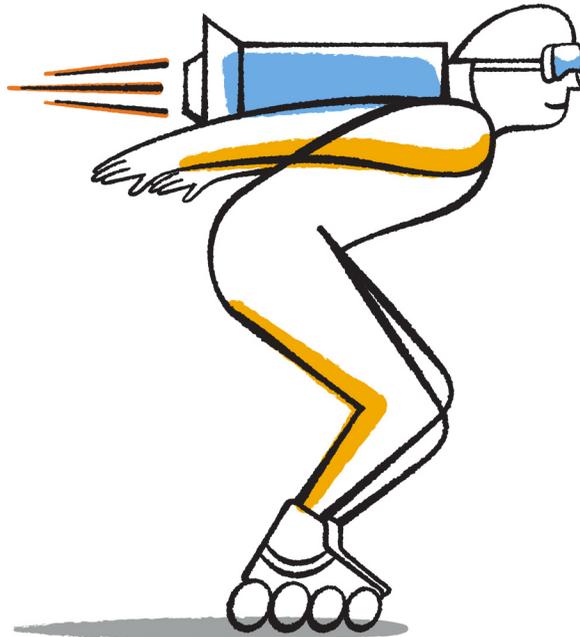




NetApp®

Clustered Data ONTAP® 8.3

FC Configuration for Red Hat® 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: 215-09654_A0
January 2015

Contents

Deciding whether to use this guide	4
FC configuration and provisioning workflow	5
Verifying that the FC configuration is supported	6
Completing the FC configuration worksheet	7
Installing the HBA utility from the HBA vendor	9
Updating the HBA driver, firmware, and BIOS	9
Installing Linux Unified Host Utilities and optimizing I/O performance	10
Recording the WWPN for each host FC port	10
Configuring DM-Multipath	12
Creating an aggregate	14
Deciding where to provision the volume	15
Verifying that the FC service is running on an existing SVM	15
Creating a LUN and its containing volume	16
Configuring FC on an existing SVM	18
Creating a new SVM	19
Zoning the FC switches by the host and LIF WWPNs	22
Discovering new SCSI devices (LUNs) and multipath devices	23
Configuring logical volumes on multipath devices and creating a file system	24
Verifying that the host can write to and read from a multipath device	25
Where to find additional information	27
Copyright information	28
Trademark information	29
How to send comments about documentation and receive update notification	30
Index	31

Deciding whether to use this guide

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

This guide is based on the following assumptions:

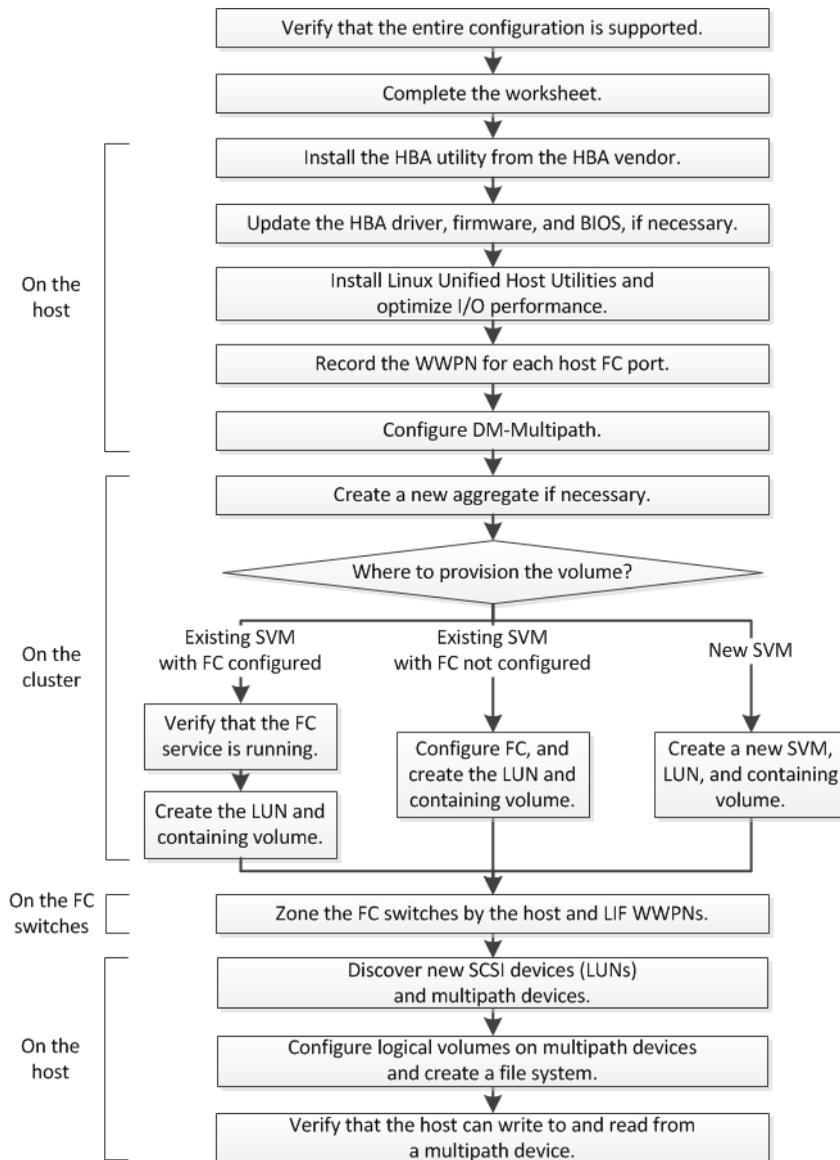
- You want to use best practices, not explore every available option.
- You do not want to read a lot of conceptual background.
- You want to use OnCommand System Manager, not the Data ONTAP command-line interface or an automated scripting tool.
- The host is running a supported version of Red Hat Enterprise Linux 6.4 or later.
- You are using traditional FC HBAs and switches.
This guide does not cover FCoE.
- 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 Data ONTAP CLI and is not covered in this guide.
- You are not configuring FC 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.

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.



Verifying that the FC configuration is supported

To ensure reliable operation, you must verify that the entire FC 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)
 - 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. 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 FC configuration worksheet

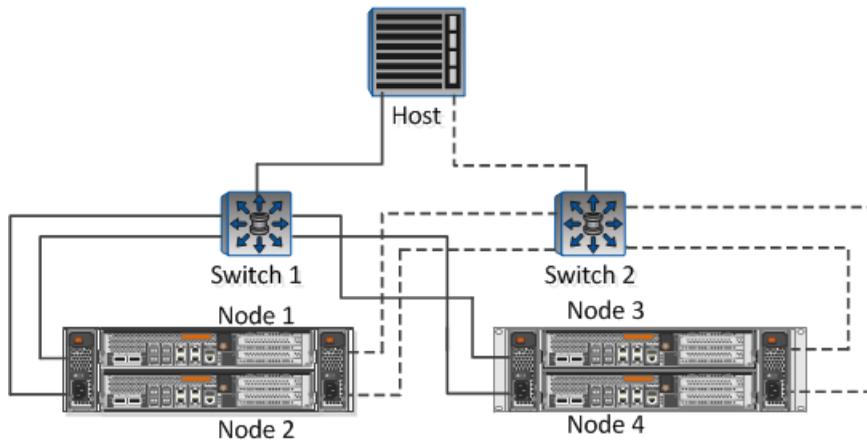
You require FC initiator and target WWPNS and storage configuration information to perform FC configuration tasks.

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 Data 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	

LIF	WWPN
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 user name (optional)	
SVM password (optional)	
SVM management LIF (optional)	Subnet:
	IP address:
	Network mask:
	Gateway:
	Home node:
	Home port:

Installing 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 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

[Emulex Downloads and Documentation](#)

[QLogic: NetApp Downloads](#)

Updating 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 *[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 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.

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 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 software following the instructions in the Linux Unified Host Utilities installation documentation.

3. For a physical host, ensure that you have set the `enterprise-storage` profile:

```
tuned-adm profile enterprise-storage
```

4. 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 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

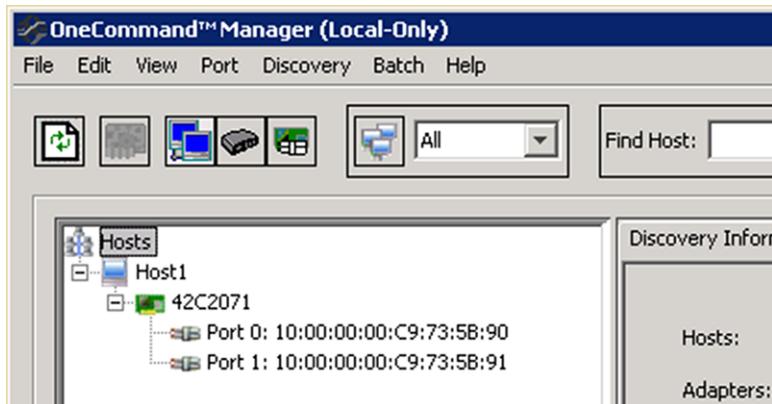
The WWPN is used for all configuration. You do not have to record the worldwide node name (WWNN).

Steps

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

Example

The following example shows Emulex 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.

Example

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

```
# sanlun fcp show adapter -v

adapter name:      host0
WWPN:           10000000c9813a14
WWNN:             20000000c9813a14
driver name:      lpfc
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:        lpfc
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

```

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

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

- a. 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.
- b. Enable NetApp recommended DM-Multipath settings by appending the following string to the end of the kernel line in the bootloader:

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

Example

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
```

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**.

Create Aggregate

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

Name:

Disk Type:

Number of Disks: Max: 8 (excluding 1 hot spare), min: 5 for RAID-DP

RAID Configuration: RAID-DP, RAID group size of 16 disks [Change](#)

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 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.

Choices

- If you want to provision volumes on an existing SVM that is already configured for FC, see [Verifying that the FC service is running on an existing SVM](#) on page 15.
- If you want to provision volumes on an existing SVM that has FC enabled but not configured, see [Configuring FC on an existing SVM](#) on page 18.

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 19.

Verifying that the FC service is running on an existing SVM

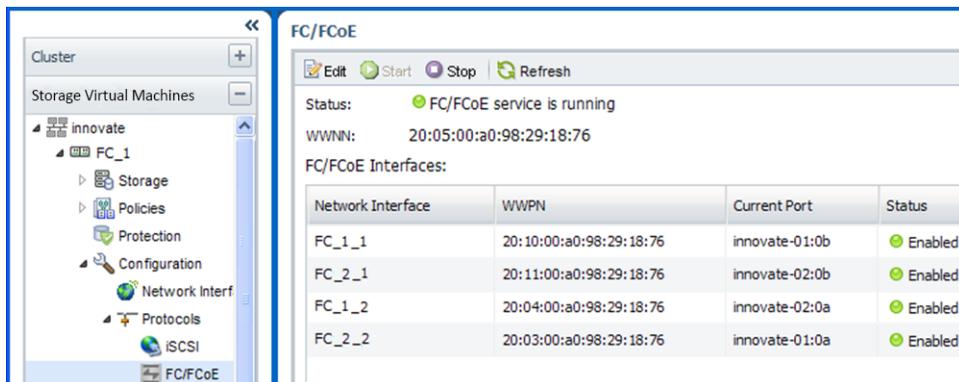
If you choose to use an existing Storage Virtual Machine (SVM), you must verify that the FC service is running on the SVM. You must also verify that FC LIFs are already created.

Before you begin

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

Steps

1. 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 then click **Configuration > Protocols > FC/FCoE**.
4. Verify that the FC service is running.



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

5. 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.

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 FC protocol enabled and the appropriate LIFs created.
- You must have recorded the 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. 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 then click **Storage > LUNs**.

4. In the **LUN Management** tab, click **Create**, and then type or select information as prompted by the wizard.
5. 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

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

Create LUN Wizard

LUN Container
You can let the wizard create a volume or you can choose an existing v...

The wizard automatically chooses the aggregate with most free space for creating LUN. But you can choose a different aggregate of your choice. You can also select a volume/qtree to create your LUN.

Create a new flexible volume in

Aggregate Name:

Volume Name:

7. On the **Initiators Mapping** page, click **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.
8. Confirm the details, and then click **Finish** to complete the wizard.

Related information

Clustered Data ONTAP 8.3 System Administration Guide for Cluster Administrators

Configuring 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 topic is intended for SVMs for which you are configuring multiple protocols, but have not yet configured FC.

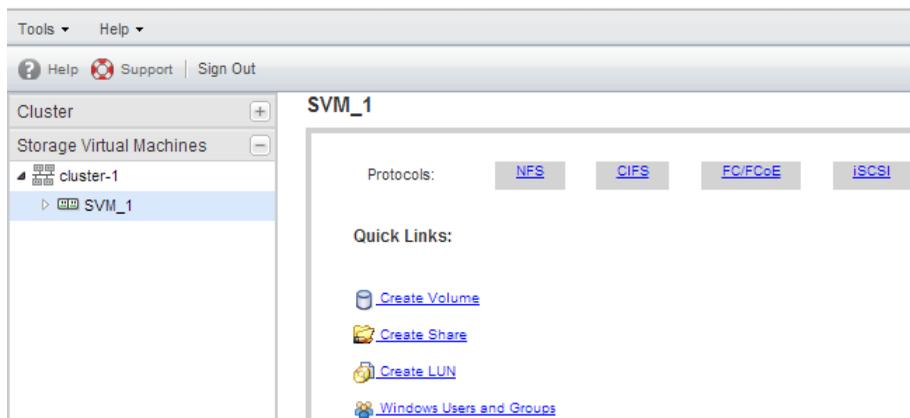
Before you begin

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

Steps

1. In the navigation pane, expand the **Storage Virtual Machines** hierarchy, and select the SVM.
2. In the SVM **Details** 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.



3. Click the **FC/FCoE** protocol link with the gray background.
The Configure FC/FCoE Protocol window is displayed.
4. 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. Click **Submit & Continue**.

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 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 a new SVM to separate data or administration from other users of the storage cluster.

Before you begin

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

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. If the **Configure iSCSI protocol** page appears because you enabled iSCSI, click **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. Click **Submit & Continue**.

Example

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](#)

Zoning 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

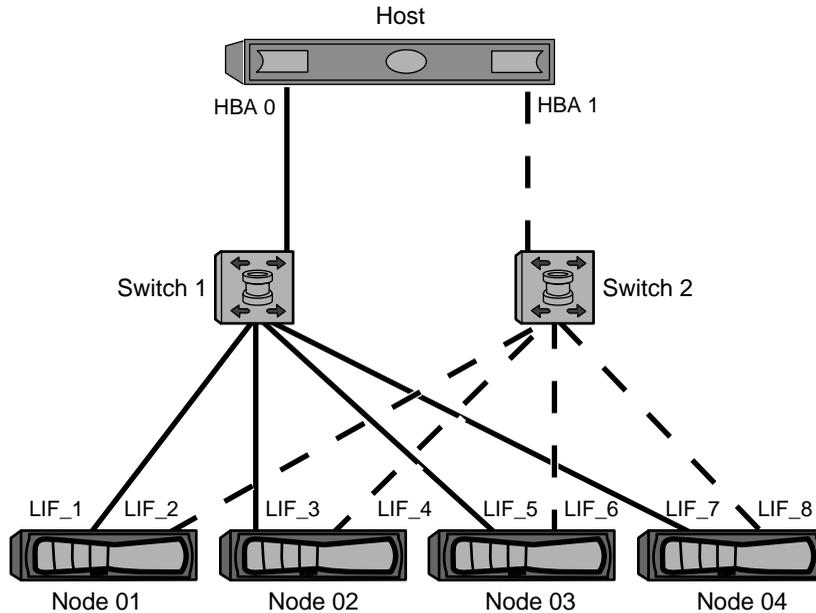
For details about zoning your switches, see the switch vendor's documentation.

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.

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 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.

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 FC service is running and re-check the DM-Multipath configuration and the FC paths to the multipath device.

Where to find additional information

After you have successfully tested FC, you can perform advanced FC 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, 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, SANtricity, SecureShare, Simplicity, Simulate ONTAP, Snap Creator, SnapCopy, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, StorageGRID, Tech OnTap, Unbound Cloud, and WAFL are trademarks or registered trademarks of NetApp, Inc., in the United States, and/or other countries. A current list of NetApp trademarks is available on the web at <http://www.netapp.com/us/legal/netapptmlist.aspx>.

Cisco and the Cisco logo are trademarks of Cisco in the U.S. and other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

How to send comments about documentation and receive update notification

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
 - FC configuration [27](#)
 - aggregates
 - creating [14](#)
 - audience
 - for this guide [4](#)
 - automatically starting
 - DM-Multipath at system boot [12](#)
- B**
- BIOS
 - updating FC HBA, in Linux hosts [9](#)
 - blacklisting
 - excluding devices in multipath.conf [12](#)
- C**
- comments
 - how to send feedback about documentation [30](#)
 - configuration
 - FC workflow [5](#)
 - where to find information about FC [27](#)
 - configuring
 - DM-Multipath [12](#)
 - FC on existing SVMs [18](#)
 - FC worksheet for [7](#)
 - creating
 - a file system [24](#)
 - aggregates [14](#)
 - logical volumes on multipath devices [24](#)
 - LUNs and their containing volumes [16](#)
 - SVMs [19](#)
- D**
- devices
 - excluding (blacklisting) in multipath.conf [12](#)
 - discovering
 - new SCSI devices and multipath devices [23](#)
 - DM-Multipath
 - configuring on Red Hat Linux hosts [12](#)
 - documentation
 - how to receive automatic notification of changes to [30](#)
 - how to send feedback about [30](#)
 - drivers
 - updating FC HBA in Linux hosts [9](#)
- E**
- editing
 - multipath.conf [12](#)
 - Emulex HBAs
 - recording the WWPN with OneCommand Manager [10](#)
 - excluding
 - devices in multipath.conf [12](#)
 - express guides
 - requirements for using FC with Red Hat Linux [4](#)
- F**
- FC
 - configuration workflow [5](#)
 - configuration worksheet for [7](#)
 - configuring on existing SVMs [18](#)
 - requirements for using FC with Red Hat Linux hosts [4](#)
 - updating HBA driver, firmware, and BIOS in Linux hosts [9](#)
 - verifying that configuration is supported [6](#)
 - where to find additional information about configuring [27](#)
 - FC ports
 - recording the WWPN [10](#)
 - FC service
 - verifying it is running on SVM [15](#)
 - FC switches
 - zoning by host and LIF WWPNs [22](#)
 - FC targets
 - creating as part of configuring SVMs [18](#)
 - creating as part of creating SVMs [19](#)
 - feedback
 - how to send comments about documentation [30](#)
 - Fibre Channel
 - See FC
 - file system

- creating on new multipath devices [24](#)
- firmware
 - updating FC HBA, in Linux hosts [9](#)
- FlexVol volumes
 - creating [16](#)
- flowcharts
 - FC configuration [5](#)
- formatting
 - a new SCSI device (LUN) [24](#)

G

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

H

- HBA utilities
 - installing [9](#)
- HBAs
 - recording the WWPN [10](#)
 - updating driver, firmware, and BIOS of FC, in Linux hosts [9](#)
- hosts
 - installing host utilities on [10](#)
 - recording the WWPN of FC ports [10](#)
 - setting the server profile [10](#)
 - verifying writing to and reading from multipath devices [25](#)

I

- igroups
 - creating and mapping the LUN to [16](#)
- information
 - how to send feedback about improving documentation [30](#)
 - where to find additional FC configuration [27](#)
- installing
 - host utilities [10](#)
- Interoperability Matrix
 - verifying supported FC configurations using [6](#)

L

- LIFs
 - creating as part of configuring SVMs [18](#)
 - creating as part of creating SVMs [19](#)
 - zoning FC switches by WWPNs of host and [22](#)

- Linux hosts, Red Hat
 - configuring DM-Multipath on [12](#)
 - packages required for DM-Multipath [6](#), [23](#)
 - updating FC HBA driver, firmware, and BIOS in [9](#)
 - verifying supported FC configurations [6](#)
- logical interfaces
 - See* LIFs
- LUNs
 - creating [16](#)
 - deciding where to provision the containing volume on a new or existing SVM [15](#)
 - discovering and creating multipath devices on [23](#)
 - See also* SCSI devices (LUNs)
 - discovering new [23](#)
 - ensuring access to with DM-Multipath [12](#)
 - how they appear to a Linux host [23](#)
 - mapping to an igroup [16](#)
 - partitioning and creating a file system on [24](#)
 - requirements for using FC with Red Hat Linux hosts [4](#)

M

- mapping
 - LUNs to an igroup [16](#)
- multipath devices
 - creating logical volumes and a file system on [24](#)
 - discovering new [23](#)
 - verifying that the host can write to and read from [25](#)
- multipath.conf
 - editing [12](#)

O

- OneCommand Manager
 - installing for Emulex HBAs [9](#)
 - recording the WWPN for Emulex HBAs [10](#)

P

- packages
 - required for DM-Multipath [6](#), [23](#)
 - required to set the server profile [10](#)
- partitions
 - creating for new SCSI devices (LUN) [24](#)
- provisioning
 - FC workflow [5](#)
 - FC worksheet for [7](#)
 - where to find information about FC [27](#)

- Q**
- QConvergeConsole
 - installing for QLogic HBAs [9](#)
 - recording the WWPN for QLogic HBAs [10](#)
 - QLogic HBAs
 - recording the WWPN with QConvergeConsole [10](#)
- R**
- read/write
 - verifying that the host can write to and read from multipath devices [25](#)
 - Red Hat hosts
 - See* Linux hosts, Red Hat
 - requirements
 - for using FC with Red Hat hosts [4](#)
 - verifying supported FC configurations [6](#)
- S**
- sanlun utility
 - using to record the WWPN [10](#)
 - SCSI devices (LUNs)
 - discovering and creating multipath devices on [23](#)
 - See also* LUNs
 - discovering new [23](#)
 - partitioning and creating a file system on [24](#)
 - server profile
 - setting [10](#)
 - starting
 - DM-Multipath [12](#)
 - storage configurations
 - FC configuration worksheet [7](#)
 - suggestions
 - how to send feedback about documentation [30](#)
 - supported configurations
 - verifying FC [6](#)
 - SVMs
 - configuring FC on existing [18](#)
 - creating [19](#)
 - deciding whether to provision a volume on a new or existing [15](#)
 - verifying FC service is running [15](#)
 - switches
 - zoning FC, by host and LIF WWPNs [22](#)
- creating FC as part of configuring SVMs [18](#)
 creating FC as part of creating SVMs [19](#)
- twitter
 - how to receive automatic notification of documentation changes [30](#)
- U**
- updates, Linux
 - See* packages
 - updating
 - FC HBA driver, firmware, and BIOS in Linux hosts [9](#)
 - utilities
 - installing HBA [9](#)
 - sanlun [10](#)
- V**
- volumes
 - creating [16](#)
 - deciding whether to provision on a new or existing SVM [15](#)
- W**
- wizards
 - running the Create LUN [16](#)
 - workflows
 - FC configuration [5](#)
 - worksheets
 - FC configuration [7](#)
 - write/read
 - verifying that the host can write to and read from multipath devices [25](#)
 - WWPNs
 - FC configuration worksheet for [7](#)
 - recording [10](#)
 - zoning FC switches by host and LIF [22](#)
- Z**
- zoning
 - FC switches, by host and LIF WWPNs [22](#)

- T**
- targets