



SANtricity® System Manager 11.30

Installing and Configuring for Linux®

Express Guide

March 2017 | 215-11222_B0
doctrcomments@netapp.com

 **NetApp®**

Contents

Deciding whether to use this Express Guide	4
Understanding the workflow	6
Verifying the configuration is supported	7
Configuring IP addresses using DHCP	9
Configuring the multipath software	10
Installing and configuring Linux Unified Host Utilities 7.0	11
Performing FC-specific tasks	12
Configuring the switches—FC	12
Determining host WWPNs and making the recommended settings—FC	12
Performing iSCSI-specific tasks	14
Configuring the switches—iSCSI	14
Configuring networking—iSCSI	14
Configuring array-side networking—iSCSI	15
Configuring host-side networking—iSCSI	17
Verifying IP network connections—iSCSI	19
Performing iSER-specific tasks	21
Configuring the switches—InfiniBand	21
Configuring network connections—iSER	21
Configuring networking for storage attached hosts—iSER	22
Performing SRP-specific tasks	25
Determining host port GUIDs and making the recommended settings—SRP	25
Configuring the switches—InfiniBand	25
Configuring network connections—SRP	26
Performing SAS-specific tasks	28
Accessing SANtricity System Manager and using Setup wizard	29
Discovering storage on the host	31
Configuring storage on the host	32
Verifying storage access on the host	34
FC worksheet	35
iSCSI worksheet	37
iSER worksheet	38
SRP worksheet	41
SAS worksheet	42
Where to find additional information	43
Copyright information	44
Trademark information	45
How to send comments about documentation and receive update notifications	46

Deciding 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.

The express method includes the following steps:

1. Setting up one of the following communication environments:
 - Fibre Channel (FC)
 - 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 the optional drive shelves and the array controllers. • You have applied power to the storage array. • You have installed all other hardware (for example, management station, 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.
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.

Component	Assumptions
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 array.
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 versions as listed on 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 versions as listed on the NetApp Interoperability Matrix Tool.

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.

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.

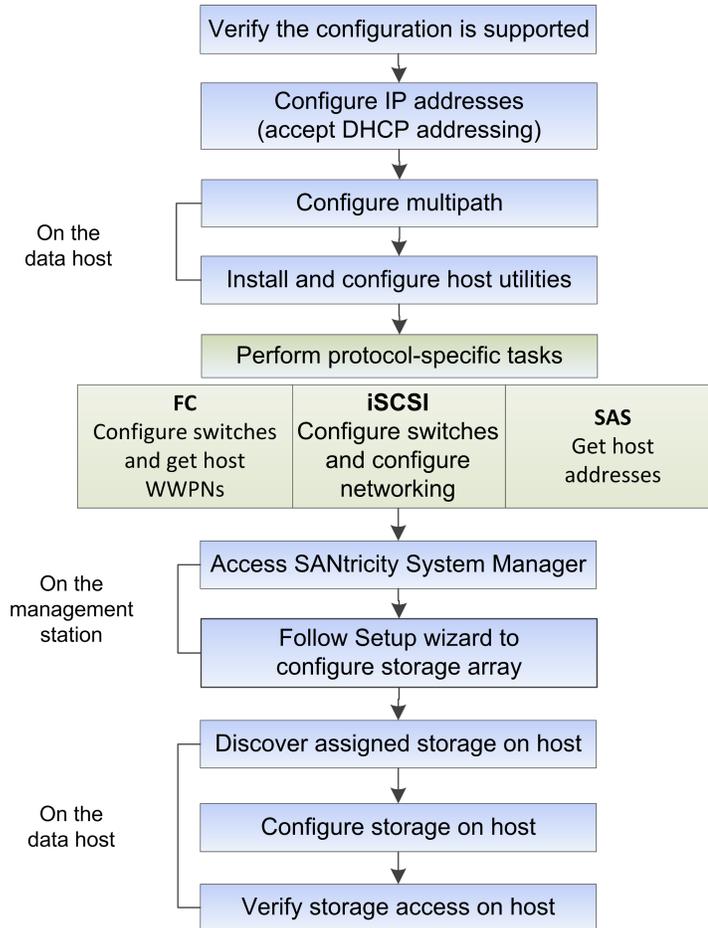
Related information

[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 System Manager to make storage available to a host.

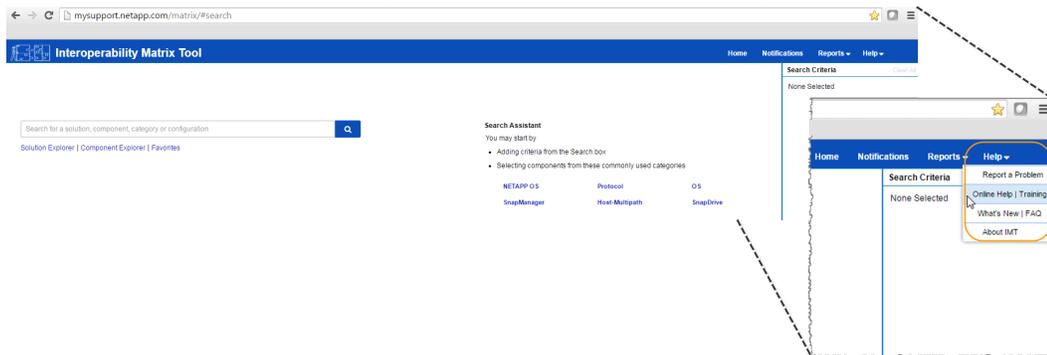


Verifying the configuration is supported

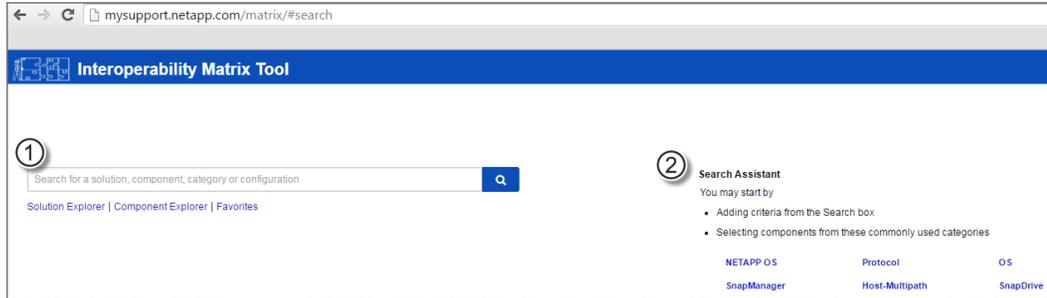
To ensure reliable operation, you create an implementation plan and then use the NetApp Interoperability Matrix Tool (IMT) to verify that the entire configuration is supported.

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.

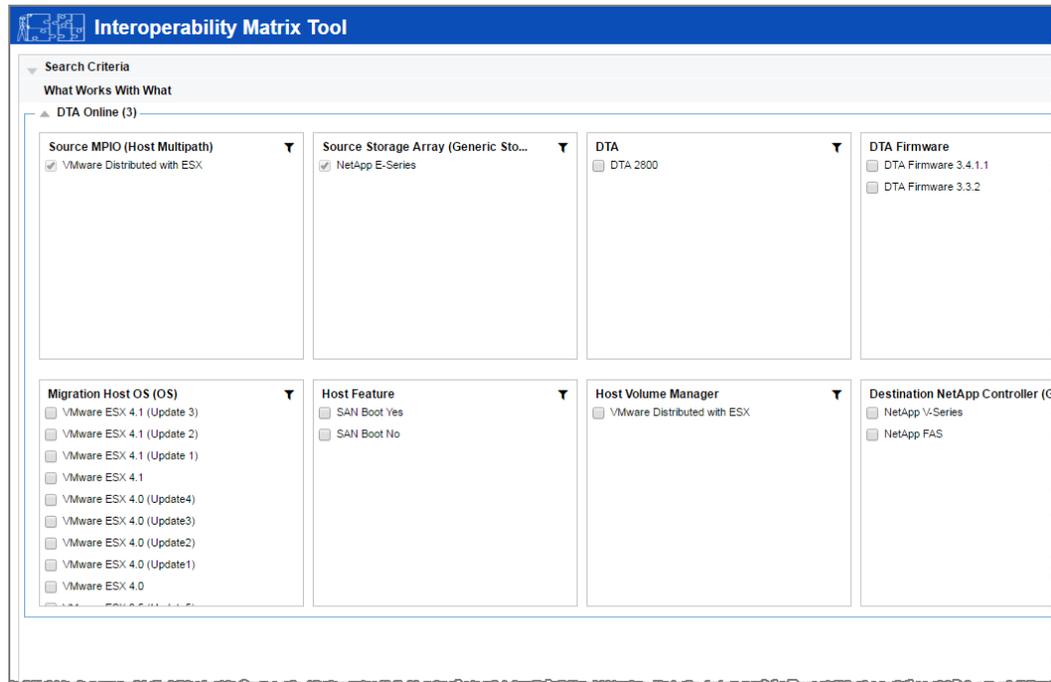


- | | |
|---|---|
| 1 | Search box: Enter a solution, component, category, or configuration for building initial criteria and inferring solutions. |
| 2 | Search Assistant: Use appropriate hints to infer solutions faster. |

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
	iSCSI	Network interface card (NIC) driver, firmware and bootcode
	SAS	Host bus adapter (HBA) driver, firmware, and bootcode

Related information

[NetApp Interoperability Matrix Tool](#)

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 assign volumes to the host.

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

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](#).

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 35.

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 35.

- 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 35.
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 35

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 37

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 15

Configuring array-side networking—iSCSI

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

Before you begin

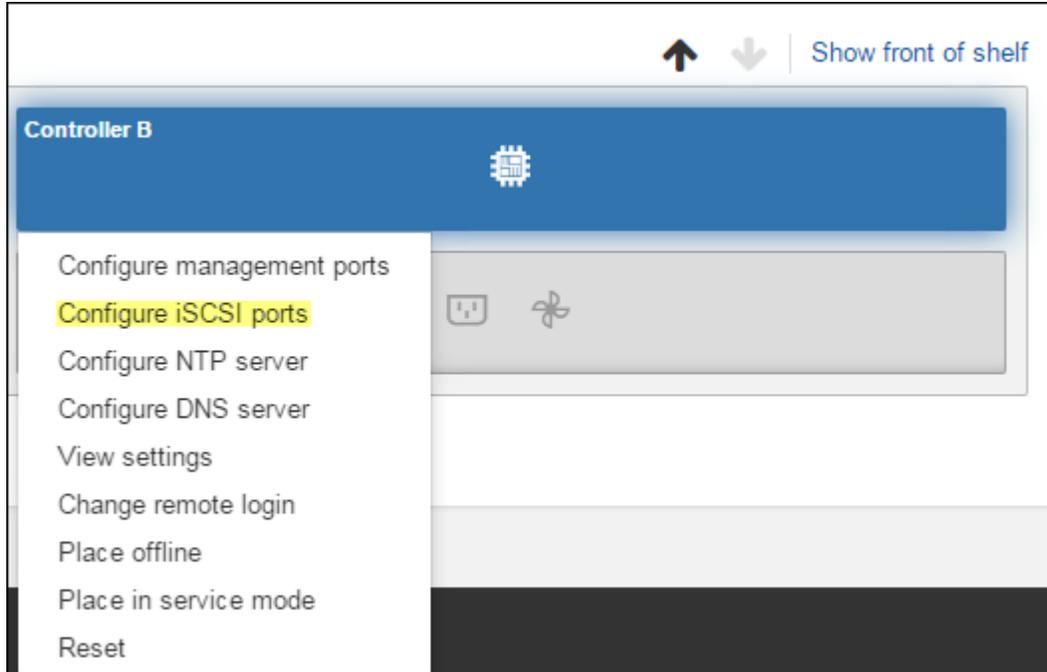
- You must know the IP address or domain name for one of the storage array controllers.
- You or your system administrator must have set up a password for the System Manager GUI.

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 DHCP-assigned address for one of the storage array controllers.
The login page for the System Manager GUI appears.
2. Enter the System Manager password.
When you open System Manager and no pools, volumes groups, workloads, or notifications have been configured, the Setup wizard launches.
3. For now, 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	Select the desired speed. The options that appear in the drop-down list depend on the maximum speed that your network can support (for example, 10 Gb/s).
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 .)	If necessary, enter a new port number. 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.
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 Maximum Transmission Unit (MTU) size is 1500 bytes per frame. You must enter a value between 1500 and 9000.
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.
Enable VLAN support (Available by clicking Show more settings .)	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 .)	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.

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

You will use System Manager to configure you storage in [Accessing SANtricity System Manager and using Setup wizard](#) on page 29. For the Setup wizard to appear again, you must close and re-open System Manager or refresh your browser.

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:2365.60080e50001bf160000000531d7be3
```

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

Example

```
iscsiadm -m node -T iqn.1992-01.com.netapp:
2365.60080e50001bf160000000531d7be3
-p 192.168.0.1:3260 -I iface0 -l
```

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

Example

```
# iscsiadm -m session
```

Related information

[NetApp Interoperability Matrix Tool](#)

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
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 38.
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 38.

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 38.
 - 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 21.
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**.

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 41.
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](#).

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
```

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 42.
- 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 42.
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 42

Related information

NetApp Interoperability Matrix Tool

Accessing SANtricity System Manager and using 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	42
Microsoft Internet Explorer	11
Microsoft Edge	2.0
Mozilla Firefox	24
Safari	6

- 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:

- 0 pools and volume groups are detected.
- 0 workloads are detected.
- 0 notifications are configured.

If the Setup wizard does not automatically appear, contact technical support.

Important: System Manager allows you to manage just your E2800 storage array. You need SANtricity Storage Manager to manage the storage enterprise and arrays other than the E2800. See the SANtricity *Storage Manager Express Guide* for your operating system for installation instructions.

Steps

- From your browser, enter the following URL:

https://<DomainNameOrIPAddress>

IPAddress is the DHCP-assigned address for one of the storage array controllers.

The login page for the System Manager GUI appears.

- Enter the System Manager password.

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.
 - **Select applications** – Specify applications, such as Exchange or SQL. System Manager optimizes storage based on application type.
 - **Define workloads** – Set up workloads, which are storage objects that support applications. You define one or more workloads per application.
 - **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 automatically 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. Create volumes by going to **Storage > Volumes > Create > Volume**.

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

Related tasks

[Configuring IP addresses using DHCP](#) on page 9

Related information

[SANtricity Storage Manager 11.30 Installing and Configuring for Linux Express Guide](#)

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.

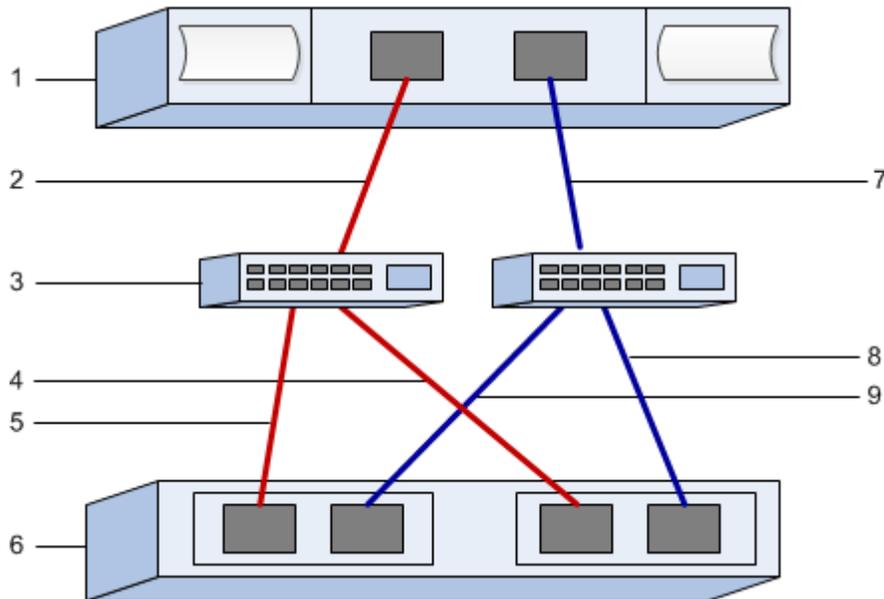
After you finish

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

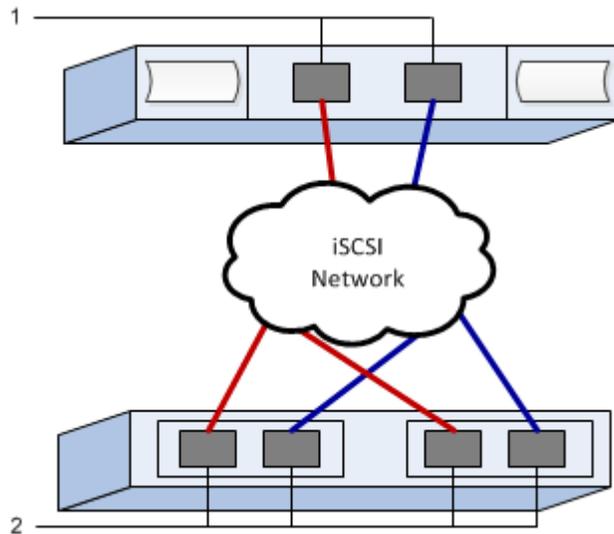
Mapping host name	
Host OS type	

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

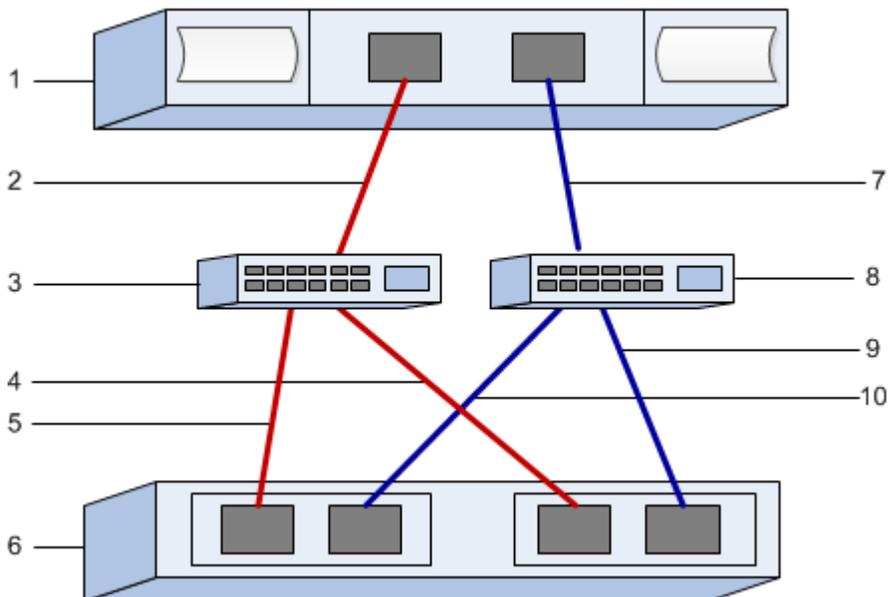
Note: The software initiator IQN is determined during the task, [Configuring storage attached hosts with iSER networking](#) on page 22.

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		

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 22.

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 22.

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	

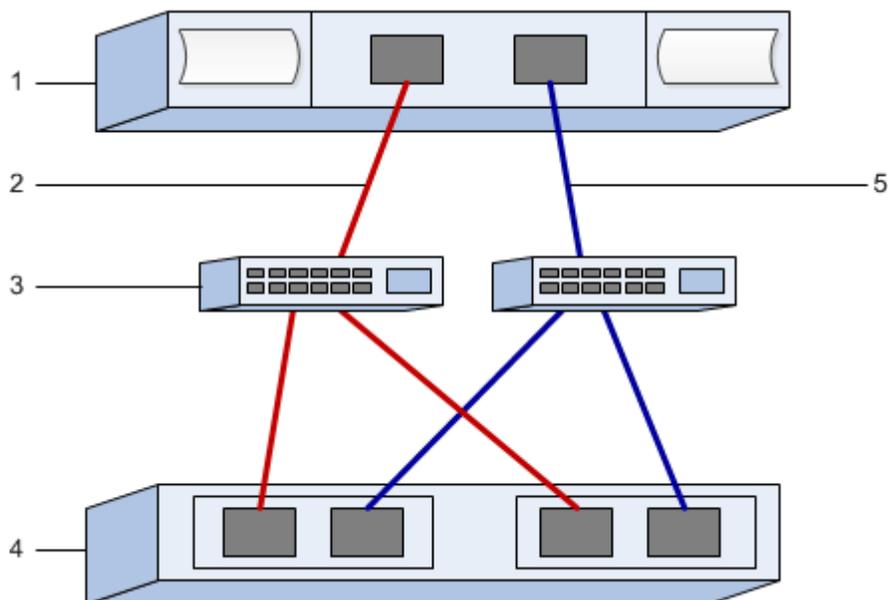
iSER: Mapping host name

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

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 25.

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.

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 for SANtricity System 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 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 [*NetApp E-Series and EF-Series Systems Documentation Center*](#).

Copyright information

Copyright © 1994–2017 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

Active IQ, AltaVault, Arch Design, ASUP, AutoSupport, Campaign Express, Clustered Data ONTAP, Customer Fitness, Data ONTAP, DataMotion, Element, Fitness, Flash Accel, Flash Cache, Flash Pool, FlexArray, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexVol, FPolicy, Fueled by SolidFire, GetSuccessful, Helix Design, LockVault, Manage ONTAP, MetroCluster, MultiStore, NetApp, NetApp Insight, OnCommand, ONTAP, ONTAPI, RAID DP, RAID-TEC, SANscreen, SANshare, SANtricity, SecureShare, Simplicity, Simulate ONTAP, Snap Creator, SnapCenter, SnapCopy, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, SolidFire, SolidFire Helix, StorageGRID, SyncMirror, 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