



SANtricity® Storage Manager 11.30

Installing and Configuring for Linux

Express Guide

August 2016 | 215-11217_A0
doccomments@netapp.com

Contents

Deciding whether to use this Express Guide	5
Understanding the workflow	8
Verifying the configuration is supported	9
Installing SANtricity Storage Manager	11
Configuring IP addresses using DHCP	12
Configuring the multipath software	13
Performing FC-specific tasks	14
Configuring the switches—FC	14
Determining host WWPNs and making the recommended settings—FC	14
Performing iSCSI-specific tasks	16
Configuring the switches—iSCSI	16
Configuring networking—iSCSI	16
Configuring array-side networking—iSCSI	17
Configuring host-side networking—iSCSI	19
Verifying IP network connections—iSCSI	21
Performing iSER-specific tasks	23
Configuring the switches—InfiniBand	23
Configuring network connections—iSER	23
Configuring networking for storage attached hosts—iSER	24
Performing SRP-specific tasks	27
Determining host port GUIDs and making the recommended settings—SRP	27
Configuring the switches—InfiniBand	27
Configuring network connections—SRP	28
Performing SAS-specific tasks	30
Installing and configuring Linux Unified Host Utilities 7.0	31
Adding a storage array to SANtricity Storage Manager	32
Accepting automatic disk pool configuration	35
Creating a volume	36
Defining a host in SANtricity Storage Manager	38
Mapping a volume to a host	41
Discovering storage on the host	43
Configuring storage on the host	44
Verifying storage access on the host	46
FC worksheet	47
iSCSI worksheet	49
iSER worksheet	50
SRP worksheet	53
SAS worksheet	54
Where to find additional information	55
Copyright information	56
Trademark information	57

**How to send comments about documentation and receive update
notifications 58**

Deciding whether to use this Express Guide

The express method for installing your storage array and SANtricity Storage Manager is appropriate for setting up a standalone Linux host to E-Series or EF-Series storage systems. 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.

The express method includes the following steps:

1. Setting up one of the following communication environments:
 - Fibre Channel (FC)
 - InfiniBand (IB) with either the iSCSI Extensions for RDMA (iSER) protocol or with the SCSI RDMA Protocol (SRP)
 - iSCSI
 - SAS
2. Creating logical volumes on the storage array and assigning a logical unit number (LUN) to each volume.
3. Making the volume LUNs 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 any optional drive shelves and the controller shelf. • You have applied power to the storage array. • You have installed all other hardware (for example, host bus adapters and switches) and made the necessary connections.
Host	<ul style="list-style-type: none"> • You have made a connection between the storage array 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.

Component	Assumptions
Storage management station	<ul style="list-style-type: none"> You are using a 1 Gb/s 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 array 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 obtained the MAC addresses for management port 1 on both storage array controllers.
Storage provisioning	<ul style="list-style-type: none"> You will not use shared volumes. You will create disk 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 versions listed on the NetApp Interoperability Matrix Tool.
Protocol: IB	<ul style="list-style-type: none"> You have determined whether you have either a 56-Gbps HIC using the iSER or SRP protocol or a 40-Gbps HIC using the SRP protocol. Note: iSER is the default IB protocol for E-Series storage arrays. To use the SRP protocol, you need to apply a feature pack. You are using NetApp-supported IB host channel adapters and switches.
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 versions listed on the NetApp Interoperability Matrix Tool.

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.

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

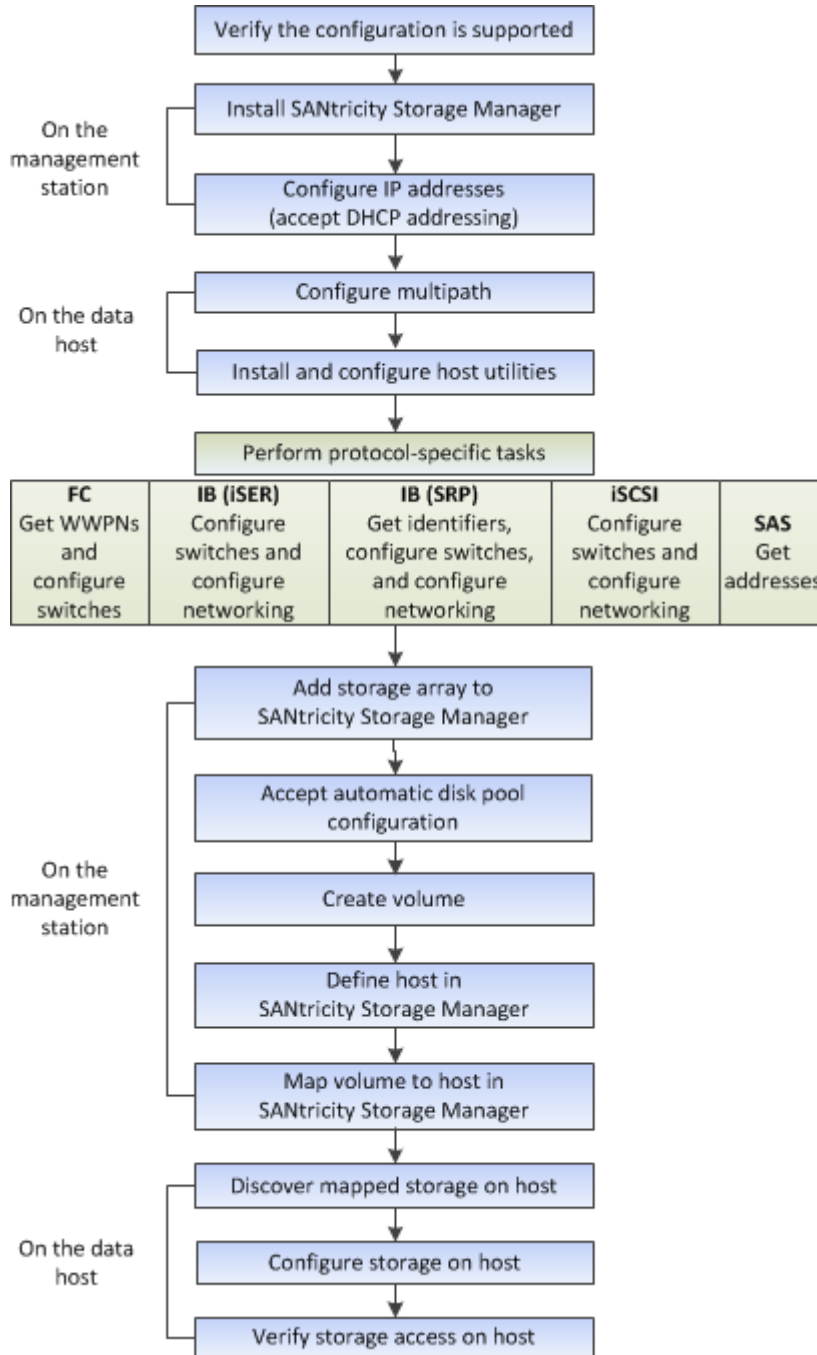
Related information

[NetApp Interoperability Matrix Tool](#)

SANtricity 11.30 Installing and Configuring for Linux Power Guide for Advanced Users

Understanding the workflow

This workflow guides you through the "express method" for configuring your storage array and SANtricity Storage Manager to make storage available to a host.

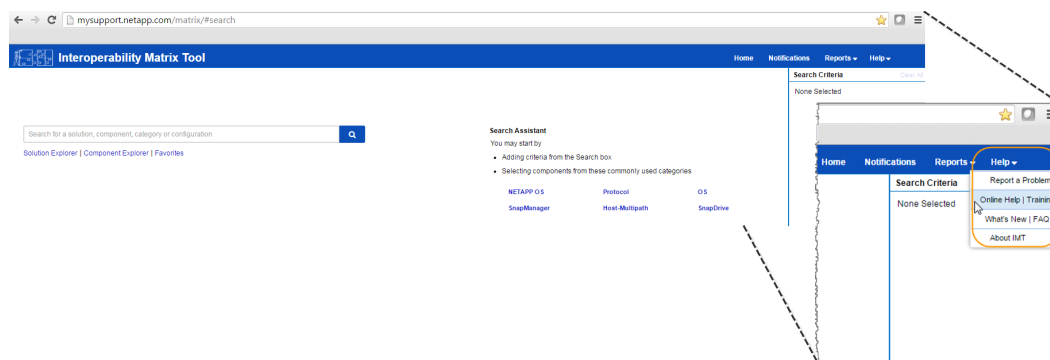


Verifying the configuration is supported

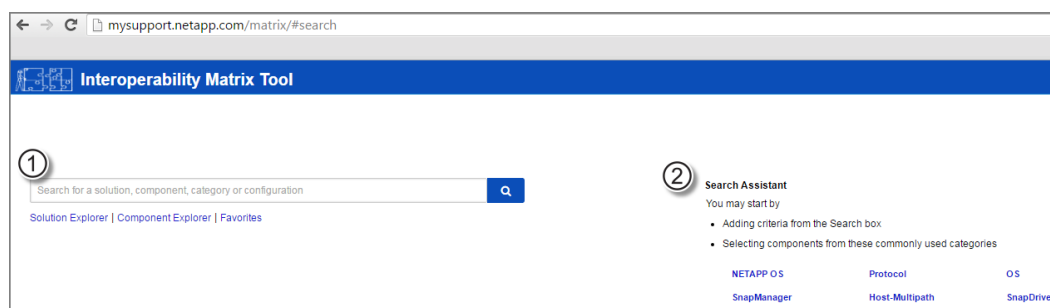
To ensure reliable operation, you create an implementation plan and then use the NetApp Interoperability Matrix Tool (IMT) to verify that the entire configuration is supported.

Steps

1. Go to the [NetApp Interoperability Matrix Tool](#).
2. Go to **Help > Online Help | Training** or **Help > What's New | FAQ** for training or refresher tools.




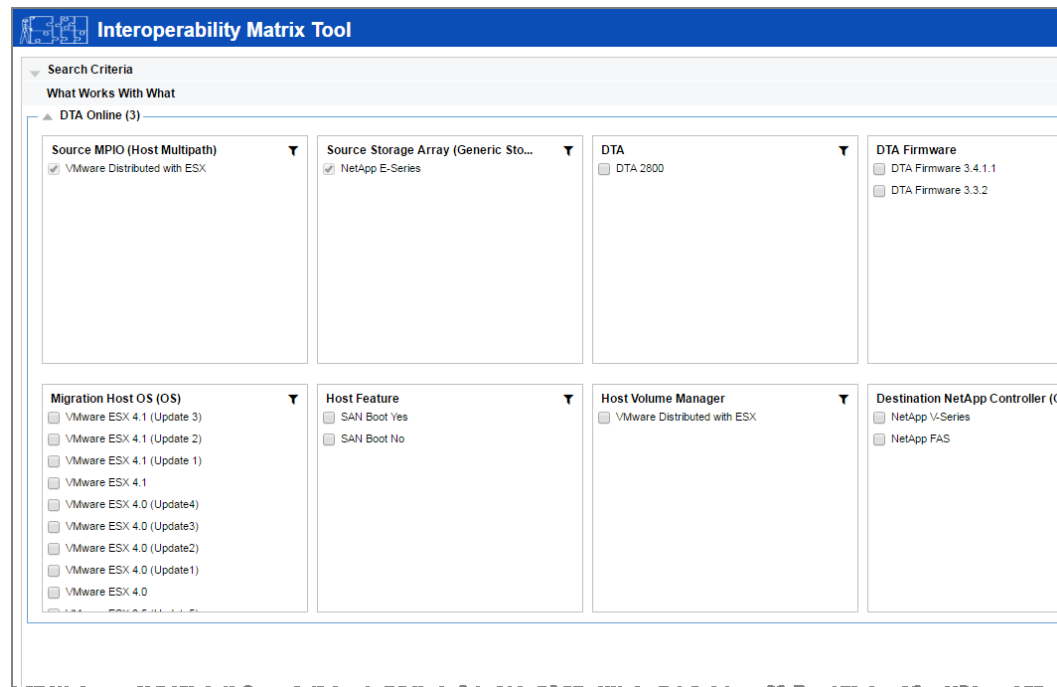
3. Use the search functions to enter the details of your configuration.



- | | |
|---|-----------------------------------------------------------------------------------------------------------------------------------|
| 1 | Search box: Enter a solution, component, category, or configuration for building initial criteria and inferring solutions. |
| 2 | Search Assistant: Use appropriate hints to infer solutions faster. |

4. Click **View What Works With What** to select from a detailed matrix of components.



Example

5. Review the information in the following tabs in the **Configuration Details** window:
 - **Notes:** Lists important information specific to your configuration. Review the alerts to identify the hot fixes that are required for your operating system.
 - **Policies & Guidelines:** Provides general guidelines for all SAN configurations.
6. As necessary, make the updates for your operating system and protocol as listed in the table.

Operating system updates	Protocol	Protocol-related updates
<i>not applicable</i>	FC	HBA driver, firmware, and bootcode
	InfiniBand	Host channel adapter driver, firmware, and bootcode
	iSCSI	Network interface card (NIC) driver, firmware and bootcode
	SAS	Host bus adapter (HBA) driver, firmware, and bootcode

Related information

[NetApp Interoperability Matrix Tool](#)

Installing SANtricity Storage Manager

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

- 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

You will install the SANtricity Storage Manager software on the management station. 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.

For more information about installation methods and customizations, see [SANtricity 11.30 Installing and Configuring for Linux Power Guide for Advanced Users](#).

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-LINUX64-11.25.0A00.0002.bin 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.

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).

Steps

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: 192.168.128.101
 - Controller B, port 1: 192.168.128.102
 - Subnet mask: 255.255.255.0
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 when you add the storage array to SANtricity Storage Manager.

Configuring 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.

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.

Example

```
# systemctl start multipathd
```

4. Configure multipath for startup persistence.

Example

```
# chkconfig multipathd on
```

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

Example

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

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

For additional configuration details, see [SANtricity 11.30 Installing and Configuring for Linux Power Guide for Advanced Users](#).

Performing FC-specific tasks

For the Fibre Channel protocol, you configure the switches and determine the host port identifiers.

Configuring the switches—FC

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.

Note: It is helpful to record the WWPNs on the [FC worksheet](#) on page 47.

About this task

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

You must zone by WWPN, not by physical port. 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.

Determining host WWPNs and making the recommended settings—FC

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
- Record the WWPNs on the [FC worksheet](#) on page 47.

- 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. Record the host identifiers on the [FC worksheet](#) on page 47.
4. Select the appropriate settings in the HBA utility.

Appropriate settings for your configuration are listed in the Notes column of the IMT.

Related concepts

[FC worksheet](#) on page 47

Related information

[NetApp Interoperability Matrix Tool](#)

Performing iSCSI-specific tasks

For the iSCSI protocol, you configure the switches and configure networking on the array side and the host side. Then you verify the IP network connections.

Configuring the switches—iSCSI

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 have enabled send and receive hardware flow control **end to end**.
- You have disabled priority flow control.
- 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.

Related concepts

[iSCSI worksheet](#) on page 49

Configuring networking—iSCSI

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.

Related tasks

[Configuring array-side networking—iSCSI](#) on page 17

Configuring array-side networking—iSCSI

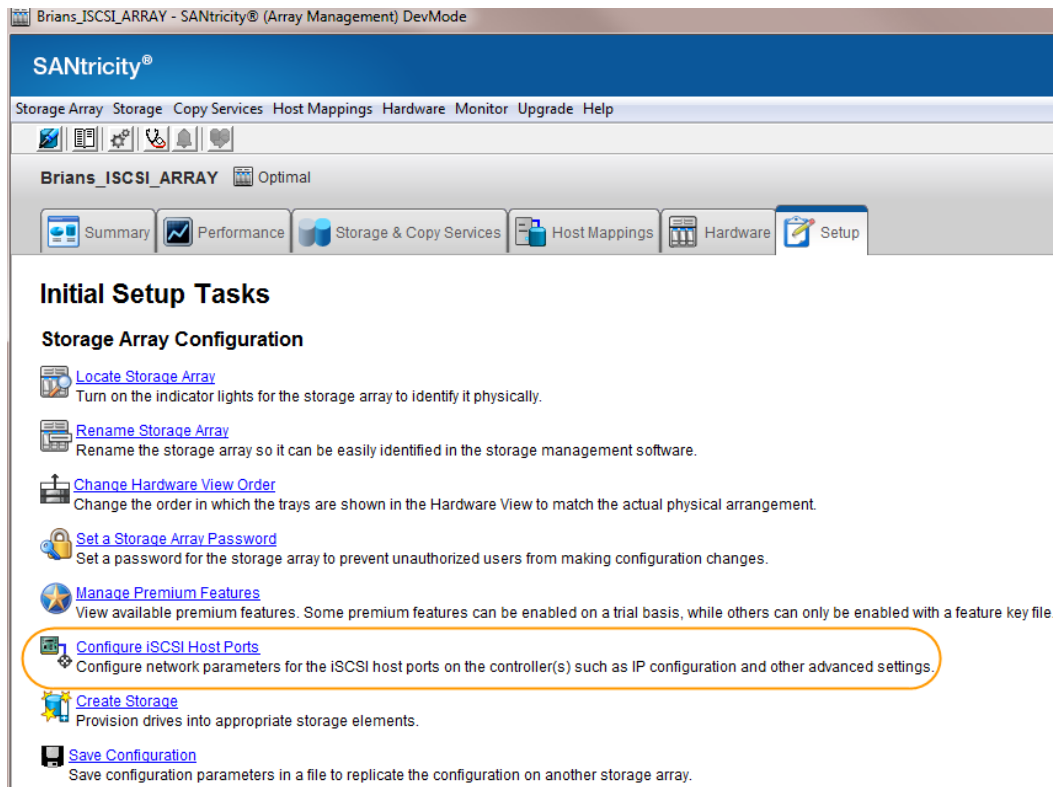
You use the Array Management Window (AMW) of SANtricity Storage Manager to configure iSCSI networking on the array side.

About this task

To configure E-Series iSCSI ports for both controllers, complete the following steps:

Steps

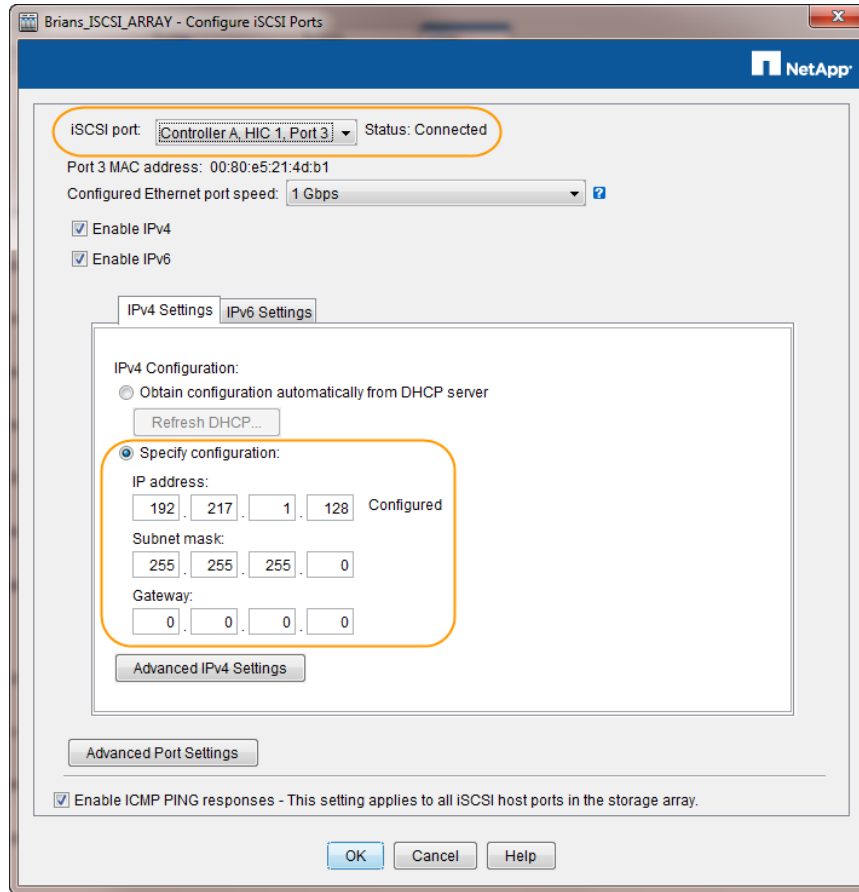
1. Open the **SANtricity Storage Manager Enterprise Management Window (EMW)**.
2. Click the **Devices** tab, and then double-click the name of the storage array in the left pane.
The **Array Management Window (AMW)** opens.
3. From the **AMW**, select the **Setup** tab.
4. Select **Configure iSCSI Host Ports**.




82011-10

5. Select the iSCSI host port to be configured and configure the IP address information as needed. Add the IP address to the worksheet for use when you configure iSCSI networking on the host side.

Note: The **Enable ICMP PING Responses** check box is checked by default and should remain selected.



6. Reconfigure the E-Series iSCSI host ports to use the same maximum transmission unit (MTU) size as the host and switch.
 - a. From the iSCSI Host Ports dialog box, select **Advanced Host Port Settings**.
 - b. Select the **Enable jumbo frames** checkbox, if appropriate.
 - c. Change the MTU size to the desired number of bytes per frame, and then select **OK**.
 - d. Select **OK** on the remaining windows as necessary to complete this procedure.



7. Repeat steps above to configure additional iSCSI ports on the E-Series controllers.

Configuring host-side networking—iSCSI

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.

Example

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

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

Red Hat Enterprise Linux 7 (RHEL 7)

Example

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

SUSE Linux Enterprise Server 12 (SLES 12)**Example**

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

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

Example

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

4. 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.

Example

```
# 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.

Example

```
# systemctl restart network
```

- d. Make sure the Linux server can ping *all* of the iSCSI target ports.
5. Configure the iSCSI interfaces by creating two iSCSI iface bindings.

Example

```
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`.

6. 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.

Example

```
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.

Example

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

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

Example

```
# iscsiadm -m session
```

Verifying IP network connections—iSCSI

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.

Performing iSER-specific tasks

For the InfiniBand iSER protocol, you configure the switches and configure the storage array and the hosts with iSER networking.

Configuring the switches—InfiniBand

Follow these guidelines and the switch vendor's documentation to configure the switches.

Before you begin

You must have administrator credentials for the InfiniBand (IB) switches.

About this task

IB fabrics require that a subnet manager is running somewhere on the fabric. If switches are available, enable the subnet manager using the management interface for the switches. If switches are not available, run a subnet manager on a host in the fabric.

For details about enabling the subnet manager, see the switch vendor's documentation.

Steps

1. Log in to the IB switch administration program.
2. Enable the subnet manager and save the configuration.

Configuring network connections—iSER

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

About this task

When you are using a 56-Gbps HIC with the iSER 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](#) on page 50.
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](#) on page 50.

Configuring networking for storage attached hosts—iSER

About this task

The InfiniBand OFED driver stack supports running both iSER and SRP simultaneously on the same ports, so no additional hardware is required. Before starting this step, see the [NetApp Interoperability Matrix Tool](#) to ensure that a NetApp-recommended OFED is installed on the system.

Steps

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

Red Hat Enterprise Linux 7 (RHEL 7)

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

SUSE Linux Enterprise Server 12 (SLES 12)

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

2. Configure IPoIB network interfaces:
 - a. Identify the InfiniBand ports that will be used. Document the HW Address (MAC address) of each port. Document the addresses in the worksheet: [iSER worksheet](#) on page 50.
 - b. (Optional; recommended) Configure persistent names for the InfiniBand network interface devices.
 - c. Configure the IP address and network information for the IPoIB interfaces identified. The specific interface configuration required may vary depending on the operating system used. Consult your vendor's operating system documentation for specific information on implementation.

Example

For RHEL 6.x Interface Name ib0, edit the `/etc/sysconfig/network-scripts/ifcfg-ib0` configuration file.

```
ONBOOT=yes
Name=ib0
Type=InfiniBand
IPADDR=192.168.xx.yy
PREFIX=24
DEFROUTE=no
HWADDR=<Port 0 HW Address Here>
```

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

Example

```
#service network restart
```

- e. Verify connectivity to the target ports. From the host, ping the IP addresses you configured in [Configuring network connections](#) on page 23.
3. Restart services to load the iSER module. 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.
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.

Example

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/ifaces/iface-ib0
iscsiadm -m iface -I iser > /var/lib/iscsi/ifaces/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
<code>iface.net_ifacename</code>	The interface device name (ex. ib0).
<code>iface.initiatorname</code>	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.

- Check the `iscsi` session status from the host:

```
iscsiadm -m session
```

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

Related information

[NetApp Interoperability Matrix Tool](#)

Performing SRP-specific tasks

For the InfiniBand SRP protocol, you configure the switches, configure networking, determine host port addresses, and make the settings recommended in the Notes column of the [NetApp Interoperability Matrix Tool](#).

Determining host port GUIDs and making the recommended settings—SRP

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](#) on page 53.
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](#).

Related information

[NetApp Interoperability Matrix Tool](#)

Configuring the switches—InfiniBand

Follow these guidelines and the switch vendor's documentation to configure the switches.

Before you begin

You must have administrator credentials for the InfiniBand (IB) switches.

About this task

IB fabrics require that a subnet manager is running somewhere on the fabric. If switches are available, enable the subnet manager using the management interface for the switches. If switches are not available, run a subnet manager on a host in the fabric.

For details about enabling the subnet manager, see the switch vendor's documentation.

Steps

1. Log in to the IB switch administration program.
2. Enable the subnet manager and save the configuration.

Configuring network connections—SRP

If your configuration uses the SRP 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 may 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.
2. Configure OFED/RDMA to load the SRP module.
3. In the OFED/RDMA configuration file, set `SRP_LOAD=yes` and `SRP_DAEMON_ENABLE=yes`.
4. Enable and start the OFED/RDMA service.

For RHEL

- To enable the InfiniBand modules to load on boot:

```
chkconfig rdma on
```

- To load the InfiniBand modules immediately:

```
service rdma start
```

For SUSE Linux

- To enable the InfiniBand modules to load on boot:

```
chkconfig openibd on
```

- To load the InfiniBand modules immediately:

```
service openibd start
```

5. Enable the SRP daemon.

For RHEL

- To enable the SRP daemon to start on boot:

```
chkconfig srpd on
```

- To start the SRP daemon immediately:

```
service srpd start
```

For SUSE Linux

- To enable the SRP daemon to start on boot:

```
chkconfig opensmd on
```

- To start the SRP daemon immediately:

```
chkconfig opensmd start
```

Related information

[*NetApp Interoperability Matrix Tool*](#)

Performing SAS-specific tasks

For the SAS protocol, you determine host port addresses and make the settings recommended in the Notes column of the [NetApp Interoperability Matrix Tool \(IMT\)](#).

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.
- It is helpful to record the SAS addresses on the [SAS worksheet](#) on page 54.
- Host I/O ports might automatically register if the host context agent is installed.

Steps

1. Download the LSI-sas2flash(6G) or sas3flash(12G) utility from your HBA vendor's web site.
2. Install the utility.
3. Record the host identifiers (SAS addresses) on the [SAS worksheet](#) on page 54.
4. Use the HBA BIOS to select the appropriate settings for your configuration.

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

Related concepts

[SAS worksheet](#) on page 54

Related information

[NetApp Interoperability Matrix Tool](#)

Installing and configuring Linux Unified Host Utilities 7.0

Linux Unified Host Utilities 7.0 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 7.0 to install.

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

2. Download the Unified Host Utilities 7.0 from [NetApp Support](#).

Related information

[Linux Unified Host Utilities 7.0 Installation Guide](#)

Adding a storage array to SANtricity Storage Manager

After you have configured the network between the management station and the array controllers, you add the controller to SANtricity Storage Manager using the Enterprise Management Window (EMW).

Steps

1. Open SANtricity Storage Manager.

Use a Linux terminal window, and run the `SMclient` command from any directory. The Linux installation also places a shortcut to launch SANtricity Storage Manager on the desktop. After using the `SMclient` command or the shortcut to open the SANtricity Storage Manager, the (EMW) is displayed.

2. On the **Select Addition Method** screen, select the **Manual** radio button, and then select **OK**.

Note: When you open SANtricity Storage Manager for the first time, the **Select Addition Method** screen prompts you to select the **Automatic** or **Manual** method to add a new storage array.

3. To add one or more new storage arrays, complete the following steps:
 - a. On the **Add New Storage Array – Manual** screen, make sure that the default **Add using Ethernet connection (out-of-band)** radio button is selected. Enter the IP address (as determined by DHCP) of port 1 for each controller (controller A and B) in the fields provided.

The screenshot shows the 'Add New Storage Array - Manual' window. It features the NetApp logo in the top right. Below the title bar, there are three links: 'What are in-band and out-of-band management connections?', 'Adding controllers with more than one Ethernet port', and 'What if my system only has one controller?'. The 'Select a management method:' section has two radio buttons. The first, 'Add using Ethernet connection (out-of-band):', is selected and circled in orange. Below it are two empty text input fields for controller addresses. The second radio button, 'Add using Host I/O connection (in-band):', is unselected. Below it is a text input field for the host address. A note below the in-band section states: 'Note: Not applicable for storage arrays managed by System Manager.' At the bottom, there are three buttons: 'Add', 'Cancel', and 'Help'.

- b. Select **Add**.

The Storage Array Added screen is displayed.

- c. To add another storage array, select **Yes** on the **Storage Array Added** screen and repeat the steps above.

Otherwise, select **No**.

4. Name the storage array so you can find it more easily in the EMW after you change the IP addresses:
 - a. In the **Setup** tab, select **Name/Rename Storage Arrays**.
 - b. In the **Select storage array:** list, confirm that the storage array you added is selected.
 - c. In the **Storage array name** field, type a name for the storage array.
 - d. Select **OK**.
 - e. On the **Name/Rename Storage Arrays** warning dialog, select **Yes** to dismiss the warning and continue.
5. Configure the network configuration information of the controllers, using information you obtain from your network administrator.
 - a. Click the **Devices** tab, and then double-click on the name of the storage array.
The Array Management Window (AMW) opens.
 - b. Select the **Hardware** tab in the Array Management Window (AMW).
 - c. Select the visual representation of one of the controllers in the **Hardware** tab.
 - d. Right-click, and then select **Configure > Management Ports**.
 - e. On the **Change Network Configuration** dialog box, select **Controller A, Port 1** in the **Ethernet port** drop-down list.
 - f. From the **Speed and duplex mode** drop-down list, select **Auto-negotiate**.
 - g. Depending on the format of your network configuration information, select the **Enable IPv4** check box, the **Enable IPv6** check box, or both check boxes.
 - h. Depending on the format you have selected, enter the network configuration information in the **IPv4 Settings** tab or the **IPv6 Settings** tab.

Note: You must obtain the network configuration information from your network administrator.
 - i. Select **Controller B, Port 1** in the **Ethernet Port** drop-down list, and repeat the steps you followed for controller A.
 - j. Select **OK**.
6. Find the identifiers of the controller's host ports by opening the AMW for the storage array and then following the appropriate steps for your protocol.

Fibre Channel

- a. From the **Monitor** menu, select **Reports > Storage Array Profile**.
- b. Select the **Hardware** tab, then select the **Controllers** tab.
- c. Scroll through the information to find the WWPNs.
- d. Record the WWPNs in the worksheet and close the dialog box.

iSCSI or InfiniBand

- a. Click the **Setup** tab.

- b. Click the link for **Manage iSCSI Settings**. Scroll through the information to find the identifiers.
- c. Record the identifiers in the worksheet and close the dialog box.

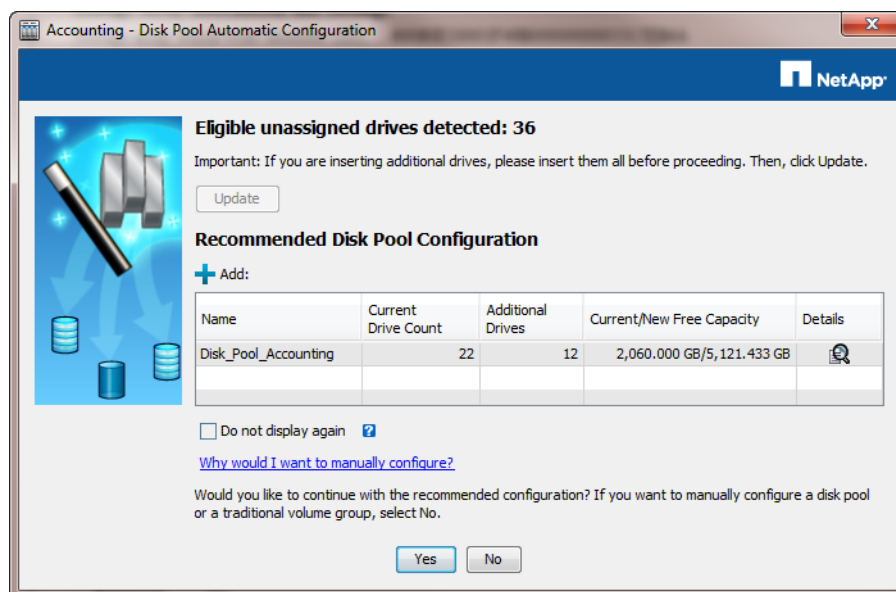
Accepting automatic disk pool configuration

When there are unassigned drives available to create new disk pools, the Disk Pool Automatic Configuration wizard is displayed when you open the Array Management Window (AMW) for your storage array. You are prompted to select the recommended disk pool configuration.

Steps

1. In the **Enterprise Management Window (EMW)**, click **Devices**, and then double-click the storage array to open the Array Management Window (AMW).

Note: When you open the AMW for a storage array for the first time, the Disk Pool Automatic Configuration wizard is displayed.
2. Accept the recommended disk pool configuration by clicking **Yes** in the **Disk Pool Automatic Configuration** dialog box.



Creating a volume

Using SANtricity Storage Manager, you create a volume—a single accessible storage area—on a storage array for the attached host to access. You create a volume from the free capacity of a disk pool.


Before you begin

You must have determined the expected capacity, usage, data protection, and performance requirements for the volume.

Steps

1. From the **Array Management Window (AMW) Storage & Copy Services** tab, expand the disk pool on the array where you want to create the new volume.
2. Right-click **Free Capacity > Create Volume**.

The following dialog box appears.



3. Configure the volume parameters.

Onscreen flyovers provide more information about particular options.

- a. If you want to create a thin volume, select the **Create thin volume** check box.

- b. From the **Units** drop-down list, select the appropriate unit for the new volume capacity.
 - c. Specify the volume capacity to be taken from the free capacity that is listed.
 - d. Enter the volume name.
 - e. From the **Map to host** drop-down list, select the **Map later** option.
4. Specify the Quality of Service attributes. Use the onscreen flyovers and the SANtricity Online Help to get more information about particular attributes.

Note: The **Use SSD cache** check box is not available in the Quality of Service section because you have chosen to map the volume to a host later. See the Array Management Window Online Help topic “Learn about SSD Cache” to decide if it is appropriate to enable SSD cache later. To enable it, from the **AMW Storage & Copy Services** tab, right-click the volume, and then select **SSD Cache > Enable**.

 - a. To enable DA protection, select the **Enable data assurance (DA) protection on the new volume** check box.

This check box appears only if the drives, the controller, and the host bus adapter are all DA-capable.
 - b. Finish selecting the Quality of Service Attributes and create the volume.
5. If you want to enable dynamic cache read prefetch, select the **Enable dynamic cache read prefetch** check box to enable it.


Note: Dynamic cache read prefetch is not available for thin volumes.
6. Select **Finish**.

Defining a host in SANtricity Storage Manager


You define a new logical host on the storage array so that volumes can be shared with the host.

Steps

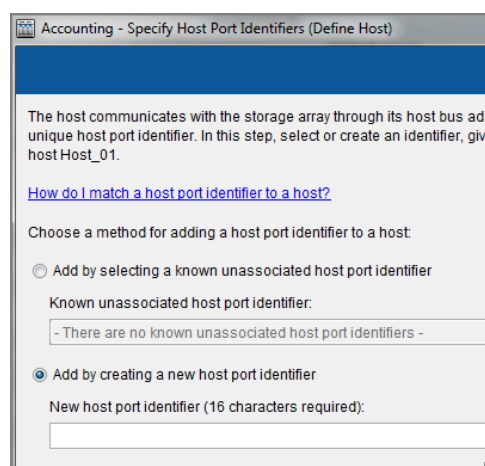
1. From the **Array Management Window (AMW)**, select the **Host Mappings** tab.
2. In the left pane, expand the storage array tree.
3. Right-click **Default Group** and select **Define > Host** to start the **Define Host** wizard.
 - a. Enter a descriptive name for the host to make it easier for administrators to manage the environment over time.
 - b. In the **Question** area of the dialog box, keep the **Yes** selected.



4. Select **Next**.
5. If the controller supports multiple host interface types, select one type from the **Choose a host interface type** drop-down list.
6. Choose a method for adding a host port identifier to the host.
 - You can add known unassociated host port identifiers by selecting the option **Add by selecting a known unassociated host port identifier**. From the drop-down list, select the identifier for the first port on the host.
In this example, a Fibre Channel host interface type is shown.



- If no identifiers are displayed, there is an issue with the path to the host, and the storage cannot be discovered. Resolve the host issue, then change the host port identifier that was not discovered by selecting the option **Add by creating a new host port identifier**. Enter the new host port identifier.



Note: The host port identifier is called a different name depending on the protocol:

- Fibre Channel and SAS: World Wide Identifier (WWID)
- InfiniBand and iSCSI: iSCSI Qualified Name (IQN)

When the identifier is displayed in the selection list, the storage array can automatically detect the path to the host.

7. Enter a descriptive alias name for the host port identifier.
8. Select **Add** to add the host port identifier to the host.
9. Repeat Step 6 on page 38 through Step 8 on page 39 for each link between the host and the storage array. You can connect and provision two to four paths between any one host and the E-Series storage array.

Note: For InfiniBand and iSCSI, additional links between host and storage will not show up as different identifiers; they will all exist behind the same IQN.

10. Select **Next**.

11. From the **Host type (operating system)** drop-down list, select host type **Linux (DM-MP) for DMMP Kernel 3.9 or earlier (RHEL6 and earlier, and SLES11); or InxDHALUA for DMMP Kernel 3.10 or later (RHEL7 and SLES11)**.
12. Select **Next**.
13. Select **No – this host will NOT share access to the same volumes with other hosts**, and then select **Next**.

Note: These instructions assume a non-clustered configuration.
14. Review the host definition information and confirm that the host name, the host type, and the host port identifiers are correct.
15. Select **Finish** to create the host.
16. Repeat Step 3 on page 38 through Step 15 on page 40 to create additional hosts as required.
17. From the **Host Mappings** tab, review the storage array tree to verify that the hosts were created.

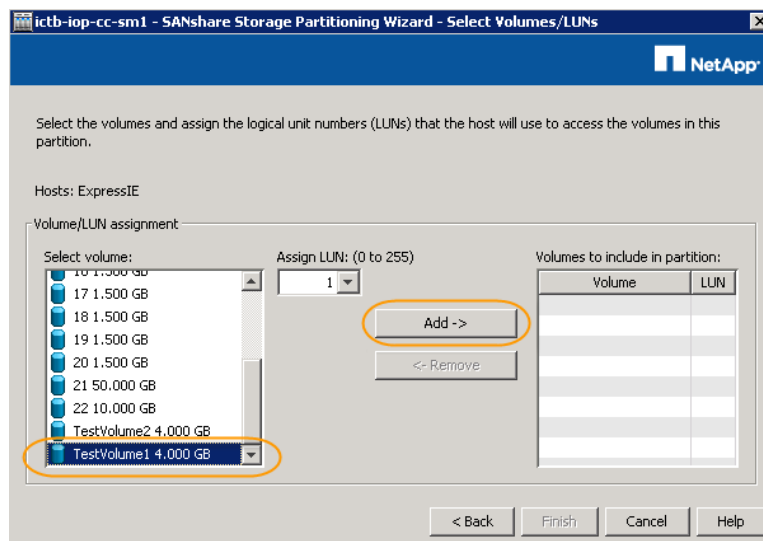
Mapping a volume to a host

Using SANtricity Storage Manager to create storage partitions, you assign a logical unit number (LUN) to a volume and map the LUN to the host.

Steps

1. From the **Array Management Window (AMW)**, select the **Host Mappings** tab.
2. Select the storage array tree, right-click the desired host, and then select **Define Storage Partition** to start the **SANshare Storage Partitioning** wizard.
3. On the **Welcome** page of the wizard, select **Next**.
4. Select **Host** to create a dedicated mapping, and then select the host name to be added to the partition.
5. Select **Next**.
6. Select an individual volume, assign a LUN to the volume, and then select **Add** to associate the volume (LUN) with the partition.

Note: When assigning LUNs for volumes attached to Linux hosts, start with LUN 1 rather than LUN 0. If an array is connected to the host before a device is mapped to LUN 0, the Linux host detects the REPORT LUNS well known logical unit as LUN 0, so that it can complete discovery. LUN 0 might not immediately map properly with a simple rescan, depending on the version of the host operating system.

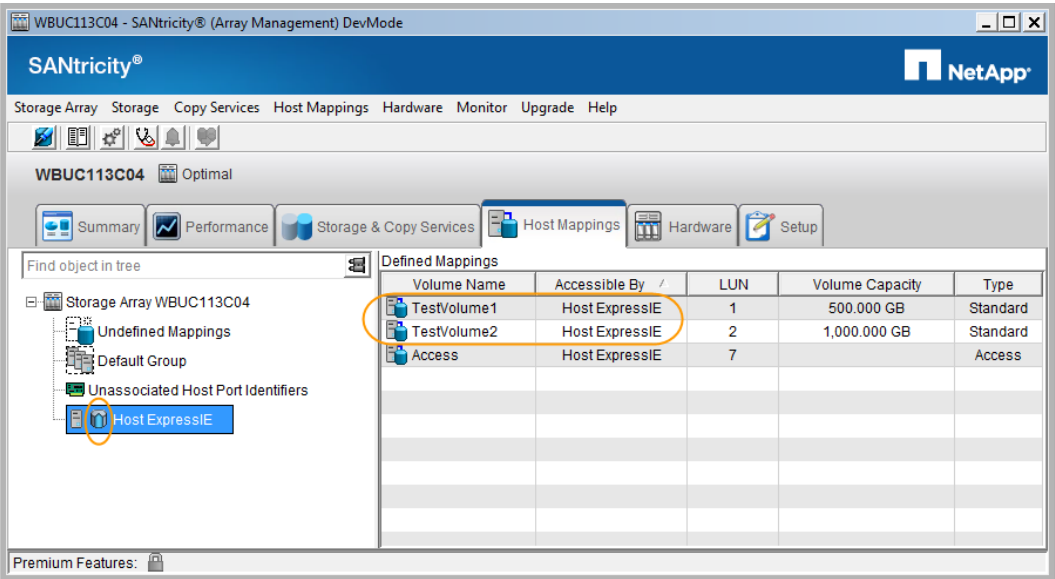


7. Repeat Step 6 until all desired volumes are added to the partition.
8. Select **Finish** to create the partition.
9. Review the storage array tree on the **Host Mappings** tab to confirm that the partition was successfully created.

Example

Note: In this example, the volumes TestVolume1 and TestVolume2 are mapped to the host ExpressIE. Only this host has access to TestVolume1 and TestVolume2. In the left pane of

the screen, the storage partition on the host is indicated by the blue slice on the storage cylinder icon.



Discovering storage on the host

LUNs on your storage system appear as disks to the Linux host. When you add new LUNs, you must manually rescan the associated disks to discover them. The host does not automatically discover new LUNs.

Steps

1. Scan the LUNs by running the `# rescan-scsi-bus.sh` command from a terminal window on the host.

This command works only if `sg3_utils` is installed. Alternatively, you can reboot the host or manually scan each SCSI host.

```
# echo "-- --"
> /sys/class/scsi_host/host3/scan
```

2. Verify disk discovery by running the `# sanlun lun show -p` command.

Example

```
# sanlun lun show
controller(7mode/E-Series)
<column headings deleted for readability>
-----
ictm1619s01c01-SRP    /dev/sddi    host14    SRP    E-Series
ictm1619s01c01-SRP    /dev/sddh    host14    SRP    E-Series
ictm1619s01c01-SRP    /dev/sddf    host14    SRP    E-Series
```

Configuring storage on the host

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.

Example

```
# 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).

Example

```
# 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.

Example

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

4. Create a folder to mount the new partition.

Example

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

5. Mount the partition.

Example

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

Verifying storage access on the host

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

Before you begin

You must have initialized the LUN 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.
4. Run the `multipath -ll` command to view the paths to the LUN, and verify that you have four paths.
5. Force a controller failover, preferably by pulling all cables from one controller, and then verify that you can still access the files on the LUN. When you are finished, reset the storage to an optimal state.


After you finish

Remove the file and folder that you copied.

FC worksheet

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	

iSCSI worksheet

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	

iSER worksheet

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

iSER: Host identifiers


Note: The software initiator IQN is determined during the task, [Configuring storage attached hosts with iSER networking](#) on page 24.

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		
n/a		
n/a		
n/a		

iSER: Recommended configuration

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



iSER: Target IQN

Document the target IQN for the storage array. You will use this information in [Configuring storage attached hosts with iSER networking](#) on page 24.

Find the Storage Array IQN name using SANtricity: **Storage Array > iSER > Manage Settings**. This information may 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: 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](#) on page 24.

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: Mapping host name

Note: The mapping host name is created during the workflow, [Defining a host in SANtricity Storage Manager](#) on page 38.

Mapping host name	
Host OS type	

SRP worksheet

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



SRP: Host identifiers

Note: The initiator GUIDs are determined in the task, [Determining host port GUIDs and making the recommended settings](#) on page 27.

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: Recommended configuration

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

SRP: Mapping host name

Note: The mapping host name is created during the workflow, [Defining a host in SANtricity Storage Manager](#) on page 38.

Mapping host name	
Host OS type	

SAS worksheet

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

Host Identifiers

Host (initiator) port connections	SAS address
Host (initiator) port connected to Controller A	
Host (initiator) port connected to Controller B	

Target Identifiers

Recommended configurations consist of two target ports.

Mappings Host

Mappings Host Name	
Host OS Type	

Where to find additional information

Use the resources listed here if you need additional information. You can also use the online help systems for the Enterprise Management Window (EMW) and the Array Management Window (AMW) of SANtricity Storage Manager.

- [*SANtricity 11.30 Installing and Configuring for Linux Power Guide for Advanced Users*](#) describes:
 - Software installation options
 - Configuration options
 - Multipath options
 - Installation on a boot device
- [*Linux Unified Host Utilities 7.0 Installation Guide*](#) describes how to use the Linux Unified Host Utilities 7.0.
- Online help describes how to use SANtricity Storage 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 [*NetApp E-Series and EF-Series Systems Documentation Center*](#).

Copyright information

Copyright © 1994–2016 NetApp, Inc. All rights reserved. Printed in the U.S.

No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

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

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

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

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

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

Trademark information

NetApp, the NetApp logo, Go Further, Faster, AltaVault, ASUP, AutoSupport, Campaign Express, Cloud ONTAP, Clustered Data ONTAP, Customer Fitness, Data ONTAP, DataMotion, Fitness, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexArray, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexVol, FPolicy, GetSuccessful, LockVault, Manage ONTAP, Mars, MetroCluster, MultiStore, NetApp Insight, OnCommand, ONTAP, ONTAPI, RAID DP, RAID-TEC, SANtricity, SecureShare, Simplicity, Simulate ONTAP, Snap Creator, SnapCenter, SnapCopy, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, StorageGRID, Tech OnTap, Unbound Cloud, and WAFL and other names are trademarks or registered trademarks of NetApp, Inc., in the United States, and/or other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. A current list of NetApp trademarks is available on the web.

<http://www.netapp.com/us/legal/netapptmlist.aspx>

How to send comments about documentation and receive update notifications

You can help us to improve the quality of our documentation by sending us your feedback. You can receive automatic notification when production-level (GA/FCS) documentation is initially released or important changes are made to existing production-level documents.

If you have suggestions for improving this document, send us your comments by email.

[*doccomments@netapp.com*](mailto:doccomments@netapp.com)

To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

If you want to be notified automatically when production-level documentation is released or important changes are made to existing production-level documents, follow Twitter account @NetAppDoc.

You can also contact us in the following ways:

- NetApp, Inc., 495 East Java Drive, Sunnyvale, CA 94089 U.S.
- Telephone: +1 (408) 822-6000
- Fax: +1 (408) 822-4501
- Support telephone: +1 (888) 463-8277