




NetApp®

Updated for 8.3.1

Clustered Data ONTAP® 8.3

iSCSI Configuration for Red Hat® Enterprise Linux® Express Guide



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
Web: www.netapp.com
Feedback: doccomments@netapp.com

Part number: 215-09655_B0
June 2015

Contents

Deciding whether to use this guide	4
iSCSI configuration and provisioning workflow	5
Verifying that the iSCSI configuration is supported	5
Completing the iSCSI configuration worksheet	6
Installing Linux Unified Host Utilities and optimizing I/O performance	8
Recording the iSCSI node name	9
Setting the iSCSI replacement timeout value	9
Starting the iSCSI service	9
Configuring DM-Multipath	10
Creating an aggregate	11
Deciding where to provision the volume	12
Verifying that the iSCSI service is running on an existing SVM	12
Creating a LUN and its containing volume	13
Configuring iSCSI on an existing SVM	14
Creating a new SVM	16
Starting the iSCSI sessions with the target	18
Discovering new SCSI devices (LUNs) and multipath devices	19
Configuring logical volumes on multipath devices and creating a file system	19
Verifying that the host can write to and read from a multipath device	20
Where to find additional information	22
Copyright information	23
Trademark information	24
How to send comments about documentation and receive update notifications	25
Index	26

Deciding whether to use this guide

This guide describes how to quickly set up the iSCSI service on a Storage Virtual Machine (SVM), provision a LUN, and make the LUN available using an iSCSI initiator on a Red Hat Enterprise Linux server.

This guide is based on the following assumptions:

- You want to use best practices, not explore every available option.
- You do not want to read a lot of conceptual background.
- You want to use OnCommand System Manager, not the Data ONTAP command-line interface or an automated scripting tool.
- You are using the iSCSI software initiator on a Red Hat Enterprise Linux host running a supported version of Red Hat Enterprise Linux Series 6.4 or higher.
- Your network uses IPv4 addressing.
- You want to assign addresses to logical interfaces using any of the following methods:
 - Automatically, from a subnet you define
 - Manually, using an address selected from an existing subnet
 - Manually, using an address that will be added to an existing subnet
- You are not configuring iSCSI SAN boot.

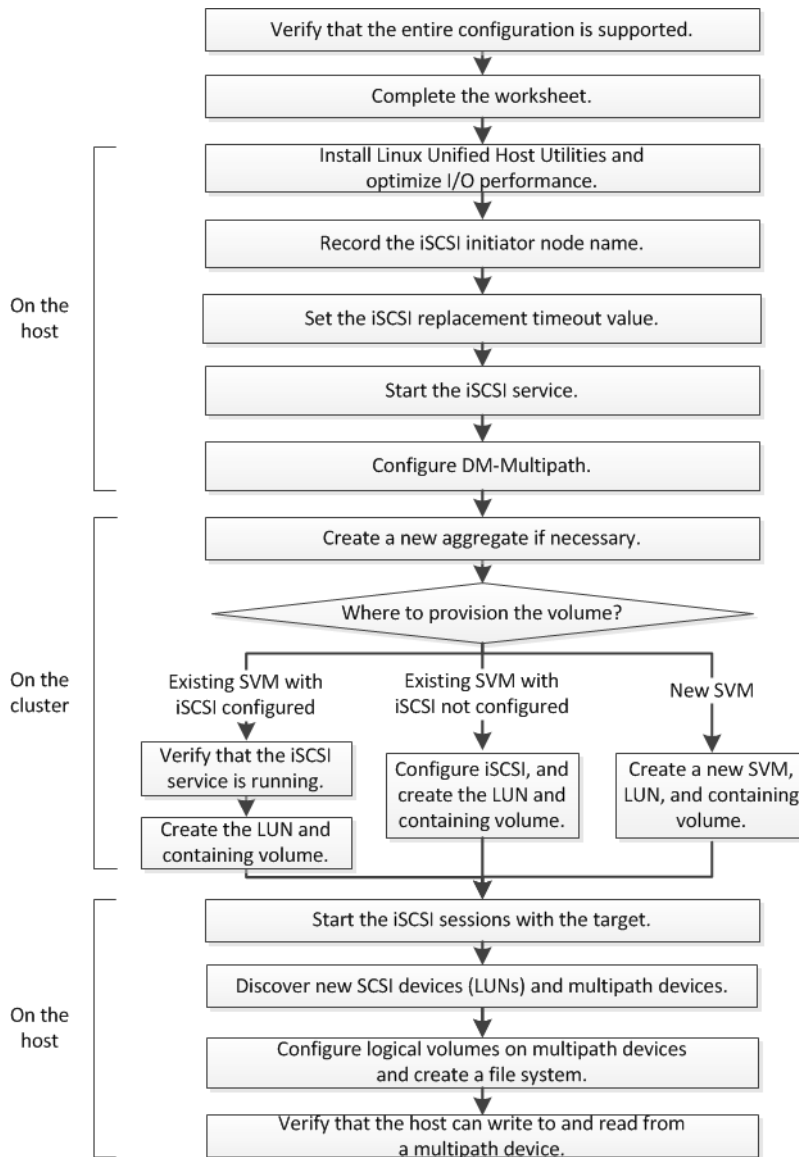
If these assumptions are not correct for your situation, you should see the following resources:

- [*Clustered Data ONTAP 8.3 SAN Administration Guide*](#)
- [*Clustered Data ONTAP 8.3 SAN Configuration Guide*](#)
- [*Linux Unified Host Utilities 7.0 Installation Guide*](#)
- [*NetApp Documentation: OnCommand Workflow Automation \(current releases\)*](#)

OnCommand Workflow Automation enables you to run prepackaged workflows that automate management tasks such as the workflows described in Express Guides.

iSCSI configuration and provisioning workflow

When you make storage available to a host using iSCSI, you provision a volume and LUN on the Storage Virtual Machine (SVM), and then connect to the LUN from the host.



Verifying that the iSCSI configuration is supported

To ensure reliable operation, you must verify that the entire iSCSI configuration is supported. The Interoperability Matrix lists the supported configurations.

Steps

1. Go to the [NetApp Interoperability Matrix Tool](#) to verify that you have a supported combination of the following components:

- Data ONTAP software
 - Host computer CPU architecture (for standard rack servers)
 - Specific processor blade model (for blade servers)
 - Storage protocol (iSCSI)
 - Linux operating system version
 - DM-Multipath package
 - Linux Unified Host Utilities
2. Click the configuration name for the selected configuration.
Details for that configuration are displayed in the Configuration Details window.
 3. Review the information in the following tabs:
 - Notes
Lists important alerts and information that are specific to your configuration.
Review the alerts to identify the packages that are required for your operating system.
 - Policies and Guidelines
Provides general guidelines for all SAN configurations.

Completing the iSCSI configuration worksheet

You require iSCSI identifiers, network addresses, and storage configuration information to perform iSCSI configuration tasks.

iSCSI identifiers


Initiator (host) iSCSI node name (IQN)	
Target alias (optional)	

Target network addresses

The Storage Virtual Machine (SVM) is the iSCSI target.

You require a subnet with two IP addresses for iSCSI data LIFs for each node in the cluster. There should be two separate networks for high availability. The specific IP addresses are assigned by Data ONTAP when you create the LIFs as part of creating the SVM.

If possible, separate iSCSI traffic on separate physical networks or on VLANs.



Subnet for LIFs: _____

Node or LIF with port to switch	IP address	Network mask	Gateway	VLAN ID	Home port
Node 1 / LIF to switch 1					
Node 2 / LIF to switch 1					
Node 3 / LIF to switch 1					
Node 4 / LIF to switch 1					
Node 1 / LIF to switch 2					
Node 2 / LIF to switch 2					
Node 3 / LIF to switch 2					
Node 4 / LIF to switch 2					

Storage configuration

If the aggregate and SVM are already created, record their names here; otherwise, you can create them as required:

Node to own LUN	
Aggregate name	
SVM name	

LUN information

LUN size	
LUN name (optional)	

LUN description (optional)	
----------------------------	--

SVM information

If you are not using an existing SVM, you require the following information to create a new one:

SVM name	
SVM IPspace	
Aggregate for SVM root volume	
SVM user name (optional)	
SVM password (optional)	
SVM management LIF (optional)	Subnet:
	IP address:
	Network mask:
	Gateway:
	Home node:
	Home port:

Installing Linux Unified Host Utilities and optimizing I/O performance

Linux Unified Host Utilities software includes the `sanlun` utility, a NetApp LUN reporting tool that you can use to display information about storage cluster nodes. You must also enable the correct server profile on the Linux host to optimize NetApp storage performance.

Before you begin

You must have identified the supported Linux Unified Host Utilities version for your configuration by using the Interoperability Matrix. You must also have the `tuned` package, which is part of your Linux OS distribution and contains the `tuned-adm` command, which you use to set the server profile on the host.

Steps

1. Download the supported Linux Unified Host Utilities version from the NetApp Support Site.
[NetApp Downloads: Software](#)
2. Install Linux Unified Host Utilities software by following the instructions in the installation documentation.
3. If the `tuned` package is not installed, enter the following command:
`yum install tuned`
4. For a physical host, ensure that you have set the `enterprise-storage` profile:
`tuned-adm profile enterprise-storage`
5. For a virtual host, ensure that you have set the `virtual-guest` profile:
`tuned-adm profile virtual-guest`

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

Recording the iSCSI node name

You must record the iSCSI initiator node name on the Linux host so that you have it when you configure the storage system.

Steps

1. Use a text editor to open the `/etc/iscsi/initiatorname.iscsi` file containing the node names.
2. Write down the node name or copy it to a text file.

Example

The iSCSI initiator node name looks something like this:

```
iqn.1994-05.com.redhat:127
```

Setting the iSCSI replacement timeout value

To support multipathing, you must modify the replacement timeout value in the iSCSI configuration file.

Step

1. Change the replacement timeout value in the `/etc/iscsi/iscsid.conf` file from the default of 120 to 5.

Example

The replacement timeout value looks like the following in `iscsid.conf`:

```
node.session.timeo.replacement_timeout = 5
```

Starting the iSCSI service

After you change the timeout value in the iSCSI configuration file, you must start the iSCSI service so that the host can discover the targets and access LUNs.

Steps

1. Start the iSCSI service.
 - a. If it's the first time you are starting the iSCSI service on the Linux host, enter:

```
service iscsid force-start
```

Note: You must run the `force-start` command the first time you start the iSCSI service. Running `/etc/init.d/iscsi start` without previously running the `force-start` command could result in an error message.
 - b. If you have previously started the iSCSI service on the Linux host, but iSCSI is not configured to start automatically at system boot, you must manually restart the service:

```
/etc/init.d/iscsi start
```

- To ensure that the iSCSI service always starts during system boot, run the following `chkconfig` commands separately (press Enter after each):

```
chkconfig iscsi on
chkconfig iscsid on
```

Configuring DM-Multipath

DM-Multipath manages multiple paths between the Linux host and the storage cluster. Configuring DM-Multipath on a LUN, which appears to the Linux host as a SCSI device, ensures that your Linux host can access its LUN on the storage cluster if a path or component fails.

Before you begin

You must have identified the required version of DM-Multipath from the [NetApp Interoperability Matrix Tool](#).

Steps

- Edit the `/etc/multipath.conf` file as follows:
 - Determine if there are non-NetApp SCSI devices to exclude (blacklist).

These are devices that are not displayed when you enter the `sanlun lun show` command.

- If there are non-NetApp SCSI devices to exclude, enter the WWID for the devices in the blacklist section of the `multipath.conf` file.

To display the WWID of a non-NetApp SCSI device, enter the following command on the device that you want to exclude, where `<SCSI_device_name>` is a device that you specify:

```
/lib/udev/scsi_id -gud /dev/<SCSI_device_name>
```

For example, if `/dev/sda` is the non-NetApp SCSI device that you want to exclude, you would type:

```
/lib/udev/scsi_id -gud /dev/sda
```

This displays the WWID of the device, which you can copy and paste in the `multipath.conf` file.

Example

In the following example showing the blacklist section of the `multipath.conf` file, the non-NetApp SCSI device with WWID `3600508e000000000753250f933cc4606` is excluded:

```
blacklist {
    wwid          3600508e000000000753250f933cc4606
    devnode      "^(ram|raw|loop|fd|md|dm-|sr|scd|st)[0-9]*"
    devnode      "^hd[a-z]"
    devnode      "^cciss.*"
}
```

- If there are no devices to exclude, remove the `wwid` line altogether from the `multipath.conf` file.
- Enable NetApp recommended DM-Multipath settings by appending the following string to the end of the kernel line in the boot loader:

```
rdloaddriver=scsi_dh_alua
```

2. Start the DM-Multipath daemon:

```
/etc/init.d/multipathd start
```

3. To ensure that the multipath daemon always starts during system boot, add the multipath service to the boot sequence:

```
chkconfig multipathd on
```

4. Reboot the Linux host.

5. Verify that the `rdloaderdriver` setting appears in the output when you run `cat /proc/cmdline`.

Example

The `rdloaderdriver` setting appears as one of several boot options that are passed to the kernel:

```
# cat /proc/cmdline
ro root=/dev/mapper/vg_ibmx3650210104-lv_root rd_NO_LUKS
LANG=en_US.UTF-8 rd_NO_MD rd_LVM_LV=vg_ibmx3650210104/lv_root
SYSEFONT=latacyrheb-sun16 rd_LVM_LV=vg_ibmx3650210104/lv_swap
crashkernel=129M@0M KEYBOARDTYPE=pc KEYTABLE=us rd_NO_DM rhgb quiet
rdloaderdriver=scsi_dh_alua
```

Creating an aggregate

If you do not want to use an existing aggregate, you can create a new aggregate to provide physical storage to the volume you are provisioning.

Steps

1. Enter the URL `https://IP-address-of-cluster-management-LIF` in a web browser and log in to System Manager using your cluster administrator credential.
2. In the navigation pane, expand the **Cluster** hierarchy and click **Storage > Aggregates**.
3. Click **Create**.
4. Follow the instructions on the screen to create the aggregate using the default RAID-DP configuration, and then click **Create**.

The screenshot shows a 'Create Aggregate' window with the following configuration:

- Name:** aggr2
- Disk Type:** SAS (with a 'Browse' button)
- Number of Disks:** 8 (with a note: 'Max: 8 (excluding 1 hot spare), min: 5 for RAID-DP')
- RAID Configuration:** RAID-DP; RAID group size of 16 disks (with a 'Change' link)
- New Usable Capacity:** 4.968 TB (Estimated)

Result

The aggregate is created with the specified configuration and added to the list of aggregates in the Aggregates window.

Deciding where to provision the volume

Before you provision a volume to contain your LUNs, you need to decide whether to add the volume to an existing Storage Virtual Machine (SVM) or to create a new SVM for the volume. You might also need to configure iSCSI on an existing SVM.

About this task

If an existing SVM is already configured with the needed protocols and has LIFs that can be accessed from the host, it is easier to use the existing SVM.

You can create a new SVM to separate data or administration from other users of the storage cluster. There is no advantage to using separate SVMs just to separate different protocols.

Choices

- If you want to provision volumes on an existing SVM that is already configured for iSCSI, see [Verifying that the iSCSI service is running on an existing SVM](#) on page 12.
- If you want to provision volumes on an existing SVM that has iSCSI enabled but not configured, see [Configuring iSCSI on an existing SVM](#) on page 14.
This is the case when you followed another Express Guide to create the SVM while configuring a different protocol.
- If you want to provision volumes on a new SVM, see [Creating a new SVM](#) on page 16.

Verifying that the iSCSI service is running on an existing SVM

If you choose to use an existing Storage Virtual Machine (SVM), you must verify that the iSCSI service is running on the SVM.

Before you begin

You must have selected an existing SVM on which you plan to create a new LUN.

Steps

1. From the OnCommand System Manager home page, double-click the appropriate storage system.
2. Expand the **Storage Virtual Machines** hierarchy in the left navigation pane.
3. In the navigation pane, select the SVM and click **Configuration > Protocols > iSCSI**.
4. Verify that the iSCSI service is running.

The screenshot displays the OnCommand System Manager interface for configuring iSCSI on an SVM. The left-hand navigation pane shows the following path: Cluster > Storage Virtual Machines > cluster-2 > iSCSI_1 > iSCSI. The main content area is titled 'iSCSI' and has two tabs: 'Service' and 'Initiator Security'. The 'Service' tab is active, showing the iSCSI service status as 'Running' (indicated by a green dot). Below the status, the following information is displayed:

- iSCSI Service: ● iSCSI service is running
- iSCSI Target Node Name: iqn.1992-08.com.netapp.sn.7a92a422fae711e384f7005056978047.vs.4
- iSCSI Target Alias: iSCSI_1

Below this information is a section titled 'iSCSI Interfaces' with a radio button to 'Enable' (selected) and a radio button to 'Disable'. A table lists the configured interfaces:

Network Interface	Target Portal Group	IP Address
iscsi_if_1	1028	10.53.33.4
iscsi_if_2	1029	10.53.33.5

- Record the iSCSI interfaces listed for the SVM.

After you finish

If the iSCSI service is not running, start the iSCSI service or create a new SVM.

If there are fewer than two iSCSI interfaces per node, update the iSCSI configuration on the SVM or create a new SVM for iSCSI.

Creating a LUN and its containing volume

You use the Create LUN wizard to create a LUN and the FlexVol volume that contains the LUN. The wizard also creates the igroup and maps the LUN to the igroup, which enables the specified host to access the LUN.

Before you begin

- There must be an aggregate with enough free space to contain the LUN.
- There must be a Storage Virtual Machine (SVM) with the iSCSI protocol enabled and the appropriate LIFs created.
- You must have recorded the iSCSI initiator node name of the host.

Starting with clustered Data ONTAP 8.3, LUNs are mapped to a subset of the initiators in the igroup to limit the number of paths from the host to the LUN.

- By default, Data ONTAP 8.3 uses Selective LUN Map to make the LUN accessible only through paths on the node owning the LUN and its HA partner.
- You still must configure all of the iSCSI LIFs on every node for LUN mobility in case the LUN is moved to another node in the cluster.
- When moving a volume or a LUN, you must modify the Selective LUN Map reporting-nodes list before moving.

About this task

If your organization has a naming convention, you should use names for the LUN, volume, and so on that fit your convention. Otherwise, you should accept the default names.

Steps

- From the OnCommand System Manager home page, double-click the appropriate storage system.
- Expand the **Storage Virtual Machines** hierarchy in the left navigation pane.
- In the navigation pane, select the SVM, and then click **Storage > LUNs**.
- In the **LUN Management** tab, click **Create**, and then type or select information as prompted by the wizard.
- On the **General Properties** page, select the LUN type **Linux** for LUNs used directly by the Linux host.

Leave the **Thin Provisioned** check box unselected.



You can specify the size of the LUN. Storage will be optimized according to the type selected.

Type:

Size:

Thin Provisioned

- On the **LUN Container** page, create a new FlexVol volume.

- On the **Initiators Mapping** page, click **Add Initiator Group**, enter the required information on the **General** tab, and then on the **Initiators** tab, enter the iSCSI initiator node name of the host that you recorded.
- Confirm the details, and then click **Finish** to complete the wizard.

Related information

[Clustered Data ONTAP 8.3 System Administration Guide](#)

Configuring iSCSI on an existing SVM

You can configure iSCSI on an existing Storage Virtual Machine (SVM) and create a LUN and its containing volume with a single wizard. The iSCSI protocol must already be enabled but not configured on the SVM. This information is intended for SVMs for which you are configuring multiple protocols, but have not yet configured iSCSI.

Before you begin

You must have enough network addresses available to create two LIFs for each node.

About this task

Starting with clustered Data ONTAP 8.3, LUNs are mapped to a subset of the initiators in the igroup to limit the number of paths from the host to the LUN.

- By default, Data ONTAP 8.3 uses Selective LUN Map to make the LUN accessible only through paths on the node owning the LUN and its HA partner.

- You still must configure all of the iSCSI LIFs on every node for LUN mobility in case the LUN is moved to another node in the cluster.
- When moving a volume or a LUN, you must modify the Selective LUN Map reporting-nodes list before moving.

Steps

1. In the navigation pane, expand the **Storage Virtual Machines** hierarchy and select the SVM.
2. In the SVM **Details** pane, verify that **iSCSI** is displayed with a gray background, which indicates that the protocol is enabled but not fully configured.

If **iSCSI** is displayed with a green background, the SVM is already configured.



3. Click the **iSCSI** protocol link with the gray background.
The Configure iSCSI Protocol window is displayed.
4. Configure the iSCSI service and LIFs from the **Configure iSCSI protocol** page:
 - a. Optional: Enter a target alias name.
 - b. Enter **2** in the **LIFs per node** field.
Two LIFs are required for each node, to ensure availability and data mobility.
 - c. Select the subnet for the LIFs.
 - d. In the **Provision a LUN for iSCSI storage** area, enter the desired LUN size, host type, and iSCSI initiator name of the host.
 - e. Click **Submit & Close**.

Example

5. Review the **Summary** page, record the LIF information, and then click **OK**.

Creating a new SVM

The Storage Virtual Machine (SVM) provides the iSCSI target through which a host accesses LUNs. When you create the SVM, you also create logical interfaces (LIFs) and the LUN and its containing volume. You can create a new SVM to separate data or administration from other users of the storage cluster.

Before you begin

- You must have enough network addresses available to create two LIFs for each node.

About this task

Starting with clustered Data ONTAP 8.3, LUNs are mapped to a subset of the initiators in the igroup to limit the number of paths from the host to the LUN.

- By default, Data ONTAP 8.3 uses Selective LUN Map to make the LUN accessible only through paths on the node owning the LUN and its HA partner.
- You still must configure all of the iSCSI LIFs on every node for LUN mobility in case the LUN is moved to another node in the cluster.
- When moving a volume or a LUN, you must modify the Selective LUN Map reporting-nodes list before moving.

Steps

1. Expand the **Storage Virtual Machines** hierarchy in the left navigation pane.
2. Click **Create**.
3. In the **Storage Virtual Machine (SVM) Setup** window, create the SVM:

SVM Details

? Specify a unique name and the data protocols for the SVM

SVM Name:

? IPspace:

? Volume Type: FlexVol volumes Infinite Volume

An SVM can contain either multiple FlexVol volumes or a single Infinite Volume.
You cannot change the volume type of the SVM after you set it.

? Data Protocols: CIFS NFS iSCSI FC/FCoE

? Default Language:

The language of the SVM determines the character set used to display the file names and data for all NAS volumes in the SVM. Therefore, you must set the language with correct value.

? Security Style:

Root Aggregate:

- a. Specify a unique name for the SVM.

The name must either be a fully qualified domain name (FQDN) or follow another convention that ensures unique names across a cluster.
- b. Select the IPspace that the SVM will belong to.

If the cluster does not use multiple IPspaces, the “Default” IPspace is used.
- c. Keep the default volume type selection.

Only FlexVol volumes are supported with SAN protocols.
- d. Select all of the protocols that you have licenses for and that you will ever use on the SVM, even if you do not want to configure all of the protocols immediately.

Selecting both NFS and CIFS when you create the SVM enables these two protocols to share the same LIFs. Adding these protocols later does not allow them to share LIFs.

If CIFS is one of the protocols you selected, then the security style is set to NTFS. Otherwise, the security style is set to UNIX.
- e. Keep the default language setting C.UTF-8.
- f. Select the desired root aggregate to contain the SVM root volume.

The aggregate for the data volume is selected separately in a later step.
- g. Click **Submit & Continue**.

The SVM is created, but protocols are not yet configured.
4. If the **Configure CIFS/NFS protocol** page appears because you enabled CIFS or NFS, click **Skip** and then configure CIFS or NFS later.
5. Configure the iSCSI service and create LIFs, and the LUN and its containing volume from the **Configure iSCSI protocol** page:
 - a. Optional: Enter a target alias name.
 - b. Select the subnet for the LIFs.
 - c. Enter **2** in the **LIFs per node** field.

Two LIFs are required for each node to ensure availability and data mobility.
 - d. Click **Submit & Continue**.

Example

6. If the **Configure FC/FCoE protocol** page appears because you enabled FC, click **Skip** and then configure FC later.
7. When the **SVM Administration** appears, configure or defer configuring a separate administrator for this SVM:
 - Click **Skip** and configure an administrator later if desired.
 - Enter the requested information, and then click **Submit & Continue**.
8. Review the **Summary** page, record the LIF information, and then click **OK**.

Related information

[NetApp Documentation: Clustered Data ONTAP Express Guides](#)

Starting the iSCSI sessions with the target

The Linux host must have an iSCSI connection with each node (target) in the storage cluster. You establish the sessions from the host.

About this task

In clustered Data ONTAP, the iSCSI host must have paths to each node in the cluster. DM-Multipath selects the best paths to use. If paths fail, DM-Multipath selects alternate paths.

Steps

1. Start the iSCSI service from the Linux host command prompt:


```
service iscsid force-start
```
2. Discover the iSCSI target:

```
iscsiadm --mode discovery --op update --type sendtargets --portal
targetIP
```

The IP address of each discovered target displays on a separate line.

3. Establish iSCSI sessions by logging in to each target:

```
iscsiadm --mode node -l all
```

4. You can view a list of the active iSCSI sessions:

```
iscsiadm --mode session
```

Discovering new SCSI devices (LUNs) and multipath devices

LUNs on the storage cluster appear to the Linux host as SCSI devices, which are I/O paths that DM-Multipath aggregates into a new device, called a multipath device. The host does not automatically discover new SCSI devices (LUNs) that you add to your system. You must manually rescan them to discover them.

Before you begin

You must have a copy of the `rescan` script, which is in the `sg3_utils` package that is part of your Linux OS distribution.

Steps

1. Discover new SCSI devices (LUNs) and create the corresponding multipath devices for the LUNs:

```
/usr/bin/rescan-scsi-bus.sh
```

2. Verify DM-Multipath configuration:

```
multipath -ll
```

Example

This displays the following type of output, listing the recommended settings for each NetApp LUN:

```
3600a0980324666546e2b443251655177 dm-2 NETAPP,LUN C-Mode
size=10G features='4 queue_if_no_path pg_init_retries 50
retain_attached_hw_handle' hwhandler='1 alua' wp=rw
|+- policy='round-robin 0' prio=50 status=active
| | - 0:0:1:0 sdb 8:16 active ready running
| | - 0:0:0:0 sda 8:0 active ready running
| | - 1:0:0:0 sde 8:64 active ready running
| | - 1:0:1:0 sdf 8:80 active ready running
|+- policy='round-robin 0' prio=10 status=enabled
| | - 0:0:3:0 sdd 8:48 active ready running
| | - 1:0:3:0 sdh 8:112 active ready running
| | - 0:0:2:0 sdc 8:32 active ready running
| | - 1:0:2:0 sdg 8:96 active ready running
```

Configuring logical volumes on multipath devices and creating a file system

When the Linux host first accesses a new SCSI device (LUN), there is no partition or file system. If you want to use a partitioned multipath device, you must first partition the underlying SCSI devices.

You might also want to create logical volumes on multipath devices, and optionally create a file system.

Before you begin

SCSI devices and corresponding multipath devices must have been discovered by the Linux host.

About this task

At a minimum DM-Multipath should be configured on the SCSI devices. Any configuration beyond DM-Multipath is optional. The following steps are general guidelines if you wish to perform additional configuration, such as partitioning, configuring logical volumes on multipath devices, and creating a file system. For details on Linux commands, refer to Red Hat Enterprise Linux documentation and man pages.

Steps

1. To use a partitioned multipath device, first partition the underlying SCSI devices using the `fdisk` or `parted` utility.
2. Create corresponding multipath partitions using the `kpartx` utility.
3. Create logical volumes on corresponding multipath devices using Logical Volume Manager (LVM) commands.
4. Create a file system, such as Ext4 or XFS, on logical volumes or multipath devices using the `mkfs` utility.

Verifying that the host can write to and read from a multipath device

Before using a multipath device, you should verify that the host can write data to the multipath device and read it back.

Before you begin

DM-Multipath must be configured on the Linux host for the storage cluster nodes you are accessing.

About this task

If the storage cluster node that the multipath device corresponds to can be failed over to its partner node, you should verify reading the data while the node is failed over. This verification might not be possible if the storage cluster is in production use.

Steps

1. Run some I/O on the specified multipath device:


```
dd if=/dev/zero of=<multipath_device_name>
```
2. Verify that I/O is running on the multipath device, specifying the number of seconds to refresh statistics (for example, every two seconds, as shown):

```
iostat 2
```

Example

When I/O is running on the multipath device, you can see the various I/O counters output change dynamically, indicating that I/O is running successfully on the multipath device.

```

avg-cpu:  %user   %nice %system %iowait  %steal   %idle
           0.00    0.00   0.01   0.00    0.00   99.98

Device:            tps    Blk_read/s    Blk_wrtn/s    Blk_read    Blk_wrtn
sda                0.18         1.34          2.29         832606     1428026
dm-0               0.37         1.29          2.29         801530     1427984
dm-1               0.00         0.00          0.00          2576        0
dm-2               0.00         0.00          0.00          1770        24
sdd                0.17         1.18         17.87        734688     11128584
sde                0.17         1.18         18.01        734560     11219016
sdf                0.00         0.00          0.00          1344        0
sdg                0.00         0.00          0.00          1344        0
dm-3               0.68         4.71         71.96       2932496     44816008
sdh                0.17         1.18         18.22        734360     11343920
sdi                0.17         1.18         17.86        734272     11124488
sdj                0.00         0.00          0.00          1344        0
sdk                0.00         0.00          0.00          1344        0
.....

```

3. If possible, run the takeover command on the storage cluster node:
storage failover takeover -ofnode <node_name>
4. Verify that I/O is still running on the multipath device by re-running the `iostat` command.
5. Run the giveback command on the storage cluster node:
storage failover giveback -ofnode <node_name>
6. Check `iostat` output to verify that I/O is still running.

After you finish

If any of the tests fail, verify that the iSCSI service is running and re-check the DM-Multipath configuration and the iSCSI paths to the multipath device.

Where to find additional information

After you have successfully tested iSCSI, you can perform advanced iSCSI configuration or add SAN or NAS protocols. There are express guides, comprehensive guides, and technical reports to help you achieve these goals.

SAN configuration

All of the following documentation is available:

- [*Clustered Data ONTAP 8.3 SAN Configuration Guide*](#)
Describes supported FC, iSCSI, and FCoE topologies for connecting host computers to storage controllers in clusters.
- [*Clustered Data ONTAP 8.3 SAN Administration Guide*](#)
Describes how to configure and manage the iSCSI, FCoE, and FC protocols for clustered SAN environments, including configuration of LUNs, igroups, and targets.
- [*Linux Unified Host Utilities 7.0 Installation Guide*](#)
Describes how to install Linux Unified Host Utilities software and configure DM-Multipath.
- [*NetApp Technical Report 4159: Red Hat Enterprise Virtualization 3.1 and NetApp Clustered Data ONTAP Storage Best Practices Guide*](#)
Describes the current best practices for Red Hat virtualized environments.
- Red Hat Enterprise Linux 6 documentation and man pages
Provides instructions for managing storage devices, file systems, and Linux commands.

Additional protocol configuration

If you want to provide additional SAN or NAS access to the SVM, you can use any of the protocol configuration express guides.

[*NetApp Documentation: Clustered Data ONTAP Express Guides*](#)

Copyright information

Copyright © 1994–2015 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 at <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 to 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

Index

A

- about this guide
 - deciding whether to use [4](#)
- additional information, where to find
 - iSCSI configuration [22](#)
- aggregates
 - creating [11](#)
- audience
 - for this guide [4](#)
- automatically starting
 - DM-Multipath at system boot [10](#)

B

- blacklisting
 - excluding devices in multipath.conf [10](#)

C

- clusters
 - starting iSCSI sessions with [18](#)
- comments
 - how to send feedback about documentation [25](#)
- configuration
 - iSCSI workflow [5](#)
 - where to find information about iSCSI [22](#)
- configuring
 - DM-Multipath [10](#)
 - iSCSI on existing SVMs [14](#)
 - iSCSI worksheet for [6](#)
- creating
 - a file system [19](#)
 - aggregates [11](#)
 - logical volumes on multipath devices [19](#)
 - LUNs and their containing volumes [13](#)
 - SVMs [16](#)

D

- devices
 - excluding (blacklisting) in multipath.conf [10](#)
- discovering
 - new SCSI devices and multipath devices [19](#)
- DM-Multipath
 - configuring on Red Hat Enterprise Linux hosts [10](#)
- documentation
 - how to receive automatic notification of changes to [25](#)
 - how to send feedback about [25](#)

E

- editing
 - iscsid.conf [9](#)
 - multipath.conf [10](#)
- excluding
 - devices in multipath.conf [10](#)

- express guides
 - requirements for using iSCSI with Red Hat Enterprise Linux [4](#)

F

- feedback
 - how to send comments about documentation [25](#)
- file system
 - creating on new multipath devices [19](#)
- FlexVol volumes
 - creating [13](#)
- flowcharts
 - iSCSI configuration [5](#)
- formatting
 - a new SCSI device (LUN) [19](#)

G

- guides, express
 - requirements for using iSCSI with Red Hat Enterprise Linux [4](#)

H

- hosts
 - installing host utilities on [8](#)
 - setting the server profile [8](#)
 - verifying writing to and reading from multipath devices [20](#)

I

- igroups
 - creating and mapping the LUN to [13](#)
- information
 - how to send feedback about improving documentation [25](#)
 - where to find additional iSCSI configuration [22](#)
- initiator node names
 - recording iSCSI [9](#)
- initiators
 - recording the iSCSI node name [9](#)
 - starting sessions with targets [18](#)
- installing
 - host utilities [8](#)
- Interoperability Matrix
 - verifying supported iSCSI configurations using [5](#)
- IQN
 - iSCSI configuration worksheet for [6](#)
 - recording iSCSI [9](#)
- iSCSI
 - configuration workflow [5](#)
 - configuration worksheet for [6](#)
 - configuring on existing SVMs [14](#)
 - connecting initiators to targets [18](#)
 - recording the initiator node name [9](#)
 - starting sessions with targets [18](#)

- verifying that configuration is supported [5](#)
 - where to find additional information about configuring [22](#)
 - iSCSI replacement timeout value
 - modifying [9](#)
 - iSCSI service
 - starting [9](#)
 - verifying running on SVM [12](#)
 - iSCSI targets
 - creating as part of configuring SVMs [14](#)
 - creating as part of creating SVMs [16](#)
 - iscsid.conf
 - modifying the replacement timeout value in [9](#)
- L**
- LIFs
 - creating as part of configuring SVMs [14](#)
 - creating as part of creating SVMs [16](#)
 - Linux hosts, Red Hat
 - packages required for DM-Multipath [5, 19](#)
 - starting iSCSI sessions with targets [18](#)
 - verifying supported iSCSI configurations [5](#)
 - Linux hosts, Red Hat Enterprise
 - configuring DM-Multipath on [10](#)
 - requirements for using iSCSI [4](#)
 - logical interfaces
 - See* LIFs
 - LUNs
 - creating [13](#)
 - deciding where to provision the containing volume on a new or existing SVM [12](#)
 - discovering and creating multipath devices on [19](#)
 - See also* SCSI devices (LUNs)
 - discovering new [19](#)
 - ensuring access to with DM-Multipath [10](#)
 - how they appear to a Linux host [19](#)
 - mapping to an igroup [13](#)
 - partitioning and creating a file system on [19](#)
 - requirements for using iSCSI with Red Hat Enterprise Linux [4](#)
- M**
- mapping
 - LUNs to an igroup [13](#)
 - multipath devices
 - creating logical volumes and a file system on [19](#)
 - discovering new [19](#)
 - verifying that the host can write to and read from [20](#)
 - multipath.conf
 - editing [10](#)
- N**
- network addresses
 - iSCSI configuration worksheet for [6](#)
 - nodes
 - starting iSCSI sessions with [18](#)

- P**
- packages
 - required for DM-Multipath [5, 19](#)
 - required to set the server profile [8](#)
 - partitions
 - creating for new SCSI devices (LUN) [19](#)
 - provisioning
 - iSCSI workflow [5](#)
 - iSCSI worksheet for [6](#)
 - where to find information about iSCSI [22](#)
- R**
- read/write
 - verifying that the host can write to and read from multipath devices [20](#)
 - Red Hat hosts
 - See* Linux hosts, Red Hat
 - replacement timeout value, iSCSI
 - modifying [9](#)
 - requirements
 - for using iSCSI with Red Hat Enterprise Linux [4](#)
 - verifying supported iSCSI configurations [5](#)
- S**
- SCSI devices (LUNs)
 - discovering and creating multipath devices on [19](#)
 - See also* LUNs
 - discovering new [19](#)
 - partitioning and creating a file system on [19](#)
 - server profile
 - setting [8](#)
 - starting
 - DM-Multipath [10](#)
 - the iSCSI service [9](#)
 - storage configurations
 - iSCSI configuration worksheet [6](#)
 - suggestions
 - how to send feedback about documentation [25](#)
 - supported configurations
 - verifying iSCSI [5](#)
 - SVMs
 - configuring iSCSI on existing [14](#)
 - creating [16](#)
 - deciding whether to provision a volume on a new or existing [12](#)
 - verifying iSCSI service is running [12](#)
- T**
- targets
 - creating iSCSI as part of configuring SVMs [14](#)
 - creating iSCSI as part of creating SVMs [16](#)
 - starting iSCSI sessions with [18](#)
 - timeout value, iSCSI
 - modifying [9](#)
 - twitter
 - how to receive automatic notification of documentation changes [25](#)

U

updates, Linux
 See packages

V

volumes
 creating [13](#)
 deciding whether to provision on a new or existing
 SVM [12](#)

W

wizards
 running the Create LUN [13](#)
workflows
 iSCSI configuration [5](#)
worksheets
 iSCSI configuration [6](#)
write/read
 verifying that the host can write to and read from
 multipath devices [20](#)