



FC configuration for Red Hat Enterprise Linux

System Manager Classic

NetApp

September 05, 2025

Table of Contents

FC configuration for Red Hat Enterprise Linux	1
FC configuration for Red Hat Enterprise Linux overview	1
Other ways to do this in ONTAP	1
FC configuration and provisioning workflow	1
Verify that the FC configuration is supported	3
Complete the FC configuration worksheet	3
Install the HBA utility from the HBA vendor	5
Update the HBA driver, firmware, and BIOS	6
Install Linux Unified Host Utilities and optimize I/O performance	6
Record the WWPN for each host FC port	7
Configure DM-Multipath	8
Create an aggregate	10
Decide where to provision the volume	11
Zone the FC switches by the host and LIF WWPNs	16
Discover new SCSI devices (LUNs) and multipath devices	17
Configure logical volumes on multipath devices and creating a file system	18
Verify that the host can write to and read from a multipath device	18

FC configuration for Red Hat Enterprise Linux

FC configuration for Red Hat Enterprise Linux overview

You can quickly set up the FC service on a storage virtual machine (SVM), provision a LUN, and make the LUN available using an FC host bus adapter (HBA) on a Red Hat Enterprise Linux server using the ONTAP System Manager classic interface (ONTAP 9.7 and earlier).

Use this procedure to set up FC service on an SVM if:

- The host is running a supported version of Red Hat Enterprise Linux 6.4 or later.
- You are using traditional FC HBAs and switches.



Fibre Channel over Ethernet (FCoE) is not covered in this procedure.

- You have at least two FC target ports available on each node in the cluster. Onboard FC and UTA2 (also called “CNA”) ports, as well as some adapters, are configurable. Configuring those ports is done in the ONTAP command-line interface (CLI) and is not covered in this procedure.
- You are not configuring FC SAN boot.

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


- [SAN administration](#)
- [SAN configuration](#)
- [Linux Unified Host Utilities 7.1 Installation](#)

Other ways to do this in ONTAP

To perform these tasks with...	Refer to...
The redesigned System Manager (available with ONTAP 9.7 and later)	Provision SAN storage for Linux servers
The ONTAP command line interface (CLI)	LUN setup workflow with the CLI

FC configuration and provisioning workflow

When you make storage available to a host using FC, you provision a volume and LUN on the storage virtual machine (SVM), and then connect to the LUN from the host.



Verify that the FC configuration is supported

For reliable operation, you must verify that the entire FC configuration is supported.

Steps

1. Go to the Interoperability Matrix to verify that you have a supported combination of the following components:
 - ONTAP software
 - Host computer CPU architecture (for standard rack servers)
 - Specific processor blade model (for blade servers)
 - FC host bus adapter (HBA) model and driver, firmware, and BIOS versions
 - Storage protocol (FC)
 - Linux operating system version
 - DM-Multipath package
 - Linux Unified Host Utilities
2. Select 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.

Complete the FC configuration worksheet

You require FC initiator and target WWPNs and storage configuration information to perform FC configuration tasks.

FC host WWPNs

Port	WWPN
Initiator (host) port connected to FC switch 1	
Initiator (host) port connected to FC switch 2	

FC target WWPNs

You require two FC data LIFs for each node in the cluster. The WWPNs are assigned by ONTAP when you create the LIFs as part of creating the storage virtual machine (SVM).

LIF	WWPN
Node 1 LIF with port connected to FC switch 1	
Node 2 LIF with port connected to FC switch 1	
Node 3 LIF with port connected to FC switch 1	
Node 4 LIF with port connected to FC switch 1	
Node 1 LIF with port connected to FC switch 2	
Node 2 LIF with port connected to FC switch 2	
Node 3 LIF with port connected to FC switch 2	
Node 4 LIF with port connected to FC 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 name	
SVM user name (optional)	
SVM password (optional)	
SVM management LIF (optional)	
	Subnet:
	IP address:
	Network mask:
	Gateway:
	Home node:
	Home port:

Install the HBA utility from the HBA vendor

The HBA utility enables you to view the worldwide port name (WWPN) of each FC port. The utility is also useful for troubleshooting FC issues.

About this task

Each HBA vendor offers an HBA utility for their FC HBAs. You must download the correct version for your host operating system and CPU.

The following is a partial list of HBA utilities:

- Emulex HBA Manager, formerly known as OneCommand Manager, for Emulex HBAs
- QLogic QConvergeConsole for QLogic HBAs

Steps

1. Download the appropriate utility from your HBA vendor's web site.
2. Run the installation program and follow the prompts to complete the installation.

Related information

[Broadcom \(Emulex\) Support Documents and Downloads](#)

[Emulex HBA Manager](#)

[QLogic: NetApp Downloads](#)

Update the HBA driver, firmware, and BIOS

If the FC host bus adapters (HBAs) in the Linux host are not running supported driver, firmware, and BIOS versions, you must update them.

Before you begin

You must have identified the supported driver, firmware, and BIOS versions for your configuration from the Interoperability Matrix tool.

[NetApp Interoperability Matrix Tool](#)

About this task

Drivers, firmware, BIOS, and HBA utilities are provided by the HBA vendors.

Steps

1. List the installed HBA driver, firmware, and BIOS versions by using the HBA utility from your HBA vendor.
2. Download and install the new driver, firmware, and BIOS as needed from the HBA vendor's support site.

Installation instructions and any required installation utilities are available with the download.

Install Linux Unified Host Utilities and optimize 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 and Host Bus Adapters (HBAs) that are on the Linux host. 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 7.1 Installation](#)

Record the WWPN for each host FC port

The worldwide port name (WWPN) is required to zone the FC switches and to create the igroups that allow the host to access its LUN.

Before you begin

You must have installed the vendor's HBA utility for the HBAs in your host and verified HBAs are running supported driver, firmware, and BIOS versions for your configuration.


About this task

The WWPN is used for all configuration. You do not have to record the worldwide node name (WWNN).

Steps

1. Run the HBA utility for your FC HBA type.
2. Select the HBA.
3. Record the WWPN of each port.

The following example shows Emulex HBA Manager, formerly known as OneCommand Manager.



Other utilities, such as QLogic QConvergeConsole, provide the equivalent information.

4. Repeat the previous step for each FC HBA in the host.

In Linux, you can also obtain the WWPN by running the `sanlun` utility.

The following example shows output from the `sanlun` command.

```

# sanlun fcp show adapter -v

adapter name: host0
**WWPN: 10000000c9813a14**
WWNN: 20000000c9813a14
driver name: lpfcc
model: LPe12002-M8
model description: Emulex LPe12002-M8 8Gb 2-port PCIe Fibre Channel
Adapter
serial number: VM84364896
hardware version: 31004549
driver version: 8.3.7.34.3p; HBAAPI(I) v2.3.b, 07-12-10
firmware version: 2.01A12 (U3D2.01A12)
Number of ports: 1
port type: Fabric
port state: Operational
supported speed: 2 GBit/sec, 4 GBit/sec, 8 GBit/sec
negotiated speed: 8 GBit/sec
OS device name: /sys/class/scsi_host/host0

adapter name: host5
**WWPN: 10000000c9813a15**
WWNN: 20000000c9813a15
driver name: lpfcc
model: LPe12002-M8
model description: Emulex LPe12002-M8 8Gb 2-port PCIe Fibre Channel
Adapter
serial number: VM84364896
hardware version: 31004549
driver version: 8.3.7.34.3p; HBAAPI(I) v2.3.b, 07-12-10
firmware version: 2.01A12 (U3D2.01A12)
Number of ports: 1
port type: Fabric
port state: Operational
supported speed: 2 GBit/sec, 4 GBit/sec, 8 GBit/sec
negotiated speed: 8 GBit/sec
OS device name: /sys/class/scsi_host/host5

```

Configure 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, enables your Linux host to 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 Interoperability Matrix tool.

NetApp Interoperability Matrix Tool



The recommended approach is that you configure two paths per node into each fabric or Ethernet network. This allows a path to fail without the node switching over to its partner. Using link aggregation (LAG) for the physical ports is not recommended with a Windows-based operating system.

Steps

1. Edit the `/etc/multipath.conf` file as follows:

- a. Determine whether 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 worldwide identifier (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 enter the following:

```
/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.

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 from the `multipath.conf` file.

- b. 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. Add the multipath service to the boot sequence so that the multipath daemon always starts during system boot:

```
chkconfig multipathd on
```

4. Reboot the Linux host.

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

The `rdloaddriver` 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 SYSFONT=latarcyrheb-sun16
rd_LVM_LV=vg_ibmx3650210104/lv_swap crashkernel=129M@0M KEYBOARDTYPE=pc
KEYTABLE=us rd_NO_DM rhgb quiet **rdloaddriver=scsi_dh_alua**
```

Create 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 which 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. Navigate to the **Aggregates** window.
3. Select **Create**.
4. Follow the instructions on the screen to create the aggregate using the default RAID-DP configuration, and then select **Create**.

Create Aggregate

To create an aggregate, select a disk type then specify the number of disks.

Name:	<input type="text" value="aggr2"/>
Disk Type:	<input type="text" value="SAS"/> Browse
Number of Disks:	<input type="text" value="8"/> <small>Max: 8 (excluding 1 hot spare), min: 5 for RAID-DP</small>
RAID Configuration:	RAID-DP; RAID group size of 16 disks Change
New Usable Capacity:	4.968 TB (Estimated)

Results

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

window.

Decide 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 FC 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.

Procedure

- If you want to provision volumes on an SVM that is already configured for FC, you must verify that the FC service is running and then create a LUN on the SVM.

[Verifying that the FC service is running on an existing SVM](#)

[Creating a LUN](#)

- If you want to provision volumes on an existing SVM that has FC enabled but not configured, configure iSCSI on the existing SVM.

[Configuring FC on an existing SVM](#)

This is the case if you did not follow this procedure to create the SVM while configuring a different protocol.

- If you want to provision volumes on a new SVM, create the SVM.

[Creating a new SVM](#)

Verify that the FC service is running on an existing SVM

If you choose to use an existing storage virtual machine (SVM), you must verify that the FC service is running on the SVM by using ONTAP System Manager. You must also verify that FC logical interfaces (LIFs) are already created.

Before you begin

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

Steps

1. Navigate to the **SVMs** window.
2. Select the required SVM.
3. Select the **SVM Settings** tab.
4. In the **Protocols** pane, select **FC/FCoE**.
5. Verify that the FC service is running.

Status: FC/FCoE service is running

WWNN: 20:05:00:a0:98:29:18:76

FC/FCoE Interfaces:

Network Interface	WWPN	Current Port	Status
FC_1_1	20:10:00:a0:98:29:18:76	innovate-01:0b	Enabled
FC_2_1	20:11:00:a0:98:29:18:76	innovate-02:0b	Enabled
FC_1_2	20:04:00:a0:98:29:18:76	innovate-02:0a	Enabled
FC_2_2	20:03:00:a0:98:29:18:76	innovate-01:0a	Enabled

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

6. Verify that there are at least two FC LIFs listed for each node.

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

Create a LUN

You use the Create LUN wizard to create a 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 FC protocol enabled and the appropriate logical interfaces (LIFs) created.
- You must have recorded the worldwide port names (WWPNs) of the host FC ports.

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

1. Navigate to the **LUNs** window.
2. Select **Create**.
3. Browse and select an SVM in which you want to create the LUNs.

The Create LUN Wizard is displayed.

4. 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:	Linux	
Size:	750	GB

Thin Provisioned

5. On the **LUN Container** page, select an existing FlexVol volume.

You must ensure that there is enough space in the volume. If sufficient space is not available in the existing volumes, you can create a new volume.

6. On the **Initiators Mapping** page, select **Add Initiator Group**, enter the required information on the **General** tab, and then on the **Initiators** tab, enter all the WWPNs of the host FC ports that you recorded.
7. Confirm the details, and then select **Finish** to complete the wizard.

Related information

System administration

Configure FC on an existing SVM

You can configure FC on an existing storage virtual machine (SVM) and create a LUN and its containing volume with a single wizard. The FC 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 FC.

Before you begin

Your FC fabric must be configured and the desired physical ports must be connected to the fabric.

Steps

1. Navigate to the **SVMs** window.
2. Select the SVM that you want to configure.
3. In the **SVMDetails** pane, verify that **FC/FCoE** is displayed with a gray background, which indicates that the protocol is enabled but not fully configured.

If **FC/FCoE** is displayed with a green background, the SVM is already configured.

Details

Protocols:	NFS CIFS FC/FCoE iSCSI
------------	--

4. Select the **FC/FCoE** protocol link with the gray background.

The Configure FC/FCoE Protocol window is displayed.

5. Configure the FC service and LIFs from the **Configure FC/FCoE protocol** page:
 - a. Select the **Configure Data LIFs for FC** check box.

b. Enter 2 in the **LIFs per node** field.

Two LIFs are required for each node, to ensure availability and data mobility.

c. Select **Submit & Close**.

Configure FC/FCoE protocol

[?](#) Configure LIFs to access the data using FC/FCoE protocol

Data Interface (LIF) Configuration

Both FC and FCoE enabled hardware found. Click on the appropriate checkbox to configure the FC and/or FCoE LIFs.

Configure Data LIFs for FC

LIFs per node: (Minimum: 1, Maximum: 2)

Provision a Lun for FCP storage (Optional)

Lun Size: GB

LUN OS Type:

Host Initiator:

Review or Edit the Interface Association

Configure Data LIFs for FCoE

6. Review the **Summary** page, record the LIF information, and then select **OK**.

Create a new SVM

The storage virtual machine (SVM) provides the FC 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 an SVM to separate the data and administration functions of a user from those of the other users in a cluster.

Before you begin

- Your FC fabric must be configured and the desired physical ports must be connected to the fabric.

Steps

1. Navigate to the **SVMs** window.
2. Select **Create**.
3. In the **Storage Virtual Machine (SVM) Setup** window, create the SVM:

Storage Virtual Machine (SVM) Setup



Enter SVM basic details

SVM Details

② Specify a unique name and the data protocols for the SVM

SVM Name:

② IPspace: ▼

② Data Protocols: CIFS NFS iSCSI FC/FCoE NVMe

② Default Language: ▼

The language of the SVM specifies the default language encoding setting for the SVM and its volumes. Using a setting that incorporates UTF-8 character encoding is recommended.

② 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 might 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. Select **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, select **Skip** and

then configure CIFS or NFS later.

5. If the **Configure iSCSI protocol** page appears because you enabled iSCSI, select **Skip** and then configure iSCSI later.
6. Configure the FC service and create LIFs, and the LUN and its containing volume from the **Configure FC/FCoE protocol** page:
 - a. Select the **Configure Data LIFs for FC** check box.
 - b. Enter **2** in the **LIFs per node** field.

Two LIFs are required for each node to ensure availability and data mobility.

- c. In the **Provision a LUN for FCP storage** area, enter the desired LUN size, host type, and WWPNs of the host initiators.
- d. Select **Submit & Continue**.

Configure FC/FCoE protocol

[?](#) Configure LIFs to access the data using FC/FCoE protocol

Data Interface (LIF) Configuration

Both FC and FCoE enabled hardware found. Click on the appropriate checkbox to configure the FC and/or FCoE LIFs.

Configure Data LIFs for FC

LIFs per node:	<input type="text" value="2"/> (Minimum: 1, Maximum: 2)	Provision a Lun for FCP storage (Optional)
		Lun Size: <input type="text"/> GB
		LUN OS Type: <input type="text" value="Linux"/>
		Host Initiator: <input type="text"/>

Review or Edit the Interface Association

Configure Data LIFs for FCoE

7. When the **SVM Administration** appears, configure or defer configuring a separate administrator for this SVM:
 - Select **Skip** and configure an administrator later if desired.
 - Enter the requested information, and then select **Submit & Continue**.
8. Review the **Summary** page, record the LIF information, and then select **OK**.

Zone the FC switches by the host and LIF WWPNs

Zoning the FC switches enables the hosts to connect to the storage and limits the number of paths. You zone the switches using the management interface of the switches.

Before you begin

- You must have administrator credentials for the switches.
- You must know the WWPN of each host initiator port and of each FC LIF for the storage virtual machine (SVM) in which you created the LUN.

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.

LUNs are mapped to a subset of the initiators in the igrup to limit the number of paths from the host to the LUN.

- By default, ONTAP 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 zone all of the FC 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.

The following illustration shows a host connected to a four-node cluster. There are two zones, one zone indicated by the solid lines and one zone indicated by the dashed lines. Each zone contains one initiator from the host and a LIF from each storage node.

You must use the WWPNs of the target LIFs, not the WWPNs of the physical FC ports on the storage nodes. The LIF WWPNs are all in the range `2x:xx:00:a0:98:xx:xx:xx`, where x is any hexadecimal digit. The physical port WWPNs are all in the range `50:0a:09:8x:xx:xx:xx:xx`.

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 initiator and all of the FC LIFs that connect to the same FC switch as the initiator.
3. Create additional zones for each FC initiator in the host.
4. Save the zones, and then activate the new zoning configuration.

Discover 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
```

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
```

Configure 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.

Verify 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
```

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.

What to do next

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

Copyright information

Copyright © 2025 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

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

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

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

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

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

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

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.