can use the SANtricity `SMdevices` utility to perform the same functions as the Unified Host Utility tool. The `SMdevices` utility is included as part of the `SMutils` package. The `SMutils` package is a collection of utilities to verify what the host sees from the storage array. It is included as part of the SANtricity software installation.

Install SANtricity Storage Manager for SMcli (SANtricity software version 11.53 or earlier)

When you install the SANtricity Storage Manager software on your management station, a graphical user interface (GUI) and a command line interface (CLI) are installed by default. These instructions assume that you will install the SANtricity Storage Manager GUI on a management station and *not* on the I/O host.

Before you begin

Important: For SANtricity software 11.60 and newer, the SANtricity Secure CLI (SMcli) is included in the SANtricity OS and downloadable through the SANtricity System Manager. For more information on how to download the SMcli through the SANtricity System Manager, refer to the *Download the command line interface (CLI)* topic under the SANtricity System Manager Online Help.

- You are using SANtricity software 11.53 or earlier.
- You must have the correct administrator or superuser privileges.
- You must have ensured that the system that will contain the SANtricity Storage Manager client or host package has the following minimum requirements:
 - **RAM:** 2 GB for Java Runtime Engine
 - **Disk space:** 5 GB
 - **OS/Architecture:** Refer to [NetApp Interoperability Matrix Tool](#) for guidance on determining the supported operating system versions and architectures.

About this task

This section describes how to install SANtricity Storage Manager on both the Windows and Linux OS platforms, because both Windows and Linux are common management station platforms when Linux is used for the data host.

Steps

1. Download the SANtricity software release from [NetApp Support Downloads](#) > **Software** > **E-Series/EF-Series SANtricity Storage Manager**.
2. Run the SANtricity installer.

Windows	Linux
Double-click the <code>SMIA*.exe</code> installation package to start the installation.	<ol style="list-style-type: none"> Go to the directory where the <code>SMIA*.bin</code> installation package is located. If the temp mount point does not have execute permissions, set the <code>IATEMPDIR</code> variable. Example: <code>IATEMPDIR=/root ./SMIA-LINUX64-11.25.0A00.0002.bin</code> Run the <code>chmod +x SMIA*.bin</code> command to grant execute permission to the file. Run the <code>./SMIA*.bin</code> command to start the installer.

- Use the installation wizard to install the software on the management station.

Access SANtricity System Manager and use the Setup wizard

You use the Setup wizard in SANtricity System Manager to configure your storage array.

Before you begin

- You have ensured that the device from which you will access SANtricity System Manager contains one of the following browsers:

Browser	Minimum version
Google Chrome	47
Microsoft Internet Explorer	11
Microsoft Edge	EdgeHTML 12
Mozilla Firefox	31
Safari	9

- You are using out-of-band management.

About this task

The wizard automatically relaunches when you open System Manager or refresh your browser and *at least one* of the following conditions is met:

- No pools and volume groups are detected.
- No workloads are detected.
- No notifications are configured.

Steps

- From your browser, enter the following URL:

`https://<DomainNameOrIPAddress>`

`IPAddress` is the address for one of the storage array controllers.

The first time SANtricity System Manager is opened on an array that has not been configured, the `Set Administrator Password` prompt appears. Role-based access management configures four local roles: admin, support, security, and monitor. The latter three roles have random passwords that cannot be guessed. After you set a password for the admin role you can change all of the passwords using the admin credentials. See *SANtricity System Manager online help* for more information on the four local user roles.

- Enter the System Manager password for the admin role in the `Set Administrator Password` and `Confirm Password` fields, and then select the **Set Password** button.

When you open System Manager and no pools, volumes groups, workloads, or notifications have been configured, the Setup wizard launches.

3. Use the Setup wizard to perform the following tasks:

- **Verify hardware (controllers and drives)** – Verify the number of controllers and drives in the storage array. Assign a name to the array.
- **Verify hosts and operating systems** – Verify the host and operating system types that the storage array can access.
- **Accept pools** – Accept the recommended pool configuration for the express installation method. A pool is a logical group of drives.
- **Configure alerts** – Allow System Manager to receive automatic notifications when a problem occurs with the storage array.
- **Enable AutoSupport** – Automatically monitor the health of your storage array and have dispatches sent to technical support.

4. If you have not already created a volume, create one by going to **Storage > Volumes > Create > Volume**.

For more information, see the online help for SANtricity System Manager.

Configure the multipath software

Multipath software provides a redundant path to the storage array in case one of the physical paths is disrupted. The multipath software presents the operating system with a single virtual device that represents the active physical paths to the storage. The multipath software also manages the failover process that updates the virtual device. You use the device mapper multipath (DM-MP) tool for Linux installations.

Before you begin

You have installed the required packages on your system.

- For Red Hat (RHEL) hosts, verify the packages are installed by running `rpm -q device-mapper-multipath`.
- For SLES hosts, verify the packages are installed by running `rpm -q multipath-tools`.

About this task

By default, DM-MP is disabled in RHEL and SLES. Complete the following steps to enable DM-MP components on the host.

If you have not already installed the operating system, use the media supplied by your operating system vendor.

Steps

1. If a `multipath.conf` file is not already created, run the `# touch /etc/multipath.conf` command.
2. Use the default multipath settings by leaving the `multipath.conf` file blank.
3. Start the multipath service.

```
# systemctl start multipathd
```

4. Save your kernel version by running the `uname -r` command.

```
# uname -r
3.10.0-327.el7.x86_64
```

You will use this information when you assign volumes to the host.

5. Do one of the following to enable the multipathd daemon on boot.

If you are using....	Do this...
RHEL 7.x and 8.x systems:	<code>systemctl enable multipathd</code>
SLES 12.x and 15.x systems:	<code>systemctl enable multipathd</code>

6. Rebuild the `initramfs` image or the `initrd` image under `/boot` directory:

If you are using....	Do this...
RHEL 7.x and 8.x systems:	<code>dracut --force --add multipath</code>
SLES 12.x and 15.x systems:	<code>dracut --force --add multipath</code>

7. Make sure that the newly created `/boot/initrams-*` image or `/boot/initrd-*` image is selected in the boot configuration file.
For example, for `grub` it is `/boot/grub/menu.lst` and for `grub2` it is `/boot/grub2/menu.cfg`.
8. Use the "Create host manually" procedure in the online help to check whether the hosts are defined. Verify that each host type is either **Linux DM-MP (Kernel 3.10 or later)** if you enable the Automatic Load Balancing feature, or **Linux DM-MP (Kernel 3.9 or earlier)** if you disable the Automatic Load Balancing feature. If necessary, change the selected host type to the appropriate setting.
9. Reboot the host.

Setting up the multipath.conf file

The `multipath.conf` file is the configuration file for the multipath daemon, `multipathd`. The `multipath.conf` file overrides the built-in configuration table for `multipathd`. Any line in the file whose first non-white-space character is `#` is considered a comment line. Empty lines are ignored.

Note: For SANtricity operating system 8.30 and newer, NetApp recommends using the default settings as provided.

Example `multipath.conf` are available in the following locations:

- For SLES, `/usr/share/doc/packages/multipath-tools/multipath.conf.synthetic`
- For RHEL, `/usr/share/doc/device-mapper-multipath-0.4.9/multipath.conf`

Determining host port GUIDs and making the recommended settings—SRP over Infiniband

The `InfiniBand-diags` package includes commands to display the globally unique ID (GUID) of each InfiniBand (IB) port. Most Linux distributions with OFED/RDMA supported through the included packages also have the `InfiniBand-diags` package, which includes commands to display information about the HCA.

Steps

1. Install the `InfiniBand-diags` package using the operating system's package management commands.
2. Run the `ibstat` command to display the port information.
3. Record the initiator's GUIDs on the [SRP worksheet](#).
4. Select the appropriate settings in the HBA utility.

Appropriate settings for your configuration are listed in the Notes column of the [NetApp Interoperability Matrix Tool](#).

Configure network connections—SRP over Infiniband

If your configuration uses the SRP over Infiniband protocol, follow the steps in this section.

About this task

To connect the Linux host to the storage array, you must enable the InfiniBand driver stack with the appropriate options. Specific settings might vary between Linux distributions. Check the [NetApp Interoperability Matrix Tool](#) for specific instructions and additional recommended settings specific to your solution.

Steps

1. Install the OFED/RDMA driver stack for your OS.

SLES

```
zypper install rdma-core
```

RHEL

```
yum install rdma-core
```

2. Configure OFED/RDMA to load the SRP module.

SLES

```
zypper install srp_daemon
```

RHEL

```
yum install srp_daemon
```

3. In the OFED/RDMA configuration file, set SRP_LOAD=yes and SRP_DAEMON_ENABLE=yes.

The RDMA configuration file is located at the following location:

```
/etc/rdma/rdma.conf
```

4. Enable and start the OFED/RDMA service.

RHEL 7.x and SLES 12.x or greater

- To enable the InfiniBand modules to load on boot:

```
systemctl enable rdma
```

- To load the InfiniBand modules immediately:

```
systemctl start rdma
```

5. Enable the SRP daemon.

RHEL 7.x and SLES 12 or greater

- To enable the SRP daemon to start on boot:

```
systemctl enable srp_daemon
```

- To start the SRP daemon immediately:

```
systemctl start srp_daemon
```

6. If you need to modify the SRP configuration, enter the following command to create `/etc/modprobe.d/ib_srp.conf`.
- ```
options ib_srp cmd_sg_entries=255 allow_ext_sg=y indirect_sg_entries=2048
```
- a. Under the `/etc/srp_daemon.conf`, add the following line.
- ```
a
    max_sect=4096
```

Create partitions and filesystems

A new LUN has no partition or file system when the Linux host first discovers it. You must format the LUN before it can be used. Optionally, you can create a file system on the LUN.

Before you begin

The host must have discovered the LUN.

In the `/dev/mapper` folder, you have run the `ls` command to see the available disks.

About this task

You can initialize the disk as a basic disk with a GUID partition table (GPT) or Master boot record (MBR).

Format the LUN with a file system such as ext4. Some applications do not require this step.

Steps

1. Retrieve the SCSI ID of the mapped disk by issuing the `sanlun lun show -p` command.
- The SCSI ID is a 33-character string of hexadecimal digits, beginning with the number 3. If user-friendly names are enabled, Device Mapper reports disks as `mpath` instead of by a SCSI ID.

```
# sanlun lun show -p

E-Series Array: ictml619s01c01-SRP(60080e50002908b40000000054efb9d2)
Volume Name:
Preferred Owner: Controller in Slot B
Current Owner: Controller in Slot B
Mode: RDAC (Active/Active)
UTM LUN: None
LUN: 116
LUN Size:
Product: E-Series
Host Device: mpathr(360080e50004300ac000007575568851d)
Multipath Policy: round-robin 0
Multipath Provider: Native

-----
-----
host      controller
path      path      /dev/      host      controller
state     type       node      adapter   target
-----
-----
up        secondary sdcx      host14     A1
up        secondary sdat      host10     A2
up        secondary sdbv      host13     B1
```

2. Create a new partition according to the method appropriate for your Linux OS release.

Typically, characters identifying the partition of a disk are appended to the SCSI ID (the number 1 or p3 for instance).

```
# parted -a optimal -s -- /dev/mapper/360080e5000321bb8000092b1535f887a mklabel  
gpt mkpart primary ext4 0% 100%
```

3. Create a file system on the partition.

The method for creating a file system varies depending on the file system chosen.

```
# mkfs.ext4 /dev/mapper/360080e5000321bb8000092b1535f887a1
```

4. Create a folder to mount the new partition.

```
# mkdir /mnt/ext4
```

5. Mount the partition.

```
# mount /dev/mapper/360080e5000321bb8000092b1535f887a1 /mnt/ext4
```

Verify storage access on the host

Before using the volume, you verify that the host can write data to the volume and read it back.

Before you begin

You must have initialized the volume and formatted it with a file system.

Steps

1. On the host, copy one or more files to the mount point of the disk.
2. Copy the files back to a different folder on the original disk.
3. Run the `diff` command to compare the copied files to the originals.

After you finish


Remove the file and folder that you copied.

Record SRP over InfiniBand-specific information for Linux

Select the SRP over InfiniBand worksheet to record your protocol-specific storage configuration information. You need this information to perform provisioning tasks.

SRP over Infiniband worksheet

You can use this worksheet to record SRP over Infiniband storage configuration information. You need this information to perform provisioning tasks.



SRP over Infiniband: Host identifiers

Note: The initiator GUIDs are determined in the task, [Determining host port GUIDs and making the recommended settings](#).

Callout No.	Host (initiator) port connections	GUID
1	Host	<i>not applicable</i>
3	Switch	<i>not applicable</i>
4	Target (storage array)	<i>not applicable</i>
2	Host port 1 to IB switch 1 ("A" path)	
5	Host port 2 to IB switch 2 ("B" path)	

SRP over Infiniband: Recommended configuration

Recommended configurations consist of two initiator ports and four target ports.

SRP over Infiniband: Mapping host name

Note: The mapping host name is created during the workflow.

Mapping host name	
Host OS type	

NVMe over InfiniBand Express Setup

You can use NVMe with the InfiniBand network protocol.

Steps

1. [Verify the Linux configuration is supported](#) on page 61
2. [Configure IP addresses using DHCP](#) on page 62
3. [Install SANtricity Storage Manager for SMcli \(SANtricity software version 11.53 or earlier\)](#) on page 63
4. [Access SANtricity System Manager and use the Setup wizard](#) on page 64
5. [Configure subnet manager](#) on page 65
6. [Set up NVMe over InfiniBand on the host side](#) on page 66
7. [Configure storage array NVMe over InfiniBand connections](#) on page 69
8. [Discover and connect to the storage from the host](#) on page 69
9. [Define a host](#) on page 72
10. [Assign a volume](#) on page 73
11. [Display the volumes visible to the host](#) on page 74
12. [Set up failover](#) on page 74
13. [Accessing NVMe volumes](#) on page 75
14. [Create filesystems](#) on page 78
15. [Verify storage access on the host](#) on page 80
16. [Record NVMe over InfiniBand-specific information for Linux](#) on page 80

Verify the Linux configuration is supported

To ensure reliable operation, you create an implementation plan and then use the NetApp Interoperability Matrix Tool (IMT) to verify that the entire configuration is supported.

Steps

1. Go to the [NetApp Interoperability Matrix Tool](#).
2. Click on the **Solution Search** tile.
3. In the **Protocols > SAN Host** area, click the **Add** button next to **E-Series SAN Host**.
4. Click **View Refine Search Criteria**.
The **Refine Search Criteria** section is displayed. In this section you may select the protocol that applies, as well as other criteria for the configuration such as Operating System, NetApp OS, and Host Multipath driver. Select the criteria you know you want for your configuration, and then see what compatible configuration elements apply. As necessary, make the updates for your operating system and protocol that are prescribed in the tool. Detailed information for your chosen configuration is accessible on the View Supported Configurations page by clicking the **right page arrow**.

Related information

[NetApp Interoperability Matrix Tool](#)

NVMe over InfiniBand restrictions

Before using NVMe over InfiniBand, review the controller, host, and recovery restrictions.

Verify your configuration

Verify your configuration, using the [NetApp Interoperability Matrix Tool](#).

About the hardware

NVMe over InfiniBand can be configured for EF300 (100GB controllers only), EF600, EF570, or E5700 controllers. The controllers must have 100GB or 200GB InfiniBand host ports.

Restrictions

The following restrictions are in effect for the 11.60 release. See the [NetApp Interoperability Matrix Tool](#) for a complete list of requirements.

Controller restrictions

- This protocol can be used only for EF300, EF600, EF570, or E5700 controllers. A minimum of 32 GB of physical memory is required to use this protocol on EF600, EF570, and E5700 controllers. For the EF300, a minimum of 16 GB of physical memory is required. If the minimum memory requirements for the controllers are not met during start of day operations, a message is displayed that helps you diagnose the problem.
- No simplex (single controller) configurations are supported.
- There is no support for mixed NVMe over InfiniBand and SCSI host interfaces.
- For EF300 controllers, no more than 64 NVMe hosts can be supported on the IB interface.

Host, host protocol and host operating system restrictions

- The host must be running the latest compatible RHEL 7, SUSE Linux Enterprise Server 12 or 15 service pack operating system. See the [NetApp Interoperability Matrix Tool](#) for a complete list of the latest requirements.
- The only supported host channel adapters are from Mellanox. See the [NetApp Interoperability Matrix Tool](#) for more information.
- The only supported host interface card (HIC) is the 100G or 200G EDR IB HIC, which also supports iSER and SRP (but iSER and SRP are not supported simultaneously).

Storage and disaster recovery restrictions

- Asynchronous and synchronous mirroring are not supported.
- Thin provisioning (the creation of thin volumes) is not supported.

Configure IP addresses using DHCP

In this express method for configuring communications between the management station and the storage array, you use Dynamic Host Configuration Protocol (DHCP) to provide IP addresses. Each storage array has either one controller (simplex) or two controllers (duplex), and each controller has two storage management ports. Each management port will be assigned an IP address.

Before you begin

You have installed and configured a DHCP server on the same subnet as the storage management ports.

About this task

The following instructions refer to a storage array with two controllers (a duplex configuration).

1. If you have not already done so, connect an Ethernet cable to the management station and to management port 1 on each controller (A and B).

The DHCP server assigns an IP address to port 1 of each controller.

Note: Do not use management port 2 on either controller. Port 2 is reserved for use by NetApp technical personnel.

Important: If you disconnect and reconnect the Ethernet cable, or if the storage array is power-cycled, DHCP assigns IP addresses again. This process occurs until static IP addresses

are configured. It is recommended that you avoid disconnecting the cable or power-cycling the array.

If the storage array cannot get DHCP-assigned IP addresses within 30 seconds, the following default IP addresses are set:

- Controller A, port 1: 169.254.128.101
 - Controller B, port 1: 169.254.128.102
 - Subnet mask: 255.255.0.0
2. Locate the MAC address label on the back of each controller, and then provide your network administrator with the MAC address for port 1 of each controller.
- Your network administrator needs the MAC addresses to determine the IP address for each controller. You will need the IP addresses to connect to your storage system through your browser.

Install SANtricity Storage Manager for SMcli (SANtricity software version 11.53 or earlier)

When you install the SANtricity Storage Manager software on your management station, a graphical user interface (GUI) and a command line interface (CLI) are installed by default. These instructions assume that you will install the SANtricity Storage Manager GUI on a management station and *not* on the I/O host.

Before you begin

Important: For SANtricity software 11.60 and newer, the SANtricity Secure CLI (SMcli) is included in the SANtricity OS and downloadable through the SANtricity System Manager. For more information on how to download the SMcli through the SANtricity System Manager, refer to the *Download the command line interface (CLI)* topic under the SANtricity System Manager Online Help.

- You are using SANtricity software 11.53 or earlier.
- You must have the correct administrator or superuser privileges.
- You must have ensured that the system that will contain the SANtricity Storage Manager client has the following minimum requirements:
 - **RAM:** 2 GB for Java Runtime Engine
 - **Disk space:** 5 GB
 - **OS/Architecture:** Refer to [NetApp Support Downloads > Software > E-Series/EF-Series SANtricity Storage Manager](#) for guidance on determining the supported operating system versions and architectures.

About this task

This section describes how to install SANtricity Storage Manager on both the Windows and Linux OS platforms, because both Windows and Linux are common management station platforms when Linux is used for the data host.

Steps

1. Download the SANtricity software release from [NetApp Support Downloads > Software > E-Series/EF-Series SANtricity Storage Manager](#).
2. Run the SANtricity installer.

Windows	Linux
Double-click the <code>SMIA*.exe</code> installation package to start the installation.	<ol style="list-style-type: none"> Go to the directory where the <code>SMIA*.bin</code> installation package is located. If the temp mount point does not have execute permissions, set the <code>IATEMPDIR</code> variable. Example: <code>IATEMPDIR=/root ./SMIA-LINUX64-11.25.0A00.0002.bin</code> Run the <code>chmod +x SMIA*.bin</code> command to grant execute permission to the file. Run the <code>./SMIA*.bin</code> command to start the installer.

- Use the installation wizard to install the software on the management station.

Access SANtricity System Manager and use the Setup wizard

You use the Setup wizard in SANtricity System Manager to configure your storage array.

Before you begin

- You have ensured that the device from which you will access SANtricity System Manager contains one of the following browsers:

Browser	Minimum version
Google Chrome	47
Microsoft Internet Explorer	11
Microsoft Edge	EdgeHTML 12
Mozilla Firefox	31
Safari	9

- You are using out-of-band management.

About this task

The wizard automatically relaunches when you open System Manager or refresh your browser and *at least one* of the following conditions is met:

- No pools and volume groups are detected.
- No workloads are detected.
- No notifications are configured.

Steps

- From your browser, enter the following URL:

`https://<DomainNameOrIPAddress>`

`IPAddress` is the address for one of the storage array controllers.

The first time SANtricity System Manager is opened on an array that has not been configured, the `Set Administrator Password` prompt appears. Role-based access management configures four local roles: admin, support, security, and monitor. The latter three roles have random passwords that cannot be guessed. After you set a password for the admin role you can change all of the passwords using the admin credentials. See *SANtricity System Manager online help* for more information on the four local user roles.

2. Enter the System Manager password for the admin role in the `Set Administrator Password` and `Confirm Password` fields, and then select the **Set Password** button. When you open System Manager and no pools, volumes groups, workloads, or notifications have been configured, the Setup wizard launches.
3. Use the Setup wizard to perform the following tasks:
 - **Verify hardware (controllers and drives)** – Verify the number of controllers and drives in the storage array. Assign a name to the array.
 - **Verify hosts and operating systems** – Verify the host and operating system types that the storage array can access.
 - **Accept pools** – Accept the recommended pool configuration for the express installation method. A pool is a logical group of drives.
 - **Configure alerts** – Allow System Manager to receive automatic notifications when a problem occurs with the storage array.
 - **Enable AutoSupport** – Automatically monitor the health of your storage array and have dispatches sent to technical support.
4. If you have not already created a volume, create one by going to **Storage > Volumes > Create > Volume**.
For more information, see the online help for SANtricity System Manager.

Configure subnet manager

Using an InfiniBand switch to run subnet manager might cause unexpected path loss during high loads. To avoid path loss, configure the subnet manager on one or more of your hosts using opensm.

Before you begin

- You are running the latest compatible RHEL 7, SUSE Linux Enterprise Server 12 or 15 service pack operating system. See the [NetApp Interoperability Matrix Tool](#) for a complete list of the latest requirements.

Steps

1. Use the `ibstat -p` command to find GUID0 and GUID1 of the HCA ports. For example:

```
# ibstat -p
0x248a070300a80a80
0x248a070300a80a81
```

2. The way that you configure Subnet Manager depends on your configuration:
 - If you are using a single switch, start and enable the `opensm` service, then add the HCA port identifier values you found in [step 2](#) to the `opensm.conf` file on each port. Repeat for the other port.
 - Edit the `/etc/rdma/opensm.conf` file to add the identifier for that port:

```
opensm -c /etc/rdma/opensm.conf

# The port GUID on which the OpenSM is running
guid 0x248a070300a80a80
```
 - If you are using the direct connect method, or if you have multiple switches, enable Subnet Manager on each port of the connected HCA on the host:
 - Add the following two lines to `/etc/rc.d/after.local` (for SUSE Linux Enterprise Server 12 and SLES 15 service pack). Substitute the values you found in [step 2](#) for

GUID0 and GUID1. For P0 and P1, use the subnet manager priorities, with 1 being the lowest and 15 the highest:

SLES example

```
opensm -B -g GUID0 -p P0 -f /var/log/opensm-ib0.log
opensm -B -g GUID1 -p P1 -f /var/log/opensm-ib1.log
```

An example of the command with value substitutions.

```
# cat /etc/rc.d/rc.local
opensm -B -g 0x248a070300a80a80 -p 15 -f /var/log/opensm-ib0.log
opensm -B -g 0x248a070300a80a81 -p 1 -f /var/log/opensm-ib1.log
```

- Add the following two lines to `/etc/rc.d/rc.local` (for RHEL 7). Substitute the values you found in [step 2](#) for GUID0 and GUID1. For P0 and P1, use the subnet manager priorities, with 1 being the lowest and 15 the highest:

RHEL example

```
opensm -B -g GUID0 -p P0 -f /var/log/opensm-ib0.log
opensm -B -g GUID1 -p P1 -f /var/log/opensm-ib1.log
```

An example of the command with value substitutions.

```
# cat /etc/rc.d/rc.local
opensm -B -g 0x248a070300a80a80 -p 15 -f /var/log/opensm-ib0.log
opensm -B -g 0x248a070300a80a81 -p 1 -f /var/log/opensm-ib1.log
```

Set up NVMe over InfiniBand on the host side

Configuring an NVMe initiator in an InfiniBand environment includes installing and configuring the `infiniband`, `nvme-cli`, and `rdma` packages, configuring initiator IP addresses, and setting up the NVMe-oF layer on the host.

Before you begin

- You are running the latest compatible RHEL 7, SUSE Linux Enterprise Server 12 and 15 service pack operating system. See the [NetApp Interoperability Matrix Tool](#) for a complete list of the latest requirements.

Steps

1. Install the `rdma`, `nvme-cli`, and `infiniband` packages:

```
# zypper install infiniband-diags
# zypper install rdma-core
# zypper install nvme-cli
```

RHEL 7

```
# yum install infiniband-diags
# yum install rdma-core
# yum install nvme-cli
```

2. Enable `ipoib`. Edit the `/etc/rdma/rdma.conf` file and modify the entry for loading `ipoib`:

```
IPOIB_LOAD=yes
```

3. Check that both ib port links are up and the State = Active:

```
# ibstat

CA 'mlx4_0'
  CA type: MT4099
  Number of ports: 2
  Firmware version: 2.40.7000
  Hardware version: 1
  Node GUID: 0x0002c90300317850
  System image GUID: 0x0002c90300317853
  Port 1:
    State: Active
    Physical state: LinkUp
    Rate: 40
    Base lid: 4
    LMC: 0
    SM lid: 4
    Capability mask: 0x0259486a
    Port GUID: 0x0002c90300317851
    Link layer: InfiniBand
  Port 2:
    State: Active
    Physical state: LinkUp
    Rate: 56
    Base lid: 5
    LMC: 0
    SM lid: 4
    Capability mask: 0x0259486a
    Port GUID: 0x0002c90300317852
    Link layer: InfiniBand
```

4. Set up IPv4 IP addresses on the ib ports.

For SUSE Linux Enterprise Server 12 and 15, create the file `/etc/sysconfig/network/ifcfg-ib0`

```
BOOTPROTO='static'
BROADCAST=
ETHTOOL_OPTIONS=
IPADDR='10.10.10.100/24'
IPOIB_MODE='connected'
MTU='65520'
NAME=
NETWORK=
REMOTE_IPADDR=
STARTMODE='auto'
```

Then, create the file `/etc/sysconfig/network/ifcfg-ib1`.

```
BOOTPROTO='static'
BROADCAST=
ETHTOOL_OPTIONS=
IPADDR='11.11.11.100/24'
IPOIB_MODE='connected'
MTU='65520'
NAME=
NETWORK=
```

```
REMOTE_IPADDR=  
STARTMODE='auto'
```

For RHEL, create the file /etc/sysconfig/network-scripts/ifcfg-ib0:

```
CONNECTED_MODE=no  
TYPE=InfiniBand  
PROXY_METHOD=none  
BROWSER_ONLY=no  
BOOTPROTO=static  
IPADDR='10.10.10.100/24'  
DEFROUTE=no  
IPV4=FAILURE_FATAL=yes  
IPV6INIT=no  
NAME=ib0  
ONBOOT=yes
```

Then, create the file /etc/sysconfig/network-scripts/ifcfg-ib1:.

```
CONNECTED_MODE=no  
TYPE=InfiniBand  
PROXY_METHOD=none  
BROWSER_ONLY=no  
BOOTPROTO=static  
IPADDR='11.11.11.100/24'  
DEFROUTE=no  
IPV4=FAILURE_FATAL=yes  
IPV6INIT=no  
NAME=ib1  
ONBOOT=yes
```

5. Enable the ib interface:

```
# ifup ib0  
# ifup ib1
```

6. Verify the IP addresses you will use to connect to the array. Run this command for both ib0 and ib1:

```
# ip addr show ib0  
# ip addr show ib1
```

As shown in the example below, the IP address for ib0 is 10.10.10.255.

```
10: ib0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 65520 qdisc pfifo_fast state UP group  
default qlen 256  
    link/infiniband 80:00:02:08:fe:80:00:00:00:00:00:00:02:c9:03:00:31:78:51 brd  
00:ff:ff:ff:ff:12:40:1b:ff:ff:00:00:00:00:00:00:00:ff:ff:ff  
    inet 10.10.10.255 brd 10.10.10.255 scope global ib0  
        valid_lft forever preferred_lft forever  
    inet6 fe80::202:c903:31:7851/64 scope link  
        valid_lft forever preferred_lft forever
```

As shown in the example below, the IP address for `ib1` is `11.11.11.255`.

```
10: ib1: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 65520 qdisc pfifo_fast state UP group
default qlen 256
    link/infiniband 80:00:02:08:fe:80:00:00:00:00:00:00:00:02:c9:03:00:31:78:51 brd
00:ff:ff:ff:ff:12:40:1b:ff:ff:00:00:00:00:00:00:00:ff:ff:ff:ff
    inet 11.11.11.255 brd 11.11.11.255 scope global ib0
        valid_lft forever preferred_lft forever
    inet6 fe80::202:c903:31:7851/64 scope link
        valid_lft forever preferred_lft forever
```

7. Set up the NVMe-oF layer on the host.
 - a. Create the following files under `/etc/modules-load.d/` to load the `nvme-rdma` kernel module and make sure the kernel module will always be on, even after a reboot:

```
# cat /etc/modules-load.d/nvme-rdma.conf
nvme-rdma
```

Configure storage array NVMe over InfiniBand connections

If your controller includes an NVMe over InfiniBand port, you can configure the IP address of each port using SANtricity System Manager.

Steps

1. Select **Hardware**.
2. If the graphic shows the drives, click **Show back of shelf**.
The graphic changes to show the controllers instead of the drives.
3. Click the controller with the NVMe over InfiniBand ports you want to configure.
The controller's context menu appears.
4. Select **Configure NVMe over InfiniBand ports**.

Note: The Configure NVMe over InfiniBand ports option appears only if System Manager detects NVMe over InfiniBand ports on the controller.

The **Configure NVMe over InfiniBand Ports** dialog box opens.

5. In the drop-down list, select the HIC port you want to configure, and then enter the IP address of the port.
6. Click **Configure**.
7. Repeat steps 5 and 6 for the other HIC ports that will be used.

Discover and connect to the storage from the host

Before making definitions of each host in SANtricity System Manager, you must discover the target controller ports from the host, and then establish NVMe connections.

Steps

1. Discover available subsystems on the NVMe-oF target for all paths using the following command:

```
nvme discover -t rdma -a target_ip_address
```

In this command, `target_ip_address` is the IP address of the target port.

Note: The `nvme discover` command discovers all controller ports in the subsystem, regardless of host access.

```
# nvme discover -t rdma -a 10.10.10.100
Discovery Log Number of Records 2, Generation counter 0
====Discovery Log Entry 0=====
trtype:  rdma
adrfam:  ipv4
subtype: nvme subsystem
treql:   not specified
portid:  0
trsvcid: 4420
subnqn:  nqn.1992-08.com.netapp:5700.600a098000af41580000000058ed54be
traddr:
10.10.10.100

rdma_prtype: infiniband
rdma_qpctype: connected
rdma_cms:    rdma-cm
rdma_pkey: 0x0000
====Discovery Log Entry 1=====
trtype:  rdma
adrfam:  ipv4
subtype: nvme subsystem
treql:   not specified
portid:  1
trsvcid: 4420
subnqn:  nqn.1992-08.com.netapp:5700.600a098000af41580000000058ed54be
traddr:
11.11.11.100

rdma_prtype: infiniband
rdma_qpctype: connected
rdma_cms:    rdma-cm
rdma_pkey: 0x0000
```

2. Repeat step 1 for any other connections.
3. Connect to the discovered subsystem on the first path using the command: `nvme connect -t rdma -n discovered_sub_nqn -a target_ip_address -Q queue_depth_setting -l controller_loss_timeout_period`

Note: The `nvme connect -t rdma -n discovered_sub_nqn -a target_ip_address -Q queue_depth_setting -l controller_loss_timeout_period` command does not persist through reboot. The NVMe connect command will need to be executed after each reboot to re-establish the NVMe connections.

Note: The nvme connections do not persist through system reboot or extended periods of the controller being unavailable.

Important: Connections are not established for any discovered port inaccessible by the host.

Important: If you specify a port number using this command, the connection fails. The default port is the only port set up for connections.

Important: The recommended queue depth setting is 1024. Override the default setting of 128 with 1024 using the `-Q 1024` command line option, as shown in the following example.

Important: The recommended controller loss timeout period in seconds is 60 minutes (3600 seconds). Override the default setting of 600 seconds with 3600 seconds using the `-l 3600` command line option, as shown in the following example.

```
# nvme connect -t rdma -a 10.10.10.100 -n nqn.1992-08.com.netapp:
5700.600a098000af41580000000058ed54be -Q 1024 -l 3600
```

- 4. Use the `nvme list` command to see a list of the NVMe devices currently connected. In the example below, it is `nvme0n1`.

```
# nvme list
```

Node	SN	Model	Namespace

/dev/nvme0n1	021648023161	NetApp E-Series	1

Usage	Format	FW Rev

5.37 GB / 5.37 GB	512 B + 0 B	0842XXXX

- 5. Connect to the discovered subsystem on the second path:

```
# nvme connect -t rdma -a 11.11.11.100 -n nqn.1992-08.com.netapp:
5700.600a098000af41580000000058ed54be -Q 1024 -l 3600
```

- 6. Use the Linux `lsblk` and `grep` commands to show additional information about each block device:

```
# lsblk | grep nvme
```

nvme0n1	259:0	0	5G	0	disk
nvme1n1	259:0	0	5G	0	disk

- 7. Use the `nvme list` command to see a new list of the NVMe devices currently connected. In the example below, it is `nvme0n1` and `nvme1n1`.

```
# nvme list
```

Node	SN	Model	Namespace

/dev/nvme0n1	021648023161	NetApp E-Series	1
/dev/nvme1n1	021648023161	NetApp E-Series	1

Usage	Format	FW Rev

5.37 GB / 5.37 GB	512 B + 0 B	0842XXXX
5.37 GB / 5.37 GB	512 B + 0 B	0842XXXX

Define a host

Using SANtricity System Manager, you define the hosts that send data to the storage array. Defining a host is one of the steps required to let the storage array know which hosts are attached to it and to allow I/O access to the volumes.

About this task

Keep these guidelines in mind when you define a host:

- You must define the host identifier ports that are associated with the host.
- Make sure that you provide the same name as the host's assigned system name.
- This operation does not succeed if the name you choose is already in use.
- The length of the name cannot exceed 30 characters.

Steps

1. Select **Storage > Hosts**.
2. Click **Create > Host**.
The Create Host dialog box appears.
3. Select the settings for the host as appropriate.

Field details

Setting	Description
Name	Type a name for the new host.
Host operating system type	<p>Select one of the following options from the drop-down list:</p> <ul style="list-style-type: none"> • SANtricity 11.60 and newer <i>Linux</i> • Pre SANtricity 11.60 <i>Linux DM-MP (Kernel 3.10 or later)</i>
Host interface type	Select the host interface type that you want to use.
Host ports	<p>Do one of the following:</p> <ul style="list-style-type: none"> • Select I/O Interface If the host ports have logged in, you can select host port identifiers from the list. This is the recommended method. • Manual add If the host ports have not logged in, look at <code>/etc/nvme/hostnqn</code> on the host to find the hostnqn identifiers and associate them with the host definition. You can manually enter the host port identifiers or copy/paste them from the <code>/etc/nvme/hostnqn</code> file (one at a time) into the Host ports field. You must add one host port identifier at a time to associate it with the host, but you can continue to select as many identifiers that are associated with the host. Each identifier is displayed in the Host ports field. If necessary, you also can remove an identifier by selecting the X next to it.

4. Click **Create**.

Result

After the host is successfully created, SANtricity System Manager creates a default name for each host port configured for the host.

The default alias is <Hostname_Port Number>. For example, the default alias for the first port created for host `IPT` is `IPT_1`.

Assign a volume

You must assign a volume (namespace) to a host or host cluster so it can be used for I/O operations. This assignment grants a host or host cluster access to one or more namespaces in a storage array.

About this task

Keep these guidelines in mind when you assign volumes:

- You can assign a volume to only one host or host cluster at a time.
- Assigned volumes are shared between controllers in the storage array.
- The same namespace ID (NSID) cannot be used twice by a host or a host cluster to access a volume. You must use a unique NSID.

Assigning a volume fails under these conditions:

- All volumes are assigned.
- The volume is already assigned to another host or host cluster.

The ability to assign a volume is unavailable under these conditions:

- No valid hosts or host clusters exist.
- All volume assignments have been defined.

All unassigned volumes are displayed, but functions for hosts with or without Data Assurance (DA) apply as follows:

- For a DA-capable host, you can select volumes that are either DA-enabled or not DA-enabled.
- For a host that is not DA-capable, if you select a volume that is DA-enabled, a warning states that the system must automatically turn off DA on the volume before assigning the volume to the host.

Steps

1. Select **Storage > Hosts**.
2. Select the host or host cluster to which you want to assign volumes, and then click **Assign Volumes**.
A dialog box appears that lists all the volumes that can be assigned. You can sort any of the columns or type something in the **Filter** box to make it easier to find particular volumes.
3. Select the checkbox next to each volume that you want to assign or select the checkbox in the table header to select all volumes.
4. Click **Assign** to complete the operation.

Result

After successfully assigning a volume or volumes to a host or a host cluster, the system performs the following actions:

- The assigned volume receives the next available NSID. The host uses the NSID to access the volume.
- The user-supplied volume name appears in volume listings associated to the host.

Display the volumes visible to the host

Use the `smdevices` tool, part of the `nvme-cli` package, to view the volumes currently visible on the host. This is an alternative to the `nvme list` command.

Step

To view information about each NVMe path to an E-Series volume, use the `nvme netapp smdevices [-o <format>]` command. The output `<format>` can be normal (the default if `-o` is not used), column, or json.

```
# nvme netapp smdevices
/dev/nvme1n1, Array Name ICTM0706SYS04, Volume Name NVMe2, NSID 1, Volume ID
000015bd5903df4a00a0980000af4462, Controller A, Access State unknown, 2.15GB
/dev/nvme1n2, Array Name ICTM0706SYS04, Volume Name NVMe3, NSID 2, Volume ID
000015c05903e24000a0980000af4462, Controller A, Access State unknown, 2.15GB
/dev/nvme1n3, Array Name ICTM0706SYS04, Volume Name NVMe4, NSID 4, Volume ID
00001bb0593a46f400a0980000af4462, Controller A, Access State unknown, 2.15GB
/dev/nvme1n4, Array Name ICTM0706SYS04, Volume Name NVMe6, NSID 6, Volume ID
00001696593b424b00a0980000af4112, Controller A, Access State unknown, 2.15GB
/dev/nvme2n1, Array Name ICTM0706SYS04, Volume Name NVMe2, NSID 1, Volume ID
000015bd5903df4a00a0980000af4462, Controller B, Access State unknown, 2.15GB
/dev/nvme2n2, Array Name ICTM0706SYS04, Volume Name NVMe3, NSID 2, Volume ID
000015c05903e24000a0980000af4462, Controller B, Access State unknown, 2.15GB
/dev/nvme2n3, Array Name ICTM0706SYS04, Volume Name NVMe4, NSID 4, Volume ID
00001bb0593a46f400a0980000af4462, Controller B, Access State unknown, 2.15GB
/dev/nvme2n4, Array Name ICTM0706SYS04, Volume Name NVMe6, NSID 6, Volume ID
00001696593b424b00a0980000af4112, Controller B, Access State unknown, 2.15GB
```

Set up failover

Multipath software provides a redundant path to the storage array in case one of the physical paths is disrupted. There are currently two methods of multipathing available for NVMe, and which you will be using is going to be dependent on which OS version you are running. For RHEL 7 and SLES 12, device mapper multipath (DMMP) will be used. For SLES 15, a native NVMe multipathing solution will be used.

Configuring the host to run failover

The SUSE Linux Enterprise Server host requires configuration changes to run failover. The failover solution uses DM-MP.

Before you begin

- You have installed the required packages on your system.
- For Red Hat (RHEL) hosts, verify the packages are installed by running `rpm -q device-mapper-multipath`
- For SLES hosts, verify the packages are installed by running `rpm -q multipath-tools`

Note: Refer to the NetApp Interoperability Matrix Tool (IMT) to ensure any required updates are installed as multipathing may not work correctly with the GA versions of SLES or RHEL.

About this task

By default, DM-MP is disabled in RHEL and SLES. Complete the following steps to enable DM-MP components on the host.

Steps

1. Add the NVMe E-Series device entry to the devices section of the `/etc/multipath.conf` file, as shown in the following example:

```
devices {
    device {
        vendor "NVME"
        product "NetApp E-Series*"
        path_grouping_policy group_by_prio
        failback immediate
        no_path_retry 30
    }
}
```

2. Configure `multipathd` to start at system boot.

```
# systemctl enable multipathd
```

3. Start `multipathd` if it is not currently running.

```
# systemctl start multipathd
```

4. Verify the status of `multipathd` to make sure it is active and running:

```
# systemctl status multipathd
```

Accessing NVMe volumes

You can configure the I/O directed to the device target based on your version of SLES (SUSE Linux version).

Accessing NVMe volumes for virtual device targets

For RHEL 7 and SLES 12, I/O is directed to virtual device targets by the Linux host. DM-MP manages the physical paths underlying these virtual targets.

Virtual devices are I/O targets

Make sure you are running I/O only to the virtual devices created by DM-MP and not to the physical device paths. If you are running I/O to the physical paths, DM-MP cannot manage a failover event and the I/O fails.

You can access these block devices through the `dm` device or the symlink in `/dev/mapper`, for example:

```
/dev/dm-1
/dev/mapper/eui.00001bc7593b7f5f00a0980000af4462
```

Example

The following example output from the `nvme list` command shows the host node name and its correlation with the namespace ID.

NODE	SN	MODEL	NAMESPACE
/dev/nvme1n1	021648023072	NetApp E-Series	10
/dev/nvme1n2	021648023072	NetApp E-Series	11
/dev/nvme1n3	021648023072	NetApp E-Series	12

```
/dev/nvme1n4 021648023072 NetApp E-Series 13
/dev/nvme2n1 021648023151 NetApp E-Series 10
/dev/nvme2n2 021648023151 NetApp E-Series 11
/dev/nvme2n3 021648023151 NetApp E-Series 12
/dev/nvme2n4 021648023151 NetApp E-Series 13
```

Column	Description
Node	The node name includes two parts: <ul style="list-style-type: none"> The notation <code>nvme1</code> represents controller A and <code>nvme2</code> represents controller B. The notation <code>n1</code>, <code>n2</code>, and so on represent the namespace identifier from the host perspective. These identifiers are repeated in the table, once for controller A and once for controller B.
Namespace	The Namespace column lists the namespace ID (NSID), which is the identifier from the storage array perspective.

In the following `multipath -ll` output, the optimized paths are shown with a `prio` value of 50, while the non-optimized paths are shown with a `prio` value of 10.

The Linux operating system routes I/O to the path group that is shown as `status=active`, while the path groups listed as `status=enabled` are available for failover.

```
eui.00001bc7593b7f500a0980000af4462 dm-0 NVME,NetApp E-Series
size=15G features='1 queue_if_no_path' hwhandler='0' wp=rw
|+- policy='service-time 0' prio=50 status=active
|  `- #:#:#:# nvme1n1 259:5 active ready running
`+- policy='service-time 0' prio=10 status=enabled
   `- #:#:#:# nvme2n1 259:9 active ready running
```

```
eui.00001bc7593b7f5f00a0980000af4462 dm-0 NVME,NetApp E-Series
size=15G features='1 queue_if_no_path' hwhandler='0' wp=rw
|+- policy='service-time 0' prio=0 status=enabled
|  `- #:#:#:# nvme1n1 259:5 failed faulty running
`+- policy='service-time 0' prio=10 status=active
   `- #:#:#:# nvme2n1 259:9 active ready running
```

Line item	Description
<code>policy='service-time 0' prio=50 status=active</code>	This line and the following line show that <code>nvme1n1</code> , which is the namespace with an NSID of 10, is optimized on the path with a <code>prio</code> value of 50 and a <code>status</code> value of <code>active</code> . This namespace is owned by controller A.
<code>policy='service-time 0' prio=10 status=enabled</code>	This line shows the failover path for namespace 10, with a <code>prio</code> value of 10 and a <code>status</code> value of <code>enabled</code> . I/O is not being directed to the namespace on this path at the moment. This namespace is owned by controller B.
<code>policy='service-time 0' prio=0 status=enabled</code>	This example shows <code>multipath -ll</code> output from a different point in time, while controller A is rebooting. The path to namespace 10 is shown as <code>failed faulty running</code> with a <code>prio</code> value of 0 and a <code>status</code> value of <code>enabled</code> .
<code>policy='service-time 0' prio=10 status=active</code>	Note that the active path refers to <code>nvme2</code> , so the I/O is being directed on this path to controller B.

Accessing NVMe volumes for physical NVMe device targets

For SLES 15, I/O is directed to the physical NVMe device targets by the Linux host. A native NVMe multipathing solution manages the physical paths underlying the single apparent physical device displayed by the host.

Note: It is best practice to use the links in `/dev/disk/by-id/` rather than `/dev/nvme0n1`, for example:

```
# ls /dev/disk/by-id/ -l lrwxrwxrwx 1 root root 13 Oct 18 15:14
nvme-
eui.0000320f5cad32cf00a0980000af4112 -> ../../nvme0n1
```

Physical NVMe devices are I/O targets

Run I/O to the physical nvme device path. There should only be one of these devices present for each namespace using the following format:

```
/dev/nvme[subsyst#]n[id#]
```

All paths are virtualized using the native multipathing solution underneath this device.

You can view your paths by running:

```
# nvme list-sys
```

Example output:

```
nvme-sys0 - NQN=nqn.1992-08.com.netapp:5700.600a098000a522500000000589aa8a6
\
+- nvme0 rdma traddr=192.4.21.131 trsvcid=4420 live
+- nvme1 rdma traddr=192.4.22.141 trsvcid=4420 live
```

If you specify a namespace device when using the 'nvme list-sys' command, it provides additional information about the paths to that namespace:

```
# nvme list-sys /dev/nvme0n1
nvme-sys0 - NQN=nqn.1992-08.com.netapp:5700.600a098000af44620000000058d5dd96
\
+- nvme0 rdma traddr=192.168.130.101 trsvcid=4420 live non-optimized
+- nvme1 rdma traddr=192.168.131.101 trsvcid=4420 live non-optimized
+- nvme2 rdma traddr=192.168.130.102 trsvcid=4420 live optimized
+- nvme3 rdma traddr=192.168.131.102 trsvcid=4420 live optimized
```

There are also hooks into the multipath commands to allow you to view your path information for native failover through them as well:

```
#multipath -ll
```

Note: To view the path information, the following must be set in `/etc/multipath.conf`:

```
defaults {
    enable_foreign nvme
}
```

Example output:

```
eui.0000a0335c05d57a00a0980000a5229d [nvme]:nvme0n9 NVMe,Netapp E-Series,08520001
size=4194304 features='n/a' hwhandler='ANA' wp=rw
|+- policy='n/a' prio=50 status=optimized
|  `-- 0:0:1 nvme0c0n1 0:0 n/a optimized    live
+- policy='n/a' prio=10 status=non-optimized
`- 0:1:1 nvme0c1n1 0:0 n/a non-optimized    live
```

Create filesystems

You create a file system on the namespace or native NVMe device and mount the filesystem.

Create filesystems (RHEL 7 and SLES 12)

For RHEL 7 and SLES 12, you create a file system on the namespace and mount the filesystem.

Steps

1. Run the `multipath -ll` command to get a list of `/dev/mapper/dm` devices.

```
# multipath -ll
```

The result of this command shows two devices, `dm-19` and `dm-16` :

```
eui.00001ffe5a94ff8500a0980000af4444 dm-19 NVME,NetApp E-Series
size=10G features='1 queue_if_no_path' hwhandler='0' wp=rw
|+- policy='service-time 0' prio=50 status=active
|  |- #:#:#:# nvme0n19 259:19  active ready running
|  `-- #:#:#:# nvme1n19 259:115 active ready running
+- policy='service-time 0' prio=10 status=enabled
|  |- #:#:#:# nvme2n19 259:51  active ready running
|  `-- #:#:#:# nvme3n19 259:83  active ready running
eui.00001fd25a94fef000a0980000af4444 dm-16 NVME,NetApp E-Series
size=16G features='1 queue_if_no_path' hwhandler='0' wp=rw
|+- policy='service-time 0' prio=50 status=active
|  |- #:#:#:# nvme0n16 259:16  active ready running
|  `-- #:#:#:# nvme1n16 259:112 active ready running
+- policy='service-time 0' prio=10 status=enabled
|  |- #:#:#:# nvme2n16 259:48  active ready running
|  `-- #:#:#:# nvme3n16 259:80  active ready running
```

2. Create a file system on the partition for each `/dev/mapper/eui-` device.

The method for creating a file system varies depending on the file system chosen. This example shows creating an `ext4` file system.

```
# mkfs.ext4 /dev/mapper/dm-19
mke2fs 1.42.11 (09-Jul-2014)
Creating filesystem with 2620928 4k blocks and 655360 inodes
Filesystem UUID: 97f987e9-47b8-47f7-b434-bf3ebbe826d0
Superblock backups stored on blocks:
    32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632

Allocating group tables: done
Writing inode tables: done
Creating journal (32768 blocks): done
Writing superblocks and filesystem accounting information: done
```

3. Create a folder to mount the new device.

```
# mkdir /mnt/ext4
```

4. Mount the device.

```
# mount /dev/mapper/eui.00001ffe5a94ff8500a0980000af4444 /mnt/ext4
```

Create filesystems (SLES 15)

For SLES 15, you create a filesystem on the native nvme device and mount the filesystem.

Steps

1. Run the `multipath -ll` command to get a list of `/dev/nvme` devices.

```
# multipath -ll
```

The result of this command shows device `nvme0n6`.

```
eui.000082dd5c05d39300a0980000a52225 [nvme]:nvme0n6 NVMe,NetApp E-Series,08520000
size=4194304 features='n/a' hwhandler='ANA' wp=rw
|+- policy='n/a' prio=50 status=optimized
|  `-- 0:0:1 nvme0c0n1 0:0 n/a optimized live
|+- policy='n/a' prio=50 status=optimized
|  `-- 0:1:1 nvme0c1n1 0:0 n/a optimized live
|+- policy='n/a' prio=10 status=non-optimized
|  `-- 0:2:1 nvme0c2n1 0:0 n/a non-optimized live
`+- policy='n/a' prio=10 status=non-optimized
   `-- 0:3:1 nvme0c3n1 0:0 n/a non-optimized live
```

2. Create a file system on the partition for each `/dev/nvme0n#` device.

The method for creating a file system varies depending on the file system chosen. This example shows creating an `ext4` file system.

```
# mkfs.ext4 /dev/disk/by-id/nvme-eui.000082dd5c05d39300a0980000a52225
mke2fs 1.42.11 (22-Oct-2019)
Creating filesystem with 2620928 4k blocks and 655360 inodes
Filesystem UUID: 97f987e9-47b8-47f7-b434-bf3ebbe826d0
Superblock backups stored on blocks:
    32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632

Allocating group tables: done
Writing inode tables: done
Creating journal (32768 blocks): done
Writing superblocks and filesystem accounting information: done
```

3. Create a folder to mount the new device.

```
# mkdir /mnt/ext4
```

4. Mount the device.

```
# mount /dev/disk/by-id/nvme-eui.000082dd5c05d39300a0980000a52225 /mnt/ext4
```

Verify storage access on the host

Before using the namespace, you verify that the host can write data to the namespace and read it back.

Before you begin

You must have initialized the namespace and formatted it with a file system.

Steps

- 1. On the host, copy one or more files to the mount point of the disk.
- 2. Copy the files back to a different folder on the original disk.
- 3. Run the `diff` command to compare the copied files to the originals.

After you finish

Remove the file and folder that you copied.

Record NVMe over InfiniBand-specific information for Linux

Select the NVMe over InfiniBand worksheet to record your protocol-specific storage configuration information. You need this information to perform provisioning tasks.

NVMe over InfiniBand worksheet

You can use this worksheet to record NVMe over Infiniband storage configuration information. You need this information to perform provisioning tasks.

NVMe over InfiniBand: Host identifiers


Note: The software initiator NQN is determined during the task, .

Locate and document the initiator NQN from each host. The NQN is typically found in the `/etc/nvme/hostnqn` file.

Callout No.	Host port connections	Host NQN
1	Host (initiator) 1	
n/a		
n/a		
n/a		
n/a		

NVMe over InfiniBand: Recommended configuration

In a direct connect topology, one or more hosts are directly connected to the subsystem. In the SANtricity OS 11.50 release, we support a single connection from each host to a subsystem controller, as shown below. In this configuration, one HCA (host channel adapter) port from each host should be on the same subnet as the E-Series controller port it is connected to, but on a different subnet from the other HCA port.



NVMe over InfiniBand: Target NQN

Document the target NQN for the storage array. You will use this information in [Configure storage array NVMe over InfiniBand connections](#) on page 69 .

Find the Storage Array NQN name using SANtricity: **Storage Array > NVMe over Infiniband > Manage Settings**. This information might be necessary when you create NVMe over Infiniband sessions from operating systems that do not support send targets discovery.

Callout No.	Array name	Target IQN
6	Array controller (target)	

NVMe over InfiniBand: Network configuration

Document the network configuration that will be used for the hosts and storage on the InfiniBand fabric. These instructions assume that two subnets will be used for full redundancy.

Your network administrator can provide the following information. You use this information in the topic, [Configure storage array NVMe over InfiniBand connections](#) on page 69.

Subnet A

Define the subnet to be used.

Network Address	Netmask

Document the NQNs to be used by the array ports and each host port.

Callout No.	Array controller (target) port connections	NQN
3	Switch	<i>not applicable</i>
5	Controller A, port 1	

Callout No.	Array controller (target) port connections	NQN
4	Controller B, port 1	
2	Host 1, port 1	
	(Optional) Host 2, port 1	

Subnet B

Define the subnet to be used.

Network Address	Netmask

Document the IQNs to be used by the array ports and each host port.

Callout No.	Array controller (target) port connections	NQN
8	Switch	<i>not applicable</i>
10	Controller A, port 2	
9	Controller B, port 2	
7	Host 1, port 2	
	(Optional) Host 2, port 2	

NVMe over InfiniBand: Mapping host name

Note: The mapping host name is created during the workflow.

Mapping host name	
Host OS type	

NVMe over RoCE Express Setup

You can use NVMe with the RDMA over Converged Ethernet (RoCE) network protocol.

Steps

1. [Verify the Linux configuration is supported](#) on page 83
2. [Configuring IP addresses using DHCP](#) on page 84
3. [Install SANtricity Storage Manager for SMcli \(SANtricity software version 11.53 or earlier\)](#) on page 85
4. [Access SANtricity System Manager and use Setup wizard](#) on page 86
5. [Configure the switch](#) on page 87
6. [Set up NVMe over RoCE on the host side](#) on page 87
7. [Configure storage array NVMe over RoCE connections](#) on page 89
8. [Discover and connect to the storage from the host](#) on page 91
9. [Define a host](#) on page 92
10. [Assign a volume](#) on page 93
11. [Display the volumes visible to the host](#) on page 94
12. [Set up failover on the host](#) on page 94
13. [Accessing NVMe Volumes](#) on page 95
14. [Create filesystems](#) on page 98
15. [Verify storage access on the host](#) on page 100
16. [Record NVMe over RoCE-specific information for Linux](#) on page 100

Verify the Linux configuration is supported

To ensure reliable operation, you create an implementation plan and then use the NetApp Interoperability Matrix Tool (IMT) to verify that the entire configuration is supported.

Steps

1. Go to the [NetApp Interoperability Matrix Tool](#).
2. Click on the **Solution Search** tile.
3. In the **Protocols > SAN Host** area, click the **Add** button next to **E-Series SAN Host**.
4. Click **View Refine Search Criteria**.
The **Refine Search Criteria** section is displayed. In this section you may select the protocol that applies, as well as other criteria for the configuration such as Operating System, NetApp OS, and Host Multipath driver. Select the criteria you know you want for your configuration, and then see what compatible configuration elements apply. As necessary, make the updates for your operating system and protocol that are prescribed in the tool. Detailed information for your chosen configuration is accessible on the View Supported Configurations page by clicking the **right page arrow**.

NVMe over RoCE restrictions

Before using NVMe over RoCE, review the controller, host, and recovery restrictions.

Verify your configuration

Verify your configuration, using the [NetApp Interoperability Matrix Tool](#).

About the hardware

NVME over RoCE can be configured for the EF300 (100GB controllers only), EF600, EF570, and E5700 controllers. The controllers must have 100GB or 200GB host port.

Restrictions

The following restrictions are in effect for the 11.60 release. See the [NetApp Interoperability Matrix Tool](#) for a complete list of requirements.

Controller restrictions

- This protocol can be used only for EF300, EF600, EF570, and E5700 controllers. A minimum of 32 GB of physical memory is required to use the protocol on EF600, EF570, and E5700 controllers. For the EF300, a minimum of 16 GB of physical memory is required. If the minimum memory requirements for the controllers are not met during start of day operations, a message is displayed that helps you diagnose the problem.
- No simplex (single-controller) configurations are supported.
- The only supported host interface card (HIC) is the 100G or 200G EDR HIC, which also supports NVMe over InfiniBand, iSER and SRP.
- There is no support for mixed NVMe over RoCE with NVMe over InfiniBand or SCSI host interfaces.

Switch restrictions



Attention: RISK OF DATA LOSS. You must enable Priority Flow Control or Global Pause Control on the switch to eliminate the risk of data loss in an NVMe over RoCE environment.

Host, host protocol, and host operating system restrictions

- See the [NetApp Interoperability Matrix Tool](#) for a complete list of requirements.
- For a list of supported host channel adapters see the [NetApp Interoperability Matrix Tool](#).

Storage and disaster recovery restrictions

- Asynchronous and synchronous mirroring are not supported.
- Thin provisioning (the creation of thin volumes) is not supported.

Configuring IP addresses using DHCP

In this express method for configuring communications between the management station and the storage array, you use Dynamic Host Configuration Protocol (DHCP) to provide IP addresses. Each storage array has either one controller (simplex) or two controllers (duplex), and each controller has two storage management ports. Each management port will be assigned an IP address.

Before you begin

You have installed and configured a DHCP server on the same subnet as the storage management ports.

About this task

The following instructions refer to a storage array with two controllers (a duplex configuration).

1. If you have not already done so, connect an Ethernet cable to the management station and to management port 1 on each controller (A and B).

The DHCP server assigns an IP address to port 1 of each controller.

Note: Do not use management port 2 on either controller. Port 2 is reserved for use by NetApp technical personnel.

Important: If you disconnect and reconnect the Ethernet cable, or if the storage array is power-cycled, DHCP assigns IP addresses again. This process occurs until static IP addresses are configured. It is recommended that you avoid disconnecting the cable or power-cycling the array.

If the storage array cannot get DHCP-assigned IP addresses within 30 seconds, the following default IP addresses are set:

- Controller A, port 1: 169.254.128.101
 - Controller B, port 1: 169.254.128.102
 - Subnet mask: 255.255.0.0
2. Locate the MAC address label on the back of each controller, and then provide your network administrator with the MAC address for port 1 of each controller.
- Your network administrator needs the MAC addresses to determine the IP address for each controller. You will need the IP addresses to connect to your storage system through your browser.

Install SANtricity Storage Manager for SMcli (SANtricity software version 11.53 or earlier)

When you install the SANtricity Storage Manager software on your management station, a graphical user interface (GUI) and a command line interface (CLI) are installed by default. These instructions assume that you will install the SANtricity Storage Manager GUI on a management station and *not* on the I/O host.

Before you begin

Important: For SANtricity software 11.60 and newer, the SANtricity Secure CLI (SMcli) is included in the SANtricity OS and downloadable through the SANtricity System Manager. For more information on how to download the SMcli through the SANtricity System Manager, refer to the *Download the command line interface (CLI)* topic under the SANtricity System Manager Online Help.

- You are using SANtricity software 11.53 or earlier.
- You must have the correct administrator or superuser privileges.
- You must have ensured that the system that will contain the SANtricity Storage Manager client has the following minimum requirements:
 - **RAM:** 2 GB for Java Runtime Engine
 - **Disk space:** 5 GB
 - **OS/Architecture:** Refer to [NetApp Interoperability Matrix Tool](#) for guidance on determining the supported operating system versions and architectures.

About this task

This section describes how to install SANtricity Storage Manager on both the Windows and Linux OS platforms, because both Windows and Linux are common management station platforms when Linux is used for the data host.

Steps

1. Download the SANtricity software release from [NetApp Support Downloads](#) > **Software** > **E-Series/EF-Series SANtricity Storage Manager**.
2. Run the SANtricity installer.

Windows	Linux
Double-click the <code>SMIA*.exe</code> installation package to start the installation.	<ol style="list-style-type: none"> Go to the directory where the <code>SMIA*.bin</code> installation package is located. If the temp mount point does not have execute permissions, set the <code>IATEMPDIR</code> variable. Example: <code>IATEMPDIR=/root ./SMIA-LINUX64-11.25.0A00.0002.bin</code> Run the <code>chmod +x SMIA*.bin</code> command to grant execute permission to the file. Run the <code>./SMIA*.bin</code> command to start the installer.

- Use the installation wizard to install the software on the management station.

Access SANtricity System Manager and use Setup wizard

You use the Setup wizard in SANtricity System Manager to configure your storage array.

Before you begin

- You have ensured that the device from which you will access SANtricity System Manager contains one of the following browsers:

Browser	Minimum version
Google Chrome	47
Microsoft Internet Explorer	11
Microsoft Edge	EdgeHTML 12
Mozilla Firefox	31
Safari	9

- You are using out-of-band management.

About this task

The wizard automatically relaunches when you open System Manager or refresh your browser and *at least one* of the following conditions is met:

- No pools and volume groups are detected.
- No workloads are detected.
- No notifications are configured.

Steps

- From your browser, enter the following URL:

`https://<DomainNameOrIPAddress>`

`IPAddress` is the address for one of the storage array controllers.

The first time SANtricity System Manager is opened on an array that has not been configured, the `Set Administrator Password` prompt appears. Role-based access management configures four local roles: admin, support, security, and monitor. The latter three roles have random passwords that cannot be guessed. After you set a password for the admin role you can change all of the passwords using the admin credentials. See *SANtricity System Manager online help* for more information on the four local user roles.

- Enter the System Manager password for the admin role in the `Set Administrator Password` and `Confirm Password` fields, and then select the **Set Password** button.

When you open System Manager and no pools, volumes groups, workloads, or notifications have been configured, the Setup wizard launches.

3. Use the Setup wizard to perform the following tasks:

- **Verify hardware (controllers and drives)** – Verify the number of controllers and drives in the storage array. Assign a name to the array.
- **Verify hosts and operating systems** – Verify the host and operating system types that the storage array can access.
- **Accept pools** – Accept the recommended pool configuration for the express installation method. A pool is a logical group of drives.
- **Configure alerts** – Allow System Manager to receive automatic notifications when a problem occurs with the storage array.
- **Enable AutoSupport** – Automatically monitor the health of your storage array and have dispatches sent to technical support.

4. If you have not already created a volume, create one by going to **Storage > Volumes > Create > Volume**.

For more information, see the online help for SANtricity System Manager.

Configure the switch

You configure the switches according to the vendor's recommendations for NVMe over RoCE. These recommendations might include both configuration directives as well as code updates.

About this task



Attention: RISK OF DATA LOSS. You must enable Priority Flow Control or Global Pause Control on the switch to eliminate the risk of data loss in an NVMe over RoCE environment.

Enable Ethernet pause frame flow control **end to end** as the best practice configuration.

Consult your network administrator for tips on selecting the best configuration for your environment.

Set up NVMe over RoCE on the host side

NVMe initiator configuration in a RoCE environment includes installing and configuring the `rdma-core` and `nvme-cli` packages, configuring initiator IP addresses, and setting up the NVMe-oF layer on the host.

Before you begin

- You are running RHEL 7 and the latest compatible SUSE Linux Enterprise Server 12 and 15 service pack operating system. See the [NetApp Interoperability Matrix Tool](#) for a complete list of the latest requirements.

Steps

1. Install the `rdma` and `nvme-cli` packages:

```
# zypper install rdma-core
# zypper install nvme-cli
```

RHEL 7

```
# yum install rdma-core
# yum install nvme-cli
```

2. Setup IPv4 IP addresses on the ethernet ports used to connect NVMe over RoCE. For each network interface, create a configuration script that contains the different variables for that interface.

The variables used in this step are based on server hardware and the network environment. The variables include the *IPADDR* and *GATEWAY*. These are example instructions for the latest SUSE Linux Enterprise Server 12 service pack:

Create the example file `/etc/sysconfig/network/ifcfg-eth4` as follows:

```
BOOTPROTO='static'
BROADCAST=
ETHTOOL_OPTIONS=
IPADDR='192.168.1.87/24'
GATEWAY='192.168.1.1'
MTU=
NAME='MT27800 Family [ConnectX-5]'
NETWORK=
REMOTE_IPADDR=
STARTMODE='auto'
```

Create the second example file `/etc/sysconfig/network/ifcfg-eth5` as follows:

```
BOOTPROTO='static'
BROADCAST=
ETHTOOL_OPTIONS=
IPADDR='192.168.2.87/24'
GATEWAY='192.168.2.1'
MTU=
NAME='MT27800 Family [ConnectX-5]'
NETWORK=
REMOTE_IPADDR=
STARTMODE='auto'
```

3. Enable the network interfaces:

```
# ifup eth4
# ifup eth5
```

4. Set up the NVMe-oF layer on the host.

- a. Create the following file under `/etc/modules-load.d/` to load the `nvme-rdma` kernel module and make sure the kernel module will always be on, even after a reboot:

```
# cat /etc/modules-load.d/nvme-rdma.conf
nvme-rdma
```


Configure storage array NVMe over RoCE connections

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

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.

About this task

You can access the NVMe over RoCE configuration from the **Hardware** page or from **Settings > System**. This task describes how to configure the ports from the Hardware page.

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

Steps

1. Select **Hardware**.
2. Click the controller with the NVMe over RoCE port you want to configure.
The controller's context menu appears.
3. Select **Configure NVMe over RoCE ports**.
The **Configure NVMe over RoCE ports** dialog box opens.
4. In the drop-down list, select the port you want to configure, and then click **Next**.
5. Select the port configuration settings you want to use, and then click **Next**.

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

Port Setting	Description
Configured ethernet port speed	<p>Select the desired speed.</p> <p>The options that appear in the drop-down list depend on the maximum speed that your network can support (for example, 10 Gbps). Possible values include:</p> <ul style="list-style-type: none">• Auto-negotiate• 10 Gbps• 25 Gbps• 40 Gbps• 50 Gbps• 100 Gbps• 200 Gbps <p>Note: When a 200Gb-capable HIC is attached with a QSFP56 cable, auto-negotiate is only available when you are connecting to Mellanox switches and/or adapters.</p> <p>Note: The configured NVMe over RoCE port speed should match the speed capability of the SFP on the selected port. All ports must be set to the same speed.</p>
Enable IPv4 and/or Enable IPv6	Select one or both options to enable support for IPv4 and IPv6 networks.

Port Setting	Description
MTU size (Available by clicking Show more port settings .)	If necessary, enter a new size in bytes for the maximum transmission unit (MTU). The default MTU size is 1500 bytes per frame. You must enter a value between 1500 and 4200.

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

Port setting	Description
Automatically obtain configuration from DHCP server	Select this option to obtain the configuration automatically.
Manually specify static configuration	Select this option, and then enter a static address in the fields. For IPv4, include the network subnet mask and gateway. For IPv6, include the routable IP addresses and router IP address. Note: If there is only one routable IP address, set the remaining address to 0:0:0:0:0:0:0:0.
Enable VLAN support (Available by clicking Show more settings .)	Important: This option is only available in an iSCSI environment. It is not available in an NVMe over RoCE environment. Select this option to enable a VLAN and enter its ID. A VLAN is a logical network that behaves like it is physically separate from other physical and virtual local area networks (LANs) supported by the same switches, the same routers, or both.
Enable ethernet priority (Available by clicking Show more settings .)	Important: This option is only available in an iSCSI environment. It is not available in an NVMe over RoCE environment. Select this option to enable the parameter that determines the priority of accessing the network. Use the slider to select a priority between 1 and 7. In a shared local area network (LAN) environment, such as Ethernet, many stations might contend for access to the network. Access is on a first-come, first-served basis. Two stations might try to access the network at the same time, which causes both stations to back off and wait before trying again. This process is minimized for switched Ethernet, where only one station is connected to a switch port.

7. Click **Finish**.

Discover and connect to the storage from the host

Before making definitions of each host in SANtricity System Manager, you must discover the target controller ports from the host, and then establish NVMe connections.

Steps

1. Discover available subsystems on the NVMe-oF target for all paths using the following command:

```
nvme discover -t rdma -a target_ip_address
```

In this command, *target_ip_address* is the IP address of the target port.

Note: The `nvme discover` command discovers all controller ports in the subsystem, regardless of host access.

```
# nvme discover -t rdma -a 192.168.1.77
Discovery Log Number of Records 2, Generation counter 0
====Discovery Log Entry 0=====
trtype:  rdma
adrfam:  ipv4
subtype: nvme subsystem
treq:    not specified
portid:  0
trsvcid: 4420
subnqn:  nqn.1992-08.com.netapp:5700.600a098000a527a7000000005ab3af94
traddr:  192.168.1.77
rdma_prtype: roce
rdma_qptype: connected
rdma_cms:   rdma-cm
rdma_pkey:  0x0000
====Discovery Log Entry 1=====
trtype:  rdma
adrfam:  ipv4
subtype: nvme subsystem
treq:    not specified
portid:  1
trsvcid: 4420
subnqn:  nqn.1992-08.com.netapp:5700.600a098000a527a7000000005ab3af94
traddr:  192.168.2.77
rdma_prtype: roce
rdma_qptype: connected
rdma_cms:   rdma-cm
rdma_pkey:  0x0000
```

2. Repeat step 1 for any other connections.
3. Connect to the discovered subsystem on the first path using the command: `nvme connect -t rdma -n discovered_sub_nqn -a target_ip_address -Q queue_depth_setting -l controller_loss_timeout_period`

Note: The `nvme connect -t rdma -n discovered_sub_nqn -a target_ip_address -Q queue_depth_setting -l controller_loss_timeout_period` command does not persist through reboot. The NVMe connect command will need to be executed after each reboot to re-establish the NVMe connections.

Important: Connections are not established for any discovered port inaccessible by the host.

Important: If you specify a port number using this command, the connection fails. The default port is the only port set up for connections.

Important: The recommended queue depth setting is 1024. Override the default setting of 128 with 1024 using the `-Q 1024` command line option, as shown in the following example.

Important: The recommended controller loss timeout period in seconds is 60 minutes (3600 seconds). Override the default setting of 600 seconds with 3600 seconds using the `-l 3600` command line option, as shown in the following example.

```
# nvme connect -t rdma -a 192.168.1.77 -n nqn.1992-08.com.netapp:
5700.600a098000a527a7000000005ab3af94 -Q 1024 -l 3600
# nvme connect -t rdma -a 192.168.2.77 -n nqn.1992-08.com.netapp:
5700.600a098000a527a7000000005ab3af94 -Q 1024 -l 3600
```

4. Repeat step 3 to connect the discovered subsystem on the second path.

Define a host

Using SANtricity System Manager, you define the hosts that send data to the storage array. Defining a host is one of the steps required to let the storage array know which hosts are attached to it and to allow I/O access to the volumes.

About this task

Keep these guidelines in mind when you define a host:

- You must define the host identifier ports that are associated with the host.
- Make sure that you provide the same name as the host's assigned system name.
- This operation does not succeed if the name you choose is already in use.
- The length of the name cannot exceed 30 characters.

Steps

1. Select **Storage > Hosts**.
2. Click **Create > Host**.
The Create Host dialog box appears.
3. Select the settings for the host as appropriate.

Field details

Setting	Description
Name	Type a name for the new host.
Host operating system type	Select one of the following options from the drop-down list: <ul style="list-style-type: none"> • SANtricity 11.60 and newer <i>Linux</i> • Pre SANtricity 11.60 <i>Linux DM-MP (Kernel 3.10 or later)</i>
Host interface type	Select the host interface type that you want to use. If the array you configure only has one available host interface type, this setting may not be available to select.

Setting	Description
Host ports	<p>Do one of the following:</p> <ul style="list-style-type: none"> • Select I/O Interface If the host ports have logged in, you can select host port identifiers from the list. This is the recommended method. • Manual add If the host ports have not logged in, look at <code>/etc/nvme/hostnqn</code> on the host to find the hostnqn identifiers and associate them with the host definition. You can manually enter the host port identifiers or copy/paste them from the <code>/etc/nvme/hostnqn</code> file (one at a time) into the Host ports field. You must add one host port identifier at a time to associate it with the host, but you can continue to select as many identifiers that are associated with the host. Each identifier is displayed in the Host ports field. If necessary, you also can remove an identifier by selecting the X next to it.

4. Click **Create**.

Result

After the host is successfully created, SANtricity System Manager creates a default name for each host port configured for the host.

The default alias is `<Hostname_Port Number>`. For example, the default alias for the first port created for host `IPT` is `IPT_1`.

Assign a volume

You must assign a volume (namespace) to a host or host cluster so it can be used for I/O operations. This assignment grants a host or host cluster access to one or more namespaces in a storage array.

About this task

Keep these guidelines in mind when you assign volumes:

- You can assign a volume to only one host or host cluster at a time.
- Assigned volumes are shared between controllers in the storage array.
- The same namespace ID (NSID) cannot be used twice by a host or a host cluster to access a volume. You must use a unique NSID.

Assigning a volume fails under these conditions:

- All volumes are assigned.
- The volume is already assigned to another host or host cluster.

The ability to assign a volume is unavailable under these conditions:

- No valid hosts or host clusters exist.
- All volume assignments have been defined.

All unassigned volumes are displayed, but functions for hosts with or without Data Assurance (DA) apply as follows:

- For a DA-capable host, you can select volumes that are either DA-enabled or not DA-enabled.

- For a host that is not DA-capable, if you select a volume that is DA-enabled, a warning states that the system must automatically turn off DA on the volume before assigning the volume to the host.

Steps

1. Select **Storage > Hosts**.
2. Select the host or host cluster to which you want to assign volumes, and then click **Assign Volumes**.
A dialog box appears that lists all the volumes that can be assigned. You can sort any of the columns or type something in the **Filter** box to make it easier to find particular volumes.
3. Select the checkbox next to each volume that you want to assign or select the checkbox in the table header to select all volumes.
4. Click **Assign** to complete the operation.

Result

After successfully assigning a volume or volumes to a host or a host cluster, the system performs the following actions:

- The assigned volume receives the next available NSID. The host uses the NSID to access the volume.
- The user-supplied volume name appears in volume listings associated to the host.

Display the volumes visible to the host

The `SMdevices` tool, part of the `nvme-cli` package, allows you to view the volumes currently visible on the host. This tool is an alternative to the `nvme list` command.

Step

To view information about each NVMe path to an E-Series volume, use the `nvme netapp smdevices [-o <format>]` command. The output `<format>` can be normal (the default if `-o` is not used), column, or json.

```
# nvme netapp smdevices
/dev/nvme1n1, Array Name ICTM0706SYS04, Volume Name NVMe2, NSID 1, Volume ID
000015bd5903df4a00a0980000af4462, Controller A, Access State unknown, 2.15GB
/dev/nvme1n2, Array Name ICTM0706SYS04, Volume Name NVMe3, NSID 2, Volume ID
000015c05903e24000a0980000af4462, Controller A, Access State unknown, 2.15GB
/dev/nvme1n3, Array Name ICTM0706SYS04, Volume Name NVMe4, NSID 4, Volume ID
00001bb0593a46f400a0980000af4462, Controller A, Access State unknown, 2.15GB
/dev/nvme1n4, Array Name ICTM0706SYS04, Volume Name NVMe6, NSID 6, Volume ID
00001696593b424b00a0980000af4112, Controller A, Access State unknown, 2.15GB
/dev/nvme2n1, Array Name ICTM0706SYS04, Volume Name NVMe2, NSID 1, Volume ID
000015bd5903df4a00a0980000af4462, Controller B, Access State unknown, 2.15GB
/dev/nvme2n2, Array Name ICTM0706SYS04, Volume Name NVMe3, NSID 2, Volume ID
000015c05903e24000a0980000af4462, Controller B, Access State unknown, 2.15GB
/dev/nvme2n3, Array Name ICTM0706SYS04, Volume Name NVMe4, NSID 4, Volume ID
00001bb0593a46f400a0980000af4462, Controller B, Access State unknown, 2.15GB
/dev/nvme2n4, Array Name ICTM0706SYS04, Volume Name NVMe6, NSID 6, Volume ID
00001696593b424b00a0980000af4112, Controller B, Access State unknown, 2.15GB
```

Set up failover on the host

Multipath software provides a redundant path to the storage array in case one of the physical paths is disrupted. There are currently two methods of multipathing available for NVMe, and which you will be using is going to be dependent on which OS version you are running. For SLES 12, device

mapper multipath (DMMP) will be used. For RHEL 7 and SLES 15, a native NVMe multipathing solution will be used.

Configuring the host to run failover

The SUSE Linux Enterprise Server hosts require additional configuration changes to run failover.

Before you begin

- You have installed the required packages on your system.
- For Red Hat (RHEL) hosts, verify the packages are installed by running `rpm -q device-mapper-multipath`
- For SLES hosts, verify the packages are installed by running `rpm -q multipath-tools`

About this task

By default, DM-MP is disabled in RHEL and SLES. Complete the following steps to enable DM-MP components on the host.

Steps

1. Add the NVMe E-Series device entry to the devices section of the `/etc/multipath.conf` file, as shown in the following example:

```
devices {
    device {
        vendor "NVME"
        product "NetApp E-Series*"
        path_grouping_policy group_by_prio
        failback immediate
        no_path_retry 30
    }
}
```

2. Configure `multipathd` to start at system boot.

```
# systemctl enable multipathd
```

3. Start `multipathd` if it is not currently running.

```
# systemctl start multipathd
```

4. Verify the status of `multipathd` to make sure it is active and running:

```
# systemctl status multipathd
```

Accessing NVMe Volumes

You can configure the I/O directed to the device target based on your Linux version.

Accessing NVMe volumes for virtual device targets

For RHEL 7 and SLES 12, I/O is directed to virtual device targets by the Linux host. DM-MP manages the physical paths underlying these virtual targets.

Virtual devices are I/O targets

Make sure you are running I/O only to the virtual devices created by DM-MP and not to the physical device paths. If you are running I/O to the physical paths, DM-MP cannot manage a failover event and the I/O fails.

You can access these block devices through the `dm` device or the `symlink` in `/dev/mapper`, for example:

```
/dev/dm-1
/dev/mapper/eui.00001bc7593b7f5f00a0980000af4462
```

Example

The following example output from the `nvme list` command shows the host node name and its correlation with the namespace ID.

```

NODE          SN          MODEL          NAMESPACE
/dev/nvme1n1  021648023072 NetApp E-Series 10
/dev/nvme1n2  021648023072 NetApp E-Series 11
/dev/nvme1n3  021648023072 NetApp E-Series 12
/dev/nvme1n4  021648023072 NetApp E-Series 13
/dev/nvme2n1  021648023151 NetApp E-Series 10
/dev/nvme2n2  021648023151 NetApp E-Series 11
/dev/nvme2n3  021648023151 NetApp E-Series 12
/dev/nvme2n4  021648023151 NetApp E-Series 13

```

Column	Description
Node	<p>The node name includes two parts:</p> <ul style="list-style-type: none"> The notation <code>nvme1</code> represents controller A and <code>nvme2</code> represents controller B. The notation <code>n1</code>, <code>n2</code>, and so on represent the namespace identifier from the host perspective. These identifiers are repeated in the table, once for controller A and once for controller B.
Namespace	The Namespace column lists the namespace ID (NSID), which is the identifier from the storage array perspective.

In the following `multipath -ll` output, the optimized paths are shown with a `prio` value of 50, while the non-optimized paths are shown with a `prio` value of 10.

The Linux operating system routes I/O to the path group that is shown as `status=active`, while the path groups listed as `status=enabled` are available for failover.

```

eui.00001bc7593b7f5f00a0980000af4462 dm-0 NVME,NetApp E-Series
size=15G features='1 queue_if_no_path' hwhandler='0' wp=rw
|+- policy='service-time 0' prio=50 status=active
|`- #:#:#:# nvme1n1 259:5 active ready running
`+- policy='service-time 0' prio=10 status=enabled
  `- #:#:#:# nvme2n1 259:9 active ready running

eui.00001bc7593b7f5f00a0980000af4462 dm-0 NVME,NetApp E-Series
size=15G features='1 queue_if_no_path' hwhandler='0' wp=rw
|+- policy='service-time 0' prio=0 status=enabled
|`- #:#:#:# nvme1n1 259:5 failed faulty running
`+- policy='service-time 0' prio=10 status=active
  `- #:#:#:# nvme2n1 259:9 active ready running

```


Line item	Description
policy='service-time 0' prio=50 status=active	This line and the following line show that <code>nvme1n1</code> , which is the namespace with an NSID of 10, is optimized on the path with a <code>prio</code> value of 50 and a <code>status</code> value of <code>active</code> . This namespace is owned by controller A.
policy='service-time 0' prio=10 status=enabled	This line shows the failover path for namespace 10, with a <code>prio</code> value of 10 and a <code>status</code> value of <code>enabled</code> . I/O is not being directed to the namespace on this path at the moment. This namespace is owned by controller B.
policy='service-time 0' prio=0 status=enabled	This example shows <code>multipath -ll</code> output from a different point in time, while controller A is rebooting. The path to namespace 10 is shown as <code>failed faulty</code> running with a <code>prio</code> value of 0 and a <code>status</code> value of <code>enabled</code> .
policy='service-time 0' prio=10 status=active	Note that the <code>active</code> path refers to <code>nvme2</code> , so the I/O is being directed on this path to controller B.

Accessing NVMe volumes for physical NVMe device targets

For SLES 15, I/O is directed to the physical NVMe device targets by the Linux host. A native NVMe multipathing solution manages the physical paths underlying the single apparent physical device displayed by the host.

Note: It is best practice to use the links in `/dev/disk/by-id/` rather than `/dev/nvme0n1`, for example:

```
# ls /dev/disk/by-id/ -l lrwxrwxrwx 1 root root 13 Oct 18 15:14
nvme-
eui.0000320f5cad32cf00a0980000af4112 -> ../../nvme0n1
```

Physical NVMe devices are I/O targets

Run I/O to the physical nvme device path. There should only be one of these devices present for each namespace using the following format:

```
/dev/nvme[sys#]n[id#]
```

All paths are virtualized using the native multipathing solution underneath this device.

You can view your paths by running:

```
# nvme list-subsys
```

Example output:

```
nvme-subsys0 - NQN=nqn.1992-08.com.netapp:5700.600a098000a522500000000589aa8a6
\
+- nvme0 rdma traddr=192.4.21.131 trsvcid=4420 live
+- nvme1 rdma traddr=192.4.22.141 trsvcid=4420 live
```

If you specify a namespace device when using the `'nvme list-subsys'` command, it provides additional information about the paths to that namespace:

```
# nvme list-subsys /dev/nvme0n1
nvme-subsys0 - NQN=nqn.1992-08.com.netapp:5700.600a098000af44620000000058d5dd96
\
+- nvme0 rdma traddr=192.168.130.101 trsvcid=4420 live non-optimized
```

```
+-- nvme1 rdma traddr=192.168.131.101 trsvcid=4420 live non-optimized
+-- nvme2 rdma traddr=192.168.130.102 trsvcid=4420 live optimized
+-- nvme3 rdma traddr=192.168.131.102 trsvcid=4420 live optimized
```

There are also hooks into the multipath commands to allow you to view your path information for native failover through them as well:

```
#multipath -ll
```

Note: To view the path information, the following must be set in `/etc/multipath.conf`:

```
defaults {
    enable_foreign nvme
}
```

Example output:

```
eui.0000a0335c05d57a00a0980000a5229d [nvme]:nvme0n9 NVMe,Netapp E-Series,08520001
size=4194304 features='n/a' hwhandler='ANA' wp=rw
|-- policy='n/a' prio=50 status=optimized
|  `-- 0:0:1 nvme0c0n1 0:0 n/a optimized    live
|-- policy='n/a' prio=10 status=non-optimized
|  `-- 0:1:1 nvme0c1n1 0:0 n/a non-optimized  live
```

Create filesystems

You create a file system on the namespace or native nvme device and mount the filesystem.

Create filesystems (RHEL 7 and SLES 12)

For RHEL 7 and SLES 12, you create a file system on the namespace and mount the filesystem.

Steps

1. Run the `multipath -ll` command to get a list of `/dev/mapper/dm` devices.

```
# multipath -ll
```

The result of this command shows two devices, `dm-19` and `dm-16` :

```
eui.00001ffe5a94ff8500a0980000af4444 dm-19 NVME,NetApp E-Series
size=10G features='1 queue_if_no_path' hwhandler='0' wp=rw
|-- policy='service-time 0' prio=50 status=active
|  |-- #:#:#:# nvme0n19 259:19  active ready running
|  |-- #:#:#:# nvme1n19 259:115 active ready running
|  `-- #:#:#:# nvme2n19 259:51  active ready running
|-- policy='service-time 0' prio=10 status=enabled
|  |-- #:#:#:# nvme0n16 259:16  active ready running
|  |-- #:#:#:# nvme1n16 259:112 active ready running
|  `-- #:#:#:# nvme2n16 259:48  active ready running
|-- policy='service-time 0' prio=10 status=enabled
|  |-- #:#:#:# nvme0n16 259:16  active ready running
|  |-- #:#:#:# nvme1n16 259:112 active ready running
|  `-- #:#:#:# nvme2n16 259:48  active ready running
|-- policy='service-time 0' prio=10 status=enabled
|  |-- #:#:#:# nvme0n16 259:16  active ready running
|  |-- #:#:#:# nvme1n16 259:112 active ready running
|  `-- #:#:#:# nvme2n16 259:48  active ready running
```

2. Create a file system on the partition for each `/dev/mapper/eui-` device.

The method for creating a file system varies depending on the file system chosen. This example shows creating an ext4 file system.

```
# mkfs.ext4 /dev/mapper/dm-19
mke2fs 1.42.11 (09-Jul-2014)
Creating filesystem with 2620928 4k blocks and 655360 inodes
Filesystem UUID: 97f987e9-47b8-47f7-b434-bf3ebbe826d0
Superblock backups stored on blocks:
    32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632

Allocating group tables: done
Writing inode tables: done
Creating journal (32768 blocks): done
Writing superblocks and filesystem accounting information: done
```

3. Create a folder to mount the new device.

```
# mkdir /mnt/ext4
```

4. Mount the device.

```
# mount /dev/mapper/eui.00001ffe5a94ff8500a0980000af4444 /mnt/ext4
```

Create filesystems (SLES 15)

For SLES 15, you create a filesystem on the native nvme device and mount the filesystem.

Steps

1. Run the `multipath -ll` command to get a list of `/dev/nvme` devices.

```
# multipath -ll
```

The result of this command shows device `nvme0n6`.

```
eui.000082dd5c05d39300a0980000a52225 [nvme]:nvme0n6 NVMe,NetApp E-Series,08520000
size=4194304 features='n/a' hwhandler='ANA' wp=rw
|+- policy='n/a' prio=50 status=optimized
|  ` 0:0:1 nvme0c0n1 0:0 n/a optimized      live
|+- policy='n/a' prio=50 status=optimized
|  ` 0:1:1 nvme0c1n1 0:0 n/a optimized      live
|+- policy='n/a' prio=10 status=non-optimized
|  ` 0:2:1 nvme0c2n1 0:0 n/a non-optimized live
`+- policy='n/a' prio=10 status=non-optimized
   ` 0:3:1 nvme0c3n1 0:0 n/a non-optimized live
```

2. Create a file system on the partition for each `/dev/nvme0n#` device.

The method for creating a file system varies depending on the file system chosen. This example shows creating an ext4 file system.

```
# mkfs.ext4 /dev/disk/by-id/nvme-eui.000082dd5c05d39300a0980000a52225
mke2fs 1.42.11 (22-Oct-2019)
Creating filesystem with 2620928 4k blocks and 655360 inodes
Filesystem UUID: 97f987e9-47b8-47f7-b434-bf3ebbe826d0
Superblock backups stored on blocks:
    32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632

Allocating group tables: done
```

```
Writing inode tables: done
Creating journal (32768 blocks): done
Writing superblocks and filesystem accounting information: done
```

3. Create a folder to mount the new device.

```
# mkdir /mnt/ext4
```

4. Mount the device.

```
# mount /dev/disk/by-id/nvme-eui.000082dd5c05d39300a0980000a52225 /mnt/ext4
```

Verify storage access on the host

Before using the namespace, you verify that the host can write data to the namespace and read it back.

Steps

1. On the host, copy one or more files to the mount point of the disk.
2. Copy the files back to a different folder on the original disk.
3. Run the `diff` command to compare the copied files to the originals.

After you finish

You remove the file and folder that you copied.

Record NVMe over RoCE-specific information for Linux


Select the NVMe over RoCE worksheet to record your protocol-specific storage configuration information. You need this information to perform provisioning tasks.

NVMe over RoCE worksheet for Linux

You can use this worksheet to record NVMe over RoCE storage configuration information. You need this information to perform provisioning tasks.

Direct connect topology

In a direct connect topology, one or more hosts are directly connected to the subsystem. In the SANtricity OS 11.50 release, we support a single connection from each host to a subsystem controller, as shown below. In this configuration, one HCA (host channel adapter) port from each host should be on the same subnet as the E-Series controller port it is connected to, but on a different subnet from the other HCA port.




An example configuration that satisfies the requirements consists of four network subnets as follows:

- Subnet 1: Host 1 HCA Port 1 and Controller 1 Host port 1
- Subnet 2: Host 1 HCA Port 2 and Controller 2 Host port 1
- Subnet 3: Host 2 HCA Port 1 and Controller 1 Host port 2
- Subnet 4: Host 2 HCA Port 2 and Controller 2 Host port 2

Switch connect topology

In a fabric topology, one or more switches are used. Refer to [NetApp Interoperability Matrix Tool](#) for a list of supported switches.



NVMe over RoCE: Host identifiers

Locate and document the initiator NQN from each host.

Host port connections	Software initiator NQN
Host (initiator) 1	
Host (initiator) 2	

NVMe over RoCE: Target NQN

Document the target NQN for the storage array.

Array name	Target NQN
Array controller (target)	

NVMe over RoCE: Target NQNs

Document the NQNs to be used by the array ports.

Array controller (target) port connections	NQN
Controller A, port 1	
Controller B, port 1	
Controller A, port 2	
Controller B, port 2	

NVMe over RoCE: Mapping host name

Note: The mapping host name is created during the workflow.

Mapping host name	
Host OS type	

NVMe over Fibre Channel Express Setup

You can use NVMe with the Fibre Channel protocol.

Steps

1. [Verify the Linux configuration is supported](#) on page 103
2. [Configuring IP addresses using DHCP](#) on page 104
3. [Install SANtricity Storage Manager for SMcli \(SANtricity software version 11.53 or earlier\)](#) on page 105
4. [Access SANtricity System Manager and use Setup wizard](#) on page 106
5. [Configure the FC switches](#) on page 107
6. [Set up NVMe over Fibre Channel on the host side](#) on page 107
7. [Define a host](#) on page 108
8. [Assign a volume](#) on page 109
9. [Display the volumes visible to the host](#) on page 110
10. [Set up failover on the host](#) on page 110
11. [Accessing NVMe volumes](#) on page 111
12. [Create partitions and filesystems](#) on page 116
13. [Verify storage access on the host](#) on page 117
14. [Record NVMe over Fibre Channel-specific information for Linux](#) on page 117

Verify the Linux configuration is supported

To ensure reliable operation, you create an implementation plan and then use the NetApp Interoperability Matrix Tool (IMT) to verify that the entire configuration is supported.

Steps

1. Go to the [NetApp Interoperability Matrix Tool](#).
2. Click on the **Solution Search** tile.
3. In the **Protocols > SAN Host** area, click the **Add** button next to **E-Series SAN Host**.
4. Click **View Refine Search Criteria**.
The **Refine Search Criteria** section is displayed. In this section you may select the protocol that applies, as well as other criteria for the configuration such as Operating System, NetApp OS, and Host Multipath driver. Select the criteria you know you want for your configuration, and then see what compatible configuration elements apply. As necessary, make the updates for your operating system and protocol that are prescribed in the tool. Detailed information for your chosen configuration is accessible on the View Supported Configurations page by clicking the **right page arrow**.

NVMe over Fibre Channel restrictions

Before using NVMe over Fibre Channel, review the controller, host, and recovery restrictions.

Verify your configuration

Verify your configuration, using the [NetApp Interoperability Matrix Tool](#).

About the hardware

NVMe over Fibre Channel can be configured for the EF300, EF600, EF570, and E5700 controllers. The controllers must have the quad 32GB host port.

Restrictions

For an up-to-date listing of all restrictions, see the [NetApp Interoperability Matrix Tool](#) for a complete list of requirements.

Controller restrictions

- This protocol can be used only for EF300, EF600, EF570, or E5700 controller with a minimum of 32 GB of physical memory. If the minimum memory requirements for the controllers are not met during start of day operations, a message is displayed that helps you diagnose the problem.
- No simplex (single-controller) configurations are supported.
- The only supported host interface card (HIC) is the quad 32GB Fibre Channel HIC.

Host, host protocol, and host operating system restrictions

- For an up-to-date listing of all compatible host, host protocol, and host operating systems, see the [NetApp Interoperability Matrix Tool](#) for a complete list of requirements.
- For a list of supported host channel adapters see the [NetApp Interoperability Matrix Tool](#).

Storage and disaster recovery restrictions

- Asynchronous and synchronous mirroring are not supported.
- Thin provisioning (the creation of thin volumes) is not supported.

Configuring IP addresses using DHCP

In this express method for configuring communications between the management station and the storage array, you use Dynamic Host Configuration Protocol (DHCP) to provide IP addresses. Each storage array has either one controller (simplex) or two controllers (duplex), and each controller has two storage management ports. Each management port will be assigned an IP address.

Before you begin

You have installed and configured a DHCP server on the same subnet as the storage management ports.

About this task

The following instructions refer to a storage array with two controllers (a duplex configuration).

1. If you have not already done so, connect an Ethernet cable to the management station and to management port 1 on each controller (A and B).

The DHCP server assigns an IP address to port 1 of each controller.

Note: Do not use management port 2 on either controller. Port 2 is reserved for use by NetApp technical personnel.

Important: If you disconnect and reconnect the Ethernet cable, or if the storage array is power-cycled, DHCP assigns IP addresses again. This process occurs until static IP addresses are configured. you should avoid disconnecting the cable or power-cycling the array.

If the storage array cannot get DHCP-assigned IP addresses within 30 seconds, the following default IP addresses are set:

- Controller A, port 1: 169.254.128.101
 - Controller B, port 1: 169.254.128.102
 - Subnet mask: 255.255.0.0
2. Locate the MAC address label on the back of each controller, and then provide your network administrator with the MAC address for port 1 of each controller.

Your network administrator needs the MAC addresses to determine the IP address for each controller. You will need the IP addresses to connect to your storage system through your browser.

Install SANtricity Storage Manager for SMcli (SANtricity software version 11.53 or earlier)

When you install the SANtricity Storage Manager software on your management station, a graphical user interface (GUI) and a command line interface (CLI) are installed by default. These instructions assume that you will install the SANtricity Storage Manager GUI on a management station and *not* on the I/O host.

Before you begin

Important: For SANtricity software 11.60 and newer, the SANtricity Secure CLI (SMcli) is included in the SANtricity OS and downloadable through the SANtricity System Manager. For more information on how to download the SMcli through the SANtricity System Manager, refer to the *Download the command line interface (CLI)* topic under the SANtricity System Manager Online Help.

- You are using SANtricity software 11.53 or earlier.
- You must have the correct administrator or superuser privileges.
- You must have ensured that the system that will contain the SANtricity Storage Manager client has the following minimum requirements:
 - **RAM:** 2 GB for Java Runtime Engine
 - **Disk space:** 5 GB
 - **OS/Architecture:** Refer to [NetApp Support Downloads > Software > E-Series/EF-Series SANtricity Storage Manager](#) for guidance on determining the supported operating system versions and architectures.

About this task

This section describes how to install SANtricity Storage Manager on both the Windows and Linux OS platforms, because both Windows and Linux are common management station platforms when Linux is used for the data host.

Steps

1. Download the SANtricity software release from [NetApp Support Downloads > Software > E-Series/EF-Series SANtricity Storage Manager](#).
2. Run the SANtricity installer.

Windows	Linux
Double-click the SMIA* .exe installation package to start the installation.	<ol style="list-style-type: none">a. Go to the directory where the SMIA* .bin installation package is located.b. If the temp mount point does not have execute permissions, set the IATEMPDIR variable. Example: IATEMPDIR=/root ./SMIA-LINUXx64-11.25.0A00.0002.binc. Run the <code>chmod +x SMIA*.bin</code> command to grant execute permission to the file.d. Run the <code>./SMIA*.bin</code> command to start the installer.

3. Use the installation wizard to install the software on the management station.

Access SANtricity System Manager and use Setup wizard

You use the Setup wizard in SANtricity System Manager to configure your storage array.

Before you begin

- You have ensured that the device from which you will access SANtricity System Manager contains one of the following browsers:

Browser	Minimum version
Google Chrome	47
Microsoft Internet Explorer	11
Microsoft Edge	EdgeHTML 12
Mozilla Firefox	31
Safari	9

- You are using out-of-band management.

About this task

The wizard automatically relaunches when you open System Manager or refresh your browser and *at least one* of the following conditions is met:

- No pools and volume groups are detected.
- No workloads are detected.
- No notifications are configured.

Steps

- From your browser, enter the following URL:

https://<DomainNameOrIPAddress>

IPAddress is the address for one of the storage array controllers.

The first time SANtricity System Manager is opened on an array that has not been configured, the **Set Administrator Password** prompt appears. Role-based access management configures four local roles: admin, support, security, and monitor. The latter three roles have random passwords that cannot be guessed. After you set a password for the admin role you can change all of the passwords using the admin credentials. See *SANtricity System Manager online help* for more information on the four local user roles.

- Enter the System Manager password for the admin role in the **Set Administrator Password** and **Confirm Password** fields, and then select the **Set Password** button. When you open System Manager and no pools, volumes groups, workloads, or notifications have been configured, the Setup wizard launches.
- Use the Setup wizard to perform the following tasks:
 - Verify hardware (controllers and drives)** – Verify the number of controllers and drives in the storage array. Assign a name to the array.
 - Verify hosts and operating systems** – Verify the host and operating system types that the storage array can access.
 - Accept pools** – Accept the recommended pool configuration for the express installation method. A pool is a logical group of drives.
 - Configure alerts** – Allow System Manager to receive automatic notifications when a problem occurs with the storage array.
 - Enable AutoSupport** – Automatically monitor the health of your storage array and have dispatches sent to technical support.

4. If you have not already created a volume, create one by going to **Storage > Volumes > Create > Volume**.

For more information, see the online help for SANtricity System Manager.

Configure the FC switches

Configuring (zoning) the Fibre Channel (FC) switches enables the hosts to connect to the storage array and limits the number of paths. You zone the switches using the management interface for the switches.

Before you begin

- You must have administrator credentials for the switches.
- You must have used your HBA utility to discover the WWPN of each host initiator port and of each controller target port connected to the switch.

About this task

For details about zoning your switches, see the switch vendor's documentation.

Each initiator port must be in a separate zone with all of its corresponding target ports.

Steps

1. Log in to the FC switch administration program, and then select the zoning configuration option.
2. Create a new zone that includes the first host initiator port and that also includes all of the target ports that connect to the same FC switch as the initiator.
3. Create additional zones for each FC host initiator port in the switch.
4. Save the zones, and then activate the new zoning configuration.

Set up NVMe over Fibre Channel on the host side

NVMe initiator configuration in a Fibre Channel environment includes installing and configuring the `nvme-cli` package, and enabling the NVMe/FC initiator on the host.

About this task

These are the instructions for SUSE Linux Enterprise Server 12 SP4 and 32GB FC HBAs.

Steps

1. Install the `nvme-cli` package:

```
# zypper install nvme-cli
```

2. Enable and start the `nvmeofc-boot-connections` service.

```
systemctl enable nvmeofc-boot-connections.service
```

```
systemctl start nvmeofc-boot-connections.service
```

3. Set `lpfc_enable_fc4_type` to 3 to enable SLES12 SP4 as an NVMe/FC initiator.

```
# cat /etc/modprobe.d/lpfc.conf
options lpfc lpfc_enable_fc4_type=3
```

4. Re-build the initrd to get the Emulex change and the boot parameter change.

```
# dracut --force
```

5. Reboot the host to reconfigure the `lpfc` driver.

```
# reboot
```

The host is rebooted and the NVMe/FC initiator is enabled on the host.

Note: After completing the host side setup, configuration of the NVMe over Fibre Channel ports occur automatically.

Define a host

Using SANtricity System Manager, you define the hosts that send data to the storage array. Defining a host is one of the steps required to let the storage array know which hosts are attached to it and to allow I/O access to the volumes.

About this task

Keep these guidelines in mind when you define a host:

- You must define the host identifier ports that are associated with the host.
- Make sure that you provide the same name as the host's assigned system name.
- This operation does not succeed if the name you choose is already in use.
- The length of the name cannot exceed 30 characters.

Steps

1. Select **Storage > Hosts**.
2. Click **Create > Host**.
The Create Host dialog box appears.
3. Select the settings for the host as appropriate.

Field details

Setting	Description
Name	Type a name for the new host.
Host operating system type	Select one of the following options from the drop-down list: <ul style="list-style-type: none">• SANtricity 11.60 and newer <i>Linux</i>• Pre SANtricity 11.60 <i>Linux DM-MP (Kernel 3.10 or later)</i>
Host interface type	Select the host interface type that you want to use. If the array you configure only has one available host interface type, this setting might not be available to select.

Setting	Description
Host ports	<p>Do one of the following:</p> <ul style="list-style-type: none"> • Select I/O Interface If the host ports have logged in, you can select host port identifiers from the list. This is the recommended method. • Manual add If the host ports have not logged in, look at <code>/etc/nvme/hostnqn</code> on the host to find the hostnqn identifiers and associate them with the host definition. You can manually enter the host port identifiers or copy/paste them from the <code>/etc/nvme/hostnqn</code> file (one at a time) into the Host ports field. You must add one host port identifier at a time to associate it with the host, but you can continue to select as many identifiers that are associated with the host. Each identifier is displayed in the Host ports field. If necessary, you also can remove an identifier by selecting the X next to it.

4. Click **Create**.

Result

After the host is successfully created, SANtricity System Manager creates a default name for each host port configured for the host.

The default alias is `<Hostname_Port Number>`. For example, the default alias for the first port created for host `IPT` is `IPT_1`.

Assign a volume

You must assign a volume (namespace) to a host or host cluster so it can be used for I/O operations. This assignment grants a host or host cluster access to one or more namespaces in a storage array.

About this task

Keep these guidelines in mind when you assign volumes:

- You can assign a volume to only one host or host cluster at a time.
- Assigned volumes are shared between controllers in the storage array.
- The same namespace ID (NSID) cannot be used twice by a host or a host cluster to access a volume. You must use a unique NSID.

Assigning a volume fails under these conditions:

- All volumes are assigned.
- The volume is already assigned to another host or host cluster.

The ability to assign a volume is unavailable under these conditions:

- No valid hosts or host clusters exist.
- All volume assignments have been defined.

All unassigned volumes are displayed, but functions for hosts with or without Data Assurance (DA) apply as follows:

- For a DA-capable host, you can select volumes that are either DA-enabled or not DA-enabled.

- For a host that is not DA-capable, if you select a volume that is DA-enabled, a warning states that the system must automatically turn off DA on the volume before assigning the volume to the host.

Steps

1. Select **Storage > Hosts**.
2. Select the host or host cluster to which you want to assign volumes, and then click **Assign Volumes**.
A dialog box appears that lists all the volumes that can be assigned. You can sort any of the columns or type something in the **Filter** box to make it easier to find particular volumes.
3. Select the checkbox next to each volume that you want to assign or select the checkbox in the table header to select all volumes.
4. Click **Assign** to complete the operation.

Result

After successfully assigning a volume or volumes to a host or a host cluster, the system performs the following actions:

- The assigned volume receives the next available NSID. The host uses the NSID to access the volume.
- The user-supplied volume name appears in volume listings associated to the host.

Display the volumes visible to the host

The `SMdevices` tool, part of the `nvme-cli` package, allows you to view the volumes currently visible on the host. This tool is an alternative to the `nvme list` command.

Step

To view information about each NVMe path to an E-Series volume, use the `nvme netapp smdevices [-o <format>]` command. The output `<format>` can be normal (the default if `-o` is not used), column, or json.

```
# nvme netapp smdevices
/dev/nvme1n1, Array Name ICTM0706SYS04, Volume Name NVMe2, NSID 1, Volume ID
000015bd5903df4a00a0980000af4462, Controller A, Access State unknown, 2.15GB
/dev/nvme1n2, Array Name ICTM0706SYS04, Volume Name NVMe3, NSID 2, Volume ID
000015c05903e24000a0980000af4462, Controller A, Access State unknown, 2.15GB
/dev/nvme1n3, Array Name ICTM0706SYS04, Volume Name NVMe4, NSID 4, Volume ID
00001bb0593a46f400a0980000af4462, Controller A, Access State unknown, 2.15GB
/dev/nvme1n4, Array Name ICTM0706SYS04, Volume Name NVMe6, NSID 6, Volume ID
00001696593b424b00a0980000af4112, Controller A, Access State unknown, 2.15GB
/dev/nvme2n1, Array Name ICTM0706SYS04, Volume Name NVMe2, NSID 1, Volume ID
000015bd5903df4a00a0980000af4462, Controller B, Access State unknown, 2.15GB
/dev/nvme2n2, Array Name ICTM0706SYS04, Volume Name NVMe3, NSID 2, Volume ID
000015c05903e24000a0980000af4462, Controller B, Access State unknown, 2.15GB
/dev/nvme2n3, Array Name ICTM0706SYS04, Volume Name NVMe4, NSID 4, Volume ID
00001bb0593a46f400a0980000af4462, Controller B, Access State unknown, 2.15GB
/dev/nvme2n4, Array Name ICTM0706SYS04, Volume Name NVMe6, NSID 6, Volume ID
00001696593b424b00a0980000af4112, Controller B, Access State unknown, 2.15GB
```

Set up failover on the host

Multipath software provides a redundant path to the storage array in case one of the physical paths is disrupted. The multipath software presents the operating system with a single virtual device that represents the active physical paths to the storage. The multipath software also manages the

failover process that updates the virtual device. You use the device mapper multipath (DM-MP) tool for Linux installations.

Configuring the host to run failover

The SUSE Linux Enterprise Server host requires configuration changes to run failover. The failover solution uses DM-MP.

Before you begin

- You have installed the required packages on your system.
- For Red Hat (RHEL) hosts, verify the packages are installed by running `rpm -q device-mapper-multipath`
- For SLES hosts, verify the packages are installed by running `rpm -q multipath-tools`

About this task

By default, DM-MP is disabled in RHEL and SLES. Complete the following steps to enable DM-MP components on the host.

Steps

1. Add the NVMe E-Series device entry to the devices section of the `/etc/multipath.conf` file, as shown in the following example:

```
devices {
    device {
        vendor "NVME"
        product "NetApp E-Series*"
        path_grouping_policy group_by_prio
        failback immediate
        no_path_retry 30
    }
}
```

2. Configure `multipathd` to start at system boot.

```
# systemctl enable multipathd
```

3. Start `multipathd` if it is not currently running.

```
# systemctl start multipathd
```

4. Verify the status of `multipathd` to make sure it is active and running:

```
# systemctl status multipathd
```

Accessing NVMe volumes

I/O is directed to virtual device targets by the Linux host. The DM-MP manages the physical paths underlying these virtual targets.

Virtual devices are I/O targets

Make sure you are running I/O only to the virtual devices created by DM-MP and not to the physical device paths. If you are running I/O to the physical paths, DM-MP cannot manage a failover event and the I/O fails.

To run raw I/O, make sure your targets are virtual devices created by DM-MP to manage the physical paths. You can access these block devices through the `dm` device or the `symlink` in `/dev/mapper`, for example:

```
/dev/dm-1  
/dev/mapper/eui.00001bc7593b7f5f00a0980000af4462
```

Example

The following example output from the `nvme list` command shows the host node name and its correlation with the namespace ID.

NODE	SN	MODEL	NAMESPACE
/dev/nvme1n1	021648023072	NetApp E-Series	10
/dev/nvme1n2	021648023072	NetApp E-Series	11
/dev/nvme1n3	021648023072	NetApp E-Series	12
/dev/nvme1n4	021648023072	NetApp E-Series	13
/dev/nvme2n1	021648023151	NetApp E-Series	10
/dev/nvme2n2	021648023151	NetApp E-Series	11
/dev/nvme2n3	021648023151	NetApp E-Series	12
/dev/nvme2n4	021648023151	NetApp E-Series	13

Column	Description
Node	The node name includes two parts: <ul style="list-style-type: none">• The notation <code>nvme1</code> represents controller A and <code>nvme2</code> represents controller B.• The notation <code>n1</code>, <code>n2</code>, and so on represent the namespace identifier from the host perspective. These identifiers are repeated in the table, once for controller A and once for controller B.
Namespace	The Namespace column lists the namespace ID (NSID), which is the identifier from the storage array perspective.

In the following `multipath -ll` output, the optimized paths are shown with a `prio` value of 50, while the non-optimized paths are shown with a `prio` value of 10.

The Linux operating system routes I/O to the path group that is shown as `status=active`, while the path groups listed as `status=enabled` are available for failover.

```
eui.00001bc7593b7f5f00a0980000af4462 dm-0 NVME,NetApp E-Series  
size=15G features='1 queue_if_no_path' hwhandler='0' wp=rw  
|+- policy='service-time 0' prio=50 status=active  
|`- #:#:#:# nvme1n1 259:5 active ready running  
`+- policy='service-time 0' prio=10 status=enabled  
  `- #:#:#:# nvme2n1 259:9 active ready running  
  
eui.00001bc7593b7f5f00a0980000af4462 dm-0 NVME,NetApp E-Series  
size=15G features='1 queue_if_no_path' hwhandler='0' wp=rw  
|+- policy='service-time 0' prio=0 status=enabled  
|`- #:#:#:# nvme1n1 259:5 failed faulty running  
`+- policy='service-time 0' prio=10 status=active  
  `- #:#:#:# nvme2n1 259:9 active ready running
```


Line item	Description
<code>policy='service-time 0' prio=50 status=active</code>	This line and the following line show that <code>nvme1n1</code> , which is the namespace with an NSID of 10, is optimized on the path with a <code>prio</code> value of 50 and a <code>status</code> value of <code>active</code> . This namespace is owned by controller A.
<code>policy='service-time 0' prio=10 status=enabled</code>	This line shows the failover path for namespace 10, with a <code>prio</code> value of 10 and a <code>status</code> value of <code>enabled</code> . I/O is not being directed to the namespace on this path at the moment. This namespace is owned by controller B.
<code>policy='service-time 0' prio=0 status=enabled</code>	This example shows <code>multipath -ll</code> output from a different point in time, while controller A is rebooting. The path to namespace 10 is shown as <code>failed faulty</code> running with a <code>prio</code> value of 0 and a <code>status</code> value of <code>enabled</code> .
<code>policy='service-time 0' prio=10 status=active</code>	Note that the <code>active</code> path refers to <code>nvme2</code> , so the I/O is being directed on this path to controller B.

Accessing NVMe volumes for virtual device targets

Virtual devices are I/O targets

Make sure you are running I/O only to the virtual devices created by DM-MP and not to the physical device paths. If you are running I/O to the physical paths, DM-MP cannot manage a failover event and the I/O fails.

You can access these block devices through the `dm` device or the `symlink` in `/dev/mapper`, for example:

```
/dev/dm-1  
/dev/mapper/eui.00001bc7593b7f5f00a0980000af4462
```

Example

The following example output from the `nvme list` command shows the host node name and its correlation with the namespace ID.

NODE	SN	MODEL	NAMESPACE
/dev/nvme1n1	021648023072	NetApp E-Series	10
/dev/nvme1n2	021648023072	NetApp E-Series	11
/dev/nvme1n3	021648023072	NetApp E-Series	12
/dev/nvme1n4	021648023072	NetApp E-Series	13
/dev/nvme2n1	021648023151	NetApp E-Series	10
/dev/nvme2n2	021648023151	NetApp E-Series	11
/dev/nvme2n3	021648023151	NetApp E-Series	12
/dev/nvme2n4	021648023151	NetApp E-Series	13

Column	Description
Node	The node name includes two parts: <ul style="list-style-type: none">The notation <code>nvme1</code> represents controller A and <code>nvme2</code> represents controller B.The notation <code>n1</code>, <code>n2</code>, and so on represent the namespace identifier from the host perspective. These identifiers are repeated in the table, once for controller A and once for controller B.

Column	Description
Namespace	The Namespace column lists the namespace ID (NSID), which is the identifier from the storage array perspective.

In the following `multipath -ll` output, the optimized paths are shown with a `prio` value of 50, while the non-optimized paths are shown with a `prio` value of 10.

The Linux operating system routes I/O to the path group that is shown as `status=active`, while the path groups listed as `status=enabled` are available for failover.

```
eui.00001bc7593b7f500a0980000af4462 dm-0 NVME,NetApp E-Series
size=15G features='1 queue_if_no_path' hwhandler='0' wp=rw
|+- policy='service-time 0' prio=50 status=active
|  `- #:#:#:# nvme1n1 259:5 active ready running
`+- policy='service-time 0' prio=10 status=enabled
   `- #:#:#:# nvme2n1 259:9 active ready running
```

```
eui.00001bc7593b7f5f00a0980000af4462 dm-0 NVME,NetApp E-Series
size=15G features='1 queue_if_no_path' hwhandler='0' wp=rw
|+- policy='service-time 0' prio=0 status=enabled
|  `- #:#:#:# nvme1n1 259:5 failed faulty running
`+- policy='service-time 0' prio=10 status=active
   `- #:#:#:# nvme2n1 259:9 active ready running
```

Line item	Description
<code>policy='service-time 0' prio=50 status=active</code>	This line and the following line show that <code>nvme1n1</code> , which is the namespace with an NSID of 10, is optimized on the path with a <code>prio</code> value of 50 and a <code>status</code> value of <code>active</code> . This namespace is owned by controller A.
<code>policy='service-time 0' prio=10 status=enabled</code>	This line shows the failover path for namespace 10, with a <code>prio</code> value of 10 and a <code>status</code> value of <code>enabled</code> . I/O is not being directed to the namespace on this path at the moment. This namespace is owned by controller B.
<code>policy='service-time 0' prio=0 status=enabled</code>	This example shows <code>multipath -ll</code> output from a different point in time, while controller A is rebooting. The path to namespace 10 is shown as <code>failed faulty running</code> with a <code>prio</code> value of 0 and a <code>status</code> value of <code>enabled</code> .
<code>policy='service-time 0' prio=10 status=active</code>	Note that the active path refers to <code>nvme2</code> , so the I/O is being directed on this path to controller B.

Accessing NVMe volumes for physical NVMe device targets

For SLES 15, I/O is directed to the physical NVMe device targets by the Linux host. A native NVMe multipathing solution manages the physical paths underlying the single apparent physical device displayed by the host.

Note: It is best practice to use the links in `/dev/disk/by-id/` rather than `/dev/nvme0n1`, for example:

```
# ls /dev/disk/by-id/ -l lrwxrwxrwx 1 root root 13 Oct 18 15:14
nvme-
eui.0000320f5cad32cf00a0980000af4112 -> ../../nvme0n1
```

Physical NVMe devices are I/O targets

Run I/O to the physical nvme device path. There should only be one of these devices present for each namespace using the following format:

```
/dev/nvme[sysfs#]n[id#]
```

All paths are virtualized using the native multipathing solution underneath this device.

You can view your paths by running:

```
# nvme list-subsys
```

Example output:

```
nvme-subsys0 - NQN=nqn.1992-08.com.netapp:5700.600a098000a522500000000589aa8a6
\
+- nvme0 rdma traddr=192.4.21.131 trsvcid=4420 live
+- nvme1 rdma traddr=192.4.22.141 trsvcid=4420 live
```

If you specify a namespace device when using the 'nvme list-subsys' command, it provides additional information about the paths to that namespace:

```
# nvme list-subsys /dev/nvme0n1
nvme-subsys0 - NQN=nqn.1992-08.com.netapp:5700.600a098000af44620000000058d5dd96
\
+- nvme0 rdma traddr=192.168.130.101 trsvcid=4420 live non-optimized
+- nvme1 rdma traddr=192.168.131.101 trsvcid=4420 live non-optimized
+- nvme2 rdma traddr=192.168.130.102 trsvcid=4420 live optimized
+- nvme3 rdma traddr=192.168.131.102 trsvcid=4420 live optimized
```

There are also hooks into the multipath commands to allow you to view your path information for native failover through them as well:

```
#multipath -ll
```

Note: To view the path information, the following must be set in `/etc/multipath.conf`:

```
defaults {
    enable_foreign nvme
}
```

Example output:

```
eui.0000a0335c05d57a00a0980000a5229d [nvme]:nvme0n9 NVMe,Netapp E-Series,08520001
size=4194304 features='n/a' hwhandler='ANA' wp=rw
|+- policy='n/a' prio=50 status=optimized
|`- 0:0:1 nvme0c0n1 0:0 n/a optimized live
`+- policy='n/a' prio=10 status=non-optimized
`- 0:1:1 nvme0c1n1 0:0 n/a non-optimized live
```

Create partitions and filesystems

You can create a partition on the multipath device, optionally create a file system on the namespace, and mount the partition.

Steps

1. Run the `multipath -ll` command to get a list of `/dev/mapper/dm` devices.

```
# multipath -ll
```

The result of this command shows two devices, `dm-19` and `dm-16` :

```
eui.00001ffe5a94ff8500a0980000af4444 dm-19 NVME,NetApp E-Series
size=10G features='1 queue_if_no_path' hwhandler='0' wp=rw
|+- policy='service-time 0' prio=50 status=active
|  |- #:#:#:# nvme0n19 259:19  active ready running
|  `-- #:#:#:# nvme1n19 259:115 active ready running
`+- policy='service-time 0' prio=10 status=enabled
   |  |- #:#:#:# nvme2n19 259:51  active ready running
   `-- #:#:#:# nvme3n19 259:83  active ready running
eui.00001fd25a94fef000a0980000af4444 dm-16 NVME,NetApp E-Series
size=16G features='1 queue_if_no_path' hwhandler='0' wp=rw
|+- policy='service-time 0' prio=50 status=active
|  |- #:#:#:# nvme0n16 259:16  active ready running
|  `-- #:#:#:# nvme1n16 259:112 active ready running
`+- policy='service-time 0' prio=10 status=enabled
   |  |- #:#:#:# nvme2n16 259:48  active ready running
   `-- #:#:#:# nvme3n16 259:80  active ready running
```

2. Create a file system on the partition for each `/dev/mapper/eui-` device.

The method for creating a file system varies depending on the file system chosen. This example shows creating an `ext4` file system.

```
# mkfs.ext4 /dev/mapper/dm-19
mke2fs 1.42.11 (09-Jul-2014)
Creating filesystem with 2620928 4k blocks and 655360 inodes
Filesystem UUID: 97f987e9-47b8-47f7-b434-bf3ebbe826d0
Superblock backups stored on blocks:
    32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632

Allocating group tables: done
Writing inode tables: done
Creating journal (32768 blocks): done
Writing superblocks and filesystem accounting information: done
```

3. Create a folder to mount the new device.

```
# mkdir /mnt/ext4
```

4. Mount the device.

```
# mount /dev/mapper/eui.00001ffe5a94ff8500a0980000af4444 /mnt/ext4
```

Verify storage access on the host

Before using the namespace, you verify that the host can write data to the namespace and read it back.

Steps

1. On the host, copy one or more files to the mount point of the disk.
2. Copy the files back to a different folder on the original disk.
3. Run the `diff` command to compare the copied files to the originals.

After you finish

You remove the file and folder that you copied.

Record NVMe over Fibre Channel-specific information for Linux

Select the NVMe over Fibre Channel worksheet to record your protocol-specific storage configuration information. You need this information to perform provisioning tasks.


NVMe over Fibre Channel worksheet for Linux

You can use this worksheet to record NVMe over Fibre Channel storage configuration information. You need this information to perform provisioning tasks.

Direct connect topology

In a direct connect topology, one or more hosts are directly connected to the controller.

Direct Connect Topology




- Host 1 HBA Port 1 and Controller A Host port 1
- Host 1 HBA Port 2 and Controller B Host port 1
- Host 2 HBA Port 1 and Controller A Host port 2
- Host 2 HBA Port 2 and Controller B Host port 2
- Host 3 HBA Port 1 and Controller A Host port 3
- Host 3 HBA Port 2 and Controller B Host port 3
- Host 4 HBA Port 1 and Controller A Host port 4
- Host 4 HBA Port 2 and Controller B Host port 4

Switch connect topology

In a fabric topology, one or more switches are used. See the to [NetApp Interoperability Matrix Tool](#) for a list of supported switches.

Fabric Topology



NVMe over Fibre Channel: Host identifiers

Locate and document the initiator NQN from each host.

Host port connections	Host NQN
Host (initiator) 1	
Host (initiator) 2	

NVMe over Fibre Channel: Target NQN

Document the target NQN for the storage array.

Array name	Target NQN
Array controller (target)	

NVMe over Fibre Channel: Target NQNs

Document the NQNs to be used by the array ports.

Array controller (target) port connections	NQN
Controller A, port 1	
Controller B, port 1	
Controller A, port 2	
Controller B, port 2	

NVMe over Fibre Channel: Mapping host name

Note: The mapping host name is created during the workflow.

Mapping host name	
Host OS type	

Where to find additional information - Linux

Use the resources listed here if you need additional information. You can also use the online help for SANtricity System Manager .

- [Linux Unified Host Utilities 7.1 Installation Guide](#) describes how to use the Linux Unified Host Utilities 7.0.
- Online help describes how to use SANtricity System Manager to complete configuration and storage management tasks. It is available within the product and as a PDF download.
- [NetApp Knowledgebase](#) (a database of articles) provides troubleshooting information, FAQs, and instructions for a wide range of NetApp products and technologies.
- For additional documentation and instructions for E-Series products, including SANtricity software, go to the .

Copyright and trademark

Copyright

Copyright © 2020 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.

Data contained herein pertains to a commercial item (as defined in FAR 2.101) and is proprietary to NetApp, Inc. The U.S. Government has a non-exclusive, non-transferrable, non-sublicensable, 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).

Trademark

NETAPP, the NETAPP logo, and the marks listed on the NetApp Trademarks page are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

<http://www.netapp.com/us/legal/netapptmlist.aspx>