



Clustered Data ONTAP® 8.3

SAN Administration Guide

February 2016 | 215-10555_B0
doccomments@netapp.com

Updated for 8.3.2

 **NetApp®**

Contents

Using space management capabilities to maximize storage utilization and availability for LUNs	7
How Data ONTAP provides the ability to use thin provisioning	7
What a thin-provisioned volume is	7
What a thin-provisioned LUN is	7
What it means to overcommit a storage object	8
Volume provisioning options	8
Recommended volume and LUN configuration combinations	9
Configuration settings for space-reserved LUNs with thick-provisioned volumes	10
Configuration settings for non-space-reserved LUNs with thin-provisioned volumes	11
Configuration settings for space-reserved LUNs with semi-thick volume provisioning	11
Determining the correct volume and LUN configuration combination for your environment	12
SAN volume configuration	14
SAN volume configuration options	14
Requirement for moving volumes in SAN environments	15
Considerations for setting fractional reserve	15
Calculating rate of data growth for LUNs	16
LUN setup workflow	18
Setting up your LUN	19
Setting up LUNs	19
Configuring switches for FCoE	19
Prerequisites for setting up LUNs	20
Verifying the license for FC or iSCSI	20
Enabling block access for SVMs with iSCSI	21
Enabling block access for SVMs with FC	21
Creating LUNs and mapping igroups	22
Enabling block access for a specific host	23
Creating port sets and binding igroups to port sets	24
LUN guidelines for SAN environments	25
Considerations when creating a LUN from an existing file	25
Guidelines for assigning LUN IDs	25
Guidelines for mapping LUNs to igroups	25
Managing LUNs	26
Selective LUN Map	26
How to determine whether SLM is enabled on a LUN map	27
Decreasing mapped LUN paths with SLM for LUNs created prior to Data ONTAP 8.3	27

Modifying SLM reporting nodes	28
Ways to limit LUN access with port sets and igroups	28
Considerations for transitioning SAN configurations	29
Capabilities and restrictions of transitioned LUNs	29
Considerations for copying LUNs	30
Increasing the size of a LUN	30
Decreasing the size of a LUN	31
Moving LUNs	31
Deleting LUNs	32
Examining configured and used space of a LUN	33
Commands for managing LUNs	33
Commands for managing port sets	34
I/O misalignments might occur on properly aligned LUNs	35
How to achieve I/O alignment using LUN OS types	36
Special I/O alignment considerations for Linux	37
Special I/O alignment considerations for Solaris LUNs	37
ESX boot LUNs report as misaligned	37
Controlling and monitoring I/O performance to LUNs by using Storage QoS	37
Tools available to effectively monitor your LUNs	38
Ways to address issues when LUNs go offline	38
What host-side space management is	39
Automatic host-side space management with SCSI thinly provisioned LUNs	40
Enabling space allocation for SCSI thinly provisioned LUNs	40
Host support for SCSI thin provisioning	41
Managing igroups	42
Commands for managing igroups	42
What igroups are	42
Example of how igroups give LUN access	43
How to specify initiator WWPNs and node names for an igroup	43
Why Data ONTAP uses ALUA	44
MPIO and ALUA	44
Managing LIFs	45
Considerations for SAN LIF migration	45
Removing a SAN LIF from a port set	45
Moving SAN LIFs	46
Deleting a LIF in a SAN environment	47
FC and FCoE LIFs on the same port need to be in separate zones	47
Considerations for adding nodes to a cluster	48
Data protection methods in SAN environments	49
Effect of moving or copying a LUN on Snapshot copies	50
Restoring a single LUN from a Snapshot copy	50
Restoring all LUNs in a volume from a Snapshot copy	51
Deleting one or more existing Snapshot copies in a volume	52
Using FlexClone LUNs to protect your data	52
Reasons for using FlexClone LUNs	53

How a FlexVol volume can reclaim free space from FlexClone files and FlexClone LUNs	53
Configuring a FlexVol volume to automatically delete FlexClone files and FlexClone LUNs	54
Cloning LUNs from an active volume	56
Creating FlexClone LUNs from a Snapshot copy in a volume	57
Preventing a specific FlexClone file or FlexClone LUN from being automatically deleted	57
Configuring and using SnapVault backups in a SAN environment	58
Accessing a read-only LUN copy from a SnapVault backup	59
Restoring a single LUN from a SnapVault backup	60
Restoring all LUNs in a volume from a SnapVault backup	62
How you can connect a host backup system to the primary storage system	64
Backing up a LUN through a host backup system	65
Ways to implement SVM disaster recovery in SAN environments	67
Managing your iSCSI service	68
Commands for managing iSCSI services	68
Configuring your network for best performance	71
Defining a security policy method for an initiator	71
Deleting an iSCSI service for an SVM	72
Getting more details in iSCSI session error recoveries	72
iSCSI service management	73
How iSCSI is implemented on the host	73
How iSCSI authentication works	73
iSNS server registration requirement	75
What iSNS is	75
What an iSNS server does	75
How SVMs interact with an iSNS server	75
About iSNS service version incompatibility	77
Registering the SVM with an iSNS server	77
Commands for managing iSNS	78
iSCSI troubleshooting tips	79
Troubleshooting iSCSI LUNs not visible on the host	79
Resolving iSCSI error messages on the storage system	80
Managing your FC service	81
Commands for managing FC protocols	81
Changing the WWPN for an FC logical interface	82
Deleting an FC service for an SVM	83
How worldwide name assignments work	83
How FC switches are identified	84
Recommended MTU configurations for FCoE jumbo frames	84
Managing systems with FC adapters	84
Commands for managing FC adapters	84
Configuring FC adapters for initiator mode	85
Configuring FC adapters for target mode	86

Displaying information about an FC target adapter	87
Changing the FC adapter speed	87
Supported port configurations for X1143A-R6 adapters	88
How to prevent loss of connectivity when using the X1133A-R6 adapter ...	89
Storage virtualization with VMware and Microsoft copy offload	90
Advantages of using a virtualized SAN environment	90
How LUN access works in a virtualized environment	90
Considerations for LIFs in cluster SAN environments	92
Improving VMware VAAI performance for ESX hosts	93
Requirements for using the VAAI environment	93
How to determine if VAAI features are supported by ESX	94
Microsoft Offloaded Data Transfer (ODX)	94
Requirements for using ODX	94
Use cases for ODX	95
Special system file requirements	96
Basic iSCSI and FC concepts	97
Protocols that hosts can use to connect to SAN storage systems	97
What Host Utilities are	97
Simplified host management with SnapDrive	97
How Data ONTAP implements an iSCSI network	98
What iSCSI is	98
What iSCSI nodes are	98
How iSCSI target nodes connect to the network	98
How iSCSI nodes are identified	98
Storage system node name	99
How the storage system checks initiator node names	100
Default port for iSCSI	100
How iSCSI communication sessions work	100
How high availability is maintained in an iSCSI SVM	100
How Data ONTAP implements an FC SAN	100
What FC is	100
What FC nodes are	100
How FC target nodes connect to the network	101
How FC nodes are identified	101
How WWPNs are used	101
Copyright information	102
Trademark information	103
How to send comments about documentation and receive update	
notifications	104
Index	105

Using space management capabilities to maximize storage utilization and availability for LUNs

Volumes and LUNs provide a set of configuration options to determine how they use and provide space. You must choose the right combination of LUN and volume configurations for your environment.

How Data ONTAP provides the ability to use thin provisioning

How you configure a storage object to use thin provisioning and how that object behaves as a result depends on the type of storage object. You can create thinly provisioned volumes and LUNs.

What a thin-provisioned volume is

A thin-provisioned volume is a volume for which storage is not set aside up-front. Instead, the storage for the volume is allocated as it is needed.

You create a thin-provisioned FlexVol volume by setting its guarantee to **none**. With a guarantee of **none**, the volume size is not limited by the aggregate size. In fact, each volume could, if required, be larger than the containing aggregate. The storage provided by the aggregate is used only as data is written to the volume.

What a thin-provisioned LUN is

The definition of a “thin-provisioned LUN” changes depending on the context. The T10 SCSI (SAN) standard uses one definition, and historically, NetApp has used a different definition.

SCSI thin-provisioned LUNs

The T10 SCSI standard defines two types of LUNs: thin provisioned and fully provisioned. Data ONTAP supports both types of T10 standard LUNs.

SCSI thin provisioning (sometimes called *T10 thin provisioning*) is a set of SCSI features and capabilities enabled by Data ONTAP. These SCSI features must be supported by your SCSI host software. SCSI thin provisioning enables host applications to support SCSI features, including LUN space reclamation and LUN space monitoring capabilities for blocks environments.

If your host software supports SCSI thin provisioning, you can use it with space-reserved and non-space-reserved LUNs, and with any volume provisioning type. See the documentation for your host software for more information about the SCSI thin provisioning capabilities it provides.

You use the Data ONTAP **space-allocation** setting to enable SCSI thin provisioning on a LUN.

NetApp thin-provisioned (non-space-reserved) LUNs

Historically, NetApp has used the term “thin-provisioned LUN” to mean a LUN for which space reservation is disabled (a non-space-reserved LUN). A non-space-reserved LUN shares some important characteristics with a thinly provisioned volume: its storage is allocated as it is used rather than at creation time, and its containing storage object can be overcommitted. The term “non-space-reserved LUN” is used for this configuration.

You use the Data ONTAP **space-reserve** setting to configure space reservation on a LUN.

What it means to overcommit a storage object

When you do not allocate the storage for a FlexVol volume or LUN up front, it enables you to overcommit the storage object that supplies its storage. Overcommitting a storage object can increase your storage efficiency but it also requires that you take an active role in monitoring your free space to prevent writes from failing due to lack of space.

A storage object is said to be overcommitted if the objects it supplies storage to are collectively larger than the amount of physical storage it can currently supply. For example, suppose you have three FlexVol volumes with volume guarantees of **none** associated with a 100 TB aggregate. If each of the FlexVol volumes has a nominal size of 40 TB, the aggregate is overcommitted. Each of the volumes could continue to receive data, providing that collectively, their physical storage requirement is not more than 100 TB. This is only possible if the volumes have a volume guarantee of **none** (they are thinly provisioned). If all of the volumes had a volume guarantee of **volume**, you could not create the third volume.

Similarly, you can overcommit a volume that contains more than one LUN, as long as the LUNs are not space-reserved.

The ability to overcommit the supplying storage object is the promise offered by not preallocating storage for an object up front, but overcommitment requires you to manage the supply of physical storage resources carefully to avoid running out of free space. This is true in any configuration, but if your supplying storage object is overcommitted, you can run out of free space by issuing a write to space that appears to have been allocated already.

Volume provisioning options

Data ONTAP provides three basic volume provisioning options: thick provisioning, thin provisioning, and semi-thick provisioning. Each option uses different ways to manage the volume space and the space requirements for Data ONTAP block sharing technologies. Understanding how the options work enables you to choose the best option for your environment.

Thick provisioning for volumes

When a thick-provisioned volume is created, Data ONTAP sets aside enough storage from the aggregate to ensure that any block in the volume can be written to at any time. When you configure a volume to use thick provisioning, you can employ any of the Data ONTAP storage efficiency capabilities, such as compression and deduplication, to offset the larger upfront storage requirements.

Thin provisioning for volumes

When a thinly provisioned volume is created, Data ONTAP does not reserve any extra space when the volume is created. As data is written to the volume, the volume requests the storage it needs from the aggregate to accommodate the write operation. Using thin-provisioned volumes enables you to overcommit your aggregate, which introduces the possibility of the volume not being able to secure the space it needs when the aggregate runs out of free space.

Semi-thick provisioning for volumes

When a volume using semi-thick provisioning is created, Data ONTAP sets aside storage space from the aggregate to account for the volume size. If the volume is running out of free space because blocks are in use by block-sharing technologies, Data ONTAP makes an effort to delete protection data objects (Snapshot copies and FlexClone LUNs) to free up the space they are holding. As long as Data ONTAP can delete the protection data objects fast enough to keep pace with the space required for overwrites, the write operations continue to succeed. This is called a “best effort” write guarantee.

You cannot employ storage efficiency technologies such as deduplication and compression on a volume that is using semi-thick provisioning.

The following table summarizes the major differences in how the three volume provisioning options can be used:

Volume provisioning	LUN space reservation	Overwrites	Protection data²	Storage efficiency³
Thick	Supported	Guaranteed ¹	Guaranteed	Supported
Thin	No effect	None	Guaranteed	Supported
Semi-thick	Supported	Best effort ¹	Best effort	Not supported

Notes

1. The ability to guarantee overwrites or provide a best-effort overwrite assurance requires that space reservation is enabled on the LUN.
2. Protection data includes Snapshot copies, and FlexClone LUNs marked for automatic deletion (backup clones).
3. Storage efficiency includes deduplication, compression, any FlexClone LUNs not marked for automatic deletion (active clones), and FlexClone subfiles (used for Copy Offload).

Recommended volume and LUN configuration combinations

There are specific combinations of FlexVol volume and LUN configurations you can use, depending on your application and administration requirements. Understanding the benefits and costs of these combinations can help you determine the right volume and LUN configuration combination for your environment.

The following volume and LUN configuration combinations are recommended:

- Space-reserved LUNs with thick volume provisioning
- Non-space-reserved LUNs with thin volume provisioning
- Space-reserved LUNs with semi-thick volume provisioning

You can use SCSI thin provisioning on your LUNs in conjunction with any of these configuration combinations.

Space-reserved LUNs with thick volume provisioning

Benefits:

- All write operations are guaranteed; they will not fail due to insufficient space.
- There are no restrictions on storage efficiency and data protection technologies on the volume.

Costs and limitations:

- Enough space must be set aside from the aggregate up front to support the thickly provisioned volume.
- Space equal to twice the size of the LUN is allocated from the volume at LUN creation time.

Non-space-reserved LUNs with thin volume provisioning

Benefits:

- There are no restrictions on storage efficiency and data protection technologies on the volume.
- Space is allocated only as it is used.

Costs and restrictions:

- Write operations are not guaranteed; they can fail if the volume runs out of free space.
- You must manage the free space in the aggregate effectively to prevent the aggregate from running out of free space.

Space-reserved LUNs with semi-thick volume provisioning**Benefits:**

Less space is reserved up front than for thick volume provisioning, and a best-effort write guarantee is still provided.

Costs and restrictions:

- Write operations can fail with this option.
You can mitigate this risk by properly balancing free space in the volume against data volatility.
- You cannot rely on retention of data protection objects such as Snapshot copies and FlexClone files and LUNs.
- You cannot use Data ONTAP block-sharing storage efficiency capabilities that cannot be automatically deleted, including deduplication, compression, and ODX/Copy Offload.

Configuration settings for space-reserved LUNs with thick-provisioned volumes

This FlexVol volume and LUN configuration combination provides the ability to use storage efficiency technologies and does not require you to actively monitor your free space, because sufficient space is allocated up front.

The following settings are required to configure a space-reserved LUN in a volume using thick provisioning:

Volume setting	Value
Guarantee	Volume
Fractional reserve	100
Snapshot reserve	Any
Snapshot autodelete	Optional
Autogrow	Optional; if enabled, aggregate free space must be actively monitored.

LUN setting	Value
Space reservation	Enabled

Technology restrictions

None

Additional considerations

None

Configuration settings for non-space-reserved LUNs with thin-provisioned volumes

This FlexVol volume and LUN configuration combination requires the smallest amount of storage to be allocated up front, but requires active free space management to prevent errors due to lack of space.

The following settings are required to configure a non-space-reserved LUN in a thin-provisioned volume:

Volume setting	Value
Guarantee	None
Fractional reserve	0
Snapshot reserve	Any
Snapshot autodelete	Optional
Autogrow	Optional

LUN setting	Value
Space reservation	Disabled

Technology restrictions

None

Additional considerations

When the volume or aggregate runs out of space, write operations to the LUN can fail.

If you do not want to actively monitor free space for both the volume and the aggregate, you should enable Autogrow for the volume and set the maximum size for the volume to the size of the aggregate. In this configuration, you must monitor aggregate free space actively, but you do not need to monitor the free space in the volume.

Configuration settings for space-reserved LUNs with semi-thick volume provisioning

This FlexVol volume and LUN configuration combination requires less storage to be allocated up front than the fully provisioned combination, but places restrictions on the efficiency technologies you can use for the volume. Overwrites are fulfilled on a best-effort basis for this configuration combination.

The following settings are required to configure a space-reserved LUN in a volume using semi-thick provisioning:

Volume setting	Value
Guarantee	Volume
Fractional reserve	0
Snapshot reserve	0
Snapshot autodelete	On, with a commitment level of destroy, a destroy list that includes all objects, the trigger set to volume, and all FlexClone LUNs enabled for automatic deletion.

Volume setting	Value
Autogrow	Optional; if enabled, aggregate free space must be actively monitored.

LUN setting	Value
Space reservation	Enabled

Technology restrictions

You cannot use the following volume storage efficiency technologies for this configuration combination:

- Compression
- Deduplication
- ODX and FlexClone Copy Offload
- FlexClone LUNs not marked for automatic deletion (active clones)
- FlexClone subfiles
- ODX/Copy Offload

Additional considerations

The following facts must be considered when employing this configuration combination:

- When the volume that supports that LUN runs low on space, protection data (FlexClone LUNs, Snapshot copies) is destroyed.
- Write operations can time out and fail when the volume runs out of free space.

Compression is enabled by default for All Flash FAS platforms. You must explicitly disable compression for any volume for which you want to use semi-thick provisioning on an All Flash FAS platform.

Determining the correct volume and LUN configuration combination for your environment

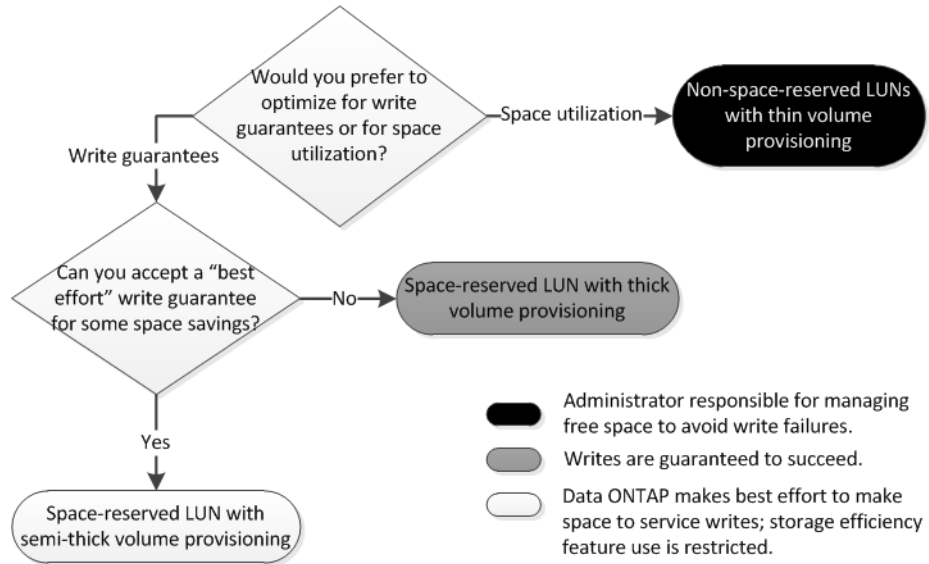
Answering a few basic questions about your environment can help you determine the best FlexVol volume and LUN configuration for your environment.

About this task

You can optimize your LUN and volume configurations for maximum storage utilization or for the security of write guarantees. You must choose which choice makes the most sense for your installation, based on your requirements for storage utilization and your ability to monitor and replenish free space quickly.

Step

1. Use the following decision tree to determine the best volume and LUN configuration combination for your environment:



SAN volume configuration

Volumes containing LUNs must be FlexVol volumes. SAN protocols can only be used with Storage Virtual Machines (SVMs) with FlexVol volumes. Infinite Volumes are not supported for SAN.

In this deliverable, “volume” always means “FlexVol volume” and “SVM” always means “SVM with FlexVol volumes”.

See the *Clustered Data ONTAP Logical Storage Management Guide* for more information about volumes.

SAN volume configuration options

You must set various options on the volume containing your LUN. The way you set the volume options determines the amount of space available to LUNs in the volume.

Autogrow

You can enable or disable Autogrow. If you enable it, autogrow allows Data ONTAP to automatically increase the size of the volume up to a maximum size that you predetermine. There must be space available in the containing aggregate to support the automatic growth of the volume. Therefore, if you enable autogrow, you must monitor the free space in the containing aggregate and add more when needed.

Autogrow cannot be triggered to support Snapshot creation. If you attempt to create a Snapshot copy and there is insufficient space on the volume, the Snapshot creation fails, even with autogrow enabled.

If autogrow is disabled, the size of your volume will remain the same.

Autoshrink

You can enable or disable Autoshrink. If you enable it, autoshrink allows Data ONTAP to automatically decrease the overall size of a volume when the amount of space consumed in the volume decreases a predetermined threshold. This increases storage efficiency by triggering volumes to automatically release unused free space.

Snapshot autodelete

Snapshot autodelete automatically deletes Snapshot copies when one of the following occurs:

- The volume is nearly full.
- The Snapshot reserve space is nearly full.
- The overwrite reserve space is full.

You can configure Snapshot autodelete to delete Snapshot copies from oldest to newest or from newest to oldest. Snapshot autodelete does not delete Snapshot copies that are linked to Snapshot copies in cloned volumes or LUNs.

If your volume needs additional space and you have enabled both autogrow and Snapshot autodelete, by default, Data ONTAP attempts to acquire the needed space by triggering autogrow first. If enough space is not acquired through autogrow, then Snapshot autodelete is triggered.

Snapshot reserve

Snapshot reserve defines the amount of space in the volume reserved for Snapshot copies. Space allocated to Snapshot reserve cannot be used for any other purpose. If all of the space allocated for Snapshot reserve is used, then Snapshot copies begin to consume additional space on the volume.

Related information

[Clustered Data ONTAP 8.3 Data Protection Guide](#)

Requirement for moving volumes in SAN environments

Before you move a volume that contains one or more LUNs, you should have a minimum of two paths per LUN (LIFs) connecting to each node in the cluster. This eliminates single points of failure and enables the system to survive component failures.

Related tasks

[Modifying SLM reporting nodes](#) on page 28

Considerations for setting fractional reserve

Fractional reserve, also called *LUN overwrite reserve*, enables you to turn off overwrite reserve for space-reserved LUNs in a FlexVol volume. This can help you maximize your storage utilization, but if your environment is negatively affected by write operations failing due to lack of space, you must understand the requirements that this configuration imposes.

The fractional reserve setting is expressed as a percentage; the only valid values are **0** and **100** percent. The fractional reserve setting is an attribute of the volume.

Setting fractional reserve to **0** increases your storage utilization. However, an application accessing data residing in the volume could experience a data outage if the volume is out of free space, even with the volume guarantee set to **volume**. With proper volume configuration and use, however, you can minimize the chance of writes failing. Data ONTAP provides a “best effort” write guarantee for volumes with fractional reserve set to **0** when *all* of the following requirements are met:

- Deduplication is not in use
- Compression is not in use
- All FlexClone LUNs are enabled for automatic deletion
This is not the default setting. You must explicitly enable automatic deletion, either at creation time or by modifying the FlexClone LUN after it is created.
- ODX and FlexClone copy offload are not in use
- Volume guarantee is set to **volume**
- LUN space reservation is **enabled**
- Volume Snapshot reserve is set to **0**
- Volume Snapshot copy automatic deletion is **enabled** with a commitment level of **destroy**, a destroy list of **lun_clone,vol_clone,cifs_share,file_clone,sfsr**, and a trigger of **volume**

This setting also ensures that FlexClone LUNs are deleted when necessary.

Note that if your rate of change is high, in rare cases the Snapshot copy automatic deletion could fall behind, resulting in the volume running out of space, even with all of the above required configuration settings in use.

In addition, you can optionally use the volume autogrow capability to decrease the likelihood of volume Snapshot copies needing to be deleted automatically. If you enable the autogrow capability, you must monitor the free space in the associated aggregate. If the aggregate becomes full enough that the volume is prevented from growing, more Snapshot copies will probably be deleted as the free space in the volume is depleted.

If you cannot meet all of the above configuration requirements and you need to ensure that the volume does not run out of space, you must set the volume's fractional reserve setting to **100**. This requires more free space up front, but guarantees that data modification operations will succeed even when the technologies listed above are in use.

The default value and allowed values for the fractional reserve setting depend on the guarantee of the volume:

Volume guarantee	Default fractional reserve	Allowed values
Volume	100	0, 100
None	0	0, 100

Calculating rate of data growth for LUNs

You need to know the rate at which your LUN data is growing over time to determine whether you should use space-reserved LUNs or non-space-reserved LUNs.

About this task

If you have a consistently high rate of data growth, then space-reserved LUNs might be a better option for you. If you have a low rate of data growth, then you should consider non-space-reserved LUNs.

You can use tools such as OnCommand Insight to calculate your rate of data growth or you can calculate it manually. The following steps are for manual calculation.

Steps

1. Set up a space-reserved LUN.
2. Monitor the data on the LUN for a set period of time, such as one week.

Make sure that your monitoring period is long enough to form a representative sample of regularly occurring increases in data growth. For instance, you might consistently have a large amount of data growth at the end of each month.
3. Each day, record in GB how much your data grows.
4. At the end of your monitoring period, add the totals for each day together, and then divide by the number of days in your monitoring period.

This calculation yields your average rate of growth.

In this example, you need a 200 GB LUN. You decide to monitor the LUN for a week and record the following daily data changes:

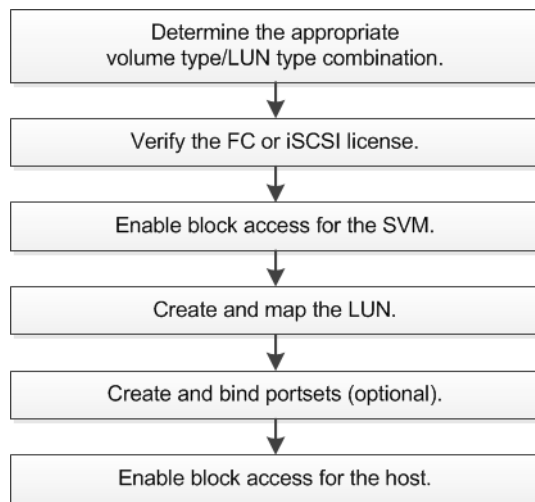
- Sunday: 20 GB
- Monday: 18 GB

- Tuesday: 17 GB
- Wednesday: 20 GB
- Thursday: 20 GB
- Friday: 23 GB
- Saturday: 22 GB

In this example, your rate of growth is $(20+18+17+20+20+23+22) / 7 = 20$ GB per day.

LUN setup workflow

To set up your LUN, you must determine the best combination of volume type and LUN type for your needs. Then you can follow a series of tasks to verify your protocol license, enable block access, create and map your LUN, and enable block access on your host. You can also optionally create and bind portsets as part of the LUN setup workflow.



Related tasks

[Determining the correct volume and LUN configuration combination for your environment](#) on page 12

[Verifying the license for FC or iSCSI](#) on page 20

[Enabling block access for SVMs with iSCSI](#) on page 21

[Enabling block access for SVMs with FC](#) on page 21

[Creating LUNs and mapping igroups](#) on page 22

[Enabling block access for a specific host](#) on page 23

[Creating port sets and binding igroups to port sets](#) on page 24

Setting up your LUN

After you have configured your volume for a specific LUN type, you must complete the steps necessary to set up your LUN. You should consider LUN setup guidelines and functions when setting up your LUNs.

Related concepts

[Recommended volume and LUN configuration combinations](#) on page 9

[Selective LUN Map](#) on page 26

Related tasks

[Determining the correct volume and LUN configuration combination for your environment](#) on page 12

[Creating LUNs and mapping igroups](#) on page 22

[Enabling block access for SVMs with iSCSI](#) on page 21

[Enabling block access for SVMs with FC](#) on page 21

[Verifying the license for FC or iSCSI](#) on page 20

[Enabling block access for a specific host](#) on page 23

[Creating port sets and binding igroups to port sets](#) on page 24

Setting up LUNs

You must complete several required tasks before you can access your LUNs including verifying your protocol license, enabling block access, creating your LUNs, and mapping your LUNs to igroups.

Configuring switches for FCoE

You must configure your switches for FCoE before your FC service can run over the existing Ethernet infrastructure.

Before you begin

- Your SAN configuration must be supported.
For more information about supported configurations, see the Interoperability Matrix.
- A Unified Target Adapter (UTA) must be installed on your storage system.
If you are using a UTA2, it must be set to `cna` mode.
- A converged network adapter (CNA) must be installed on your host.

Steps

1. Use your switch documentation to configure your switches for FCoE.
2. Use the `dcg show` command to verify that the DCB settings for each node in the cluster have been correctly configured.

Example

```
run -node node1 -command dcb show
```

DCB settings are configured on the switch. Consult your switch documentation if the settings are incorrect.

- Use the `fc adapter show` command to verify that the FCoE login is working when the FC target port online status is `true`.

Example

```
cluster1::> fc adapter show -fields
node,adapter,status,state,speed,fabric-established,physical-protocol
```

If the FC target port online status is `false`, consult your switch documentation.

Related information

[NetApp Interoperability Matrix Tool](#)

[NetApp Technical Report 3800: Fibre Channel over Ethernet \(FCoE\) End-to-End Deployment Guide](#)

[Cisco MDS 9000 NX-OS and SAN-OS Software Configuration Guides](#)

[Brocade products](#)

Prerequisites for setting up LUNs

Setting up LUNs involves creating the LUN, creating an igroup, and mapping the LUN to the igroup. Your system must meet certain prerequisites before you can set up your LUNs.

- The Interoperability Matrix must list your SAN configuration as supported.
- Your SAN environment must meet the SAN host and controller configuration limits specified in the *Clustered Data ONTAP SAN Configuration Guide* for your version of Data ONTAP software.
- A supported version of Host Utilities must be installed.
For more information, see the Host Utilities documentation.
- You must have SAN LIFs on the LUN owning node and the owning node's HA partner.
- If you are setting up an FC LUN, you should be familiar with the FC best practices reviewed in *TR-4017: Fibre Channel SAN Best Practices*.

Related information

[NetApp Interoperability Matrix Tool](#)

[Clustered Data ONTAP 8.3 SAN Configuration Guide](#)

[NetApp Documentation: Host Utilities \(current releases\)](#)

[NetApp Technical Report 4017: Fibre Channel SAN Best Practices](#)

Verifying the license for FC or iSCSI

Before you can enable block access for a Storage Virtual Machine (SVM) with FC or iSCSI, you must have a license.

Steps

- Use the `system license show` command to verify that you have a license for FC or iSCSI.

Example

```
system license show
```

Package	Type	Description	Expiration
Base	site	Cluster Base License	-
NFS	site	NFS License	-

CIFS	site	CIFS License	-
iSCSI	site	iSCSI License	-
FCP	site	FCP License	-

- If you do not have a license for FC or iSCSI, use the `license add` command.

Example

```
license add -license-code your_license_code
```

Enabling block access for SVMs with iSCSI

To enable block access, you must assign an iSCSI protocol to your Storage Virtual Machine (SVM) and create LIFs for that SVM.

About this task

You need a minimum of one iSCSI LIF per node for each SVM serving data with the iSCSI protocol. For redundancy, you should create at least two LIFs per node.

Steps

- Enable the SVMs to listen for iSCSI traffic:


```
vserver iscsi create -vserver vserver_name -target-alias vserver_name
```
- Create a LIF for the SVMs on each node serving iSCSI:


```
network interface create -vserver vserver_name -lif lif_name -role data
-data-protocol iscsi -home-node node_name -home-port port_name -address
ip_address -netmask netmask
```
- Create a LIF on each node's HA partner:


```
network interface create -vserver vserver_name -lif lif_name -role data
-data-protocol iscsi -home-node node_name -home-port port_name -address
ip_address -netmask netmask
```
- Verify that you set up your LIFs correctly:


```
network interface show -vserver vserver_name
```
- From your host, create iSCSI sessions to your LIFs.

Related concepts

[Considerations for LIFs in cluster SAN environments](#) on page 92

[What CHAP authentication is](#) on page 74

Related tasks

[Defining a security policy method for an initiator](#) on page 71

Enabling block access for SVMs with FC

To enable block access, you must create LIFs for a Storage Virtual Machine (SVM) and assign the FC protocol to those LIFs.

Before you begin

You must have an FC license and it must be enabled. If the FC license is not enabled, the LIFs and SVMs appear to be online but the operational status is `down`. The FC service must be enabled for your LIFs and SVMs to be operational.

About this task

You need a minimum of one FC LIF per node for each SVM serving data with the FC protocol. For redundancy, you should create at least two LIFs per node.

Steps

1. Enable FC service on the SVM:

```
vserver fcp create -vserver vserver_name -status-admin up
```

2. Create a LIF for the SVMs on each node serving FC:

```
network interface create -vserver vserver_name -lif lif_name -role data  
-data-protocol fcp -home-node node_name -home-port port
```

The `-role` parameter should be `data` and the `data-protocol` parameter should be `fcp`.

3. Create a LIF on each node's HA partner:

```
network interface create -vserver vserver_name -lif lif_name -role data  
-data-protocol fcp -home-node node_name -home-port port
```

The `-role` parameter should be `data` and the `data-protocol` parameter should be `fcp`.

4. Verify that your LIFs have been created and that their operational status is `online`:

```
network interface show -vserver vserver_name -lif lif_name
```

Related concepts

[Considerations for LIFs in cluster SAN environments](#) on page 92

Related information

[NetApp Support](#)

[NetApp Interoperability Matrix Tool](#)

Creating LUNs and mapping igroups

As part of configuring your SAN environment, you must create LUNs, create your initiator groups (igroups), and map your LUNs to your igroups.

Before you begin

- You must have created your aggregates, volumes, and Storage Virtual Machines (SVMs).
- You must have enabled block access with FC or iSCSI.
- You must have created SAN LIFs on the LUN owning node and the owning node's HA partner.

About this task

When you create a LUN, you must specify the LUN OS type. The actual size of the LUN might vary slightly based on the OS type of the LUN. The LUN OS type cannot be modified after the LUN is created.

The metadata for each LUN requires approximately 64 KB of space in the containing aggregate. When you create a LUN, you must ensure that the containing aggregate has enough space for the LUN's metadata. If the aggregate does not contain enough space for the LUN's metadata, some hosts might not be able to access the LUN. If necessary, you can grow your LUN up to 10 times its original size. For example, if you create a 100 GB LUN, you can grow that LUN to 1,000 GB. You cannot exceed 16 TB, which is the maximum LUN size limit.

ALUA is always enabled during LUN creation. You cannot change the ALUA setting.

Steps

1. Create your LUNs:

```
lun create -vserver vserver_name -volume volume_name -lun lun_name -size
lun_size -ostype lun_ostype -space-reserve enabled|disabled
```

If your host operating system is Windows 2012, you must use the `windows_2008` ostype. Space-reserve is enabled by default. If you want a non-space-reserved LUN, you must set the `space-reserve` option to `disabled`.

2. Create your igroups:

```
igroup create -vserver vserver_name -igroup igroup_name -protocol fcp|
iscsi -ostype lun_ostype -initiator initiator_name
```

If your host operating system is Windows 2012, you must use the `windows_2008` ostype.

3. Map your LUNs to igroups:

```
lun mapping create -vserver vserver_name -volume volume_name -lun
lun_name -igroup igroup_name
```

4. Verify that your LUNs are configured correctly:

```
lun show -vserver vserver_name
```

Related concepts

[Guidelines for mapping LUNs to igroups](#) on page 25

[Selective LUN Map](#) on page 26

[Recommended volume and LUN configuration combinations](#) on page 9

Related tasks

[Enabling block access for SVMs with iSCSI](#) on page 21

[Enabling block access for SVMs with FC](#) on page 21

[Determining the correct volume and LUN configuration combination for your environment](#) on page 12

Related information

[Clustered Data ONTAP 8.3 Physical Storage Management Guide](#)

[Clustered Data ONTAP 8.3 Logical Storage Management Guide](#)

[Clustered Data ONTAP 8.3 Network Management Guide](#)

Enabling block access for a specific host

You must enable block access on your specific host so that your initiators can access your targets.

Before you begin

- You must have network connectivity between the host and the LIFs on the SVM.
- Your FC or iSCSI service must be on and operational.
- You must have LUNs that are mapped to initiator groups (igroups).

Steps

1. Follow steps in your host documentation for enabling block access on your specific hosts.
2. Use the Host Utilities to complete the FC or iSCSI mapping and to discover your LUNs on the host.

Related information

[NetApp Documentation: Host Utilities \(current releases\)](#)

Creating port sets and binding igroups to port sets

In addition to using Selective LUN Map (SLM), you can create a port set and bind the port set to an igroup to further limit which LIFs can be used by an initiator to access a LUN. If you do not bind a port set to an igroup, then all of the initiators in the igroup can access mapped LUNs through all of the LIFs on the node owning the LUN and the owning node's HA partner.

Before you begin

You must have at least one LIF and one igroup.

Unless you are using interface groups (ifgrps), two LIFs are recommended for redundancy for both iSCSI and FC. Only one LIF is recommended for ifgrps.

About this task

It is advantageous to use ports sets with SLM when you have multiple targets on a node and you want to restrict access of a certain target to a certain initiator. Without port sets, all targets on the node will be accessible by all of the initiators with access to the LUN through the node owning the LUN and the owning node's HA partner.

Steps

1. Create a port set containing the appropriate LIFs:

```
portset create -vserver vserver_name -portset portset_name -protocol
protocol -port-name port_name
```

If you are using FC, specify the `protocol` parameter as `fc`. If you are using iSCSI, specify the `protocol` parameter as `iscsi`.

Example

2. Bind the igroup to the port set:

```
lun igroup bind -vserver vserver_name -igroup igroup_name -portset
portset_name
```

3. Verify that your port sets and LIFs are correct:

```
portset show -vserver vserver_name
```

Vserver	Portset	Protocol	Port Names	Igroups
vs3	portset0	iscsi	lif0,lif1	igroup1

Related concepts

[Selective LUN Map](#) on page 26

LUN guidelines for SAN environments

Before you begin setting up your LUN, you need to review LUN guidelines, considerations, and supported configurations.

Considerations when creating a LUN from an existing file

When you create a LUN from an existing file, you have multiple options for the file path.

The LUN should exist in a volume root or a qtree root. The following LUN and file path examples show the current options:

- **LUN and file are in the same volume root.** For example, the file path is `/vol/vol1/file1` and the LUN path must be `/vol/vol1/lun1`.
- **LUN and file are in the same qtree root.** For example, the file path is `/vol/vol1/mtree1/file1` and the LUN path must be `/vol/vol1/mtree1/lun1`.
- **LUN is in a volume root and the file is one directory below the same volume root.** For example, the file path is `/vol/vol1/subdir1/file1` and the LUN path must be `/vol/vol1/lun1`.
- **LUN is in a qtree root and the file is one directory below the same qtree root.** For example, the file path is `/vol/vol1/mtree1/subdir1/file1` and the LUN path must be `/vol/vol1/mtree1/lun1`.

Guidelines for assigning LUN IDs

You can assign a number for the LUN ID, or you can accept the default LUN ID. However, your Host Utilities documentation might have additional guidelines about how to assign LUN IDs.

Typically, the default LUN ID begins with 0 and is assigned in increments of 1 for each additional mapped LUN. The host associates the LUN ID with the location and path name of the LUN. The range of valid LUN ID numbers depends on the host. For detailed information, see the documentation provided with your Host Utilities.

Related concepts

[What Host Utilities are](#) on page 97

Related information

[NetApp Documentation: Host Utilities \(current releases\)](#)

Guidelines for mapping LUNs to igroups

There are several important guidelines that you must follow when mapping LUNs to an igroup.

- You can map a LUN only once to an igroup.
- You can map a LUN only once to a specific initiator through the igroup.
- You can add a single initiator to multiple igroups, but the initiator can be mapped to a LUN only once.
- You cannot use the same LUN ID for two LUNs mapped to the same igroup.
- You should use the same protocol type for igroups and port sets.

Managing LUNs

After you create your LUNs, you can manage them in a number of ways. For example, you can control LUN availability, unmap a LUN from an igroup, delete a LUN, and increase the LUN size.

Selective LUN Map

Selective LUN Map (SLM) reduces the number of paths from the host to the LUN. With SLM, when a new LUN map is created, the LUN is accessible only through paths on the node owning the LUN and its HA partner.

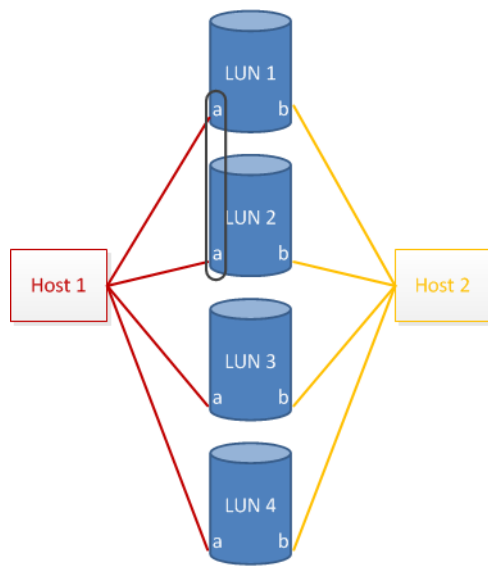
SLM enables management of a single igroup per host and also supports nondisruptive LUN move operations that do not require portset manipulation or LUN remapping.

Portsets can be used with SLM just as in previous versions of Data ONTAP to further restrict access of certain targets to certain initiators. When using SLM with portsets, LUNs will be accessible on the set of LIFs in the portset on the node that owns the LUN and on that node's HA partner.

Beginning with Data ONTAP 8.3 SLM is enabled by default on all new LUN maps. For LUNs created prior to Data ONTAP 8.3, you can manually apply SLM by using the `lun mapping remove-reporting-nodes` command to remove the LUN reporting nodes and restrict LUN access to the LUN owning node and its HA partner.

Example

An SVM has 4 nodes and 16 LIFs per node. The storage administrator creates a portset with 2 LIFs from each node, binds the portset to an igroup, and maps a LUN to the igroup. The LUN is only accessible on the node that owns the LUN and that node's HA partner through the 4 LIFS specified in the portset.



Related tasks

[Decreasing mapped LUN paths with SLM for LUNs created prior to Data ONTAP 8.3](#) on page 27

[Modifying SLM reporting nodes](#) on page 28

How to determine whether SLM is enabled on a LUN map

If your environment has a combination of LUNs created in Data ONTAP 8.3 and LUNs transitioned into Data ONTAP 8.3 from previous versions, you might need to determine whether Selective LUN Map (SLM) is enabled on a specific LUN.

You can use the information displayed in the output of the `lun mapping show -fields reporting-nodes, node` command to determine whether SLM is enabled on your LUN map. If SLM is not enabled, "-" will be displayed in the cells under the `reporting-nodes` column of the command output. If SLM is enabled, the list of nodes displayed under the `nodes` column will be duplicated in the `reporting-nodes` column.

Decreasing mapped LUN paths with SLM for LUNs created prior to Data ONTAP 8.3

You can reduce the number of paths mapped to your LUNs by using Selective LUN Map (SLM). Reducing the number of LUN paths simplifies host-side multipath management and ensures that the number of host paths to the LUN does not exceed the maximum supported by Data ONTAP.

Before you begin

Your Storage Virtual Machine (SVM) must be configured with at least one SAN LIF on the desired protocol on the node owning the LUN and node's HA partner.

For FC configurations, two LIFs per node must be configured on separate physical ports that are each connected to different switches or fabrics. For iSCSI configurations, a single iSCSI LIF must be configured on an ifgrp spanning at least two physical ports. This maintains an active/optimized path to the LUN if a single port or switch fails.

About this task

SLM is enabled by default on all new LUN maps created in Data ONTAP 8.3. You must manually apply SLM to LUNs created prior to Data ONTAP 8.3 by manually adding the reporting nodes to the LUN map. Adding the reporting nodes limits the number of host paths to the LUN owning node and its HA partner.

Steps

1. Limit the nodes with paths to the LUN to the node owning the LUN and the owning node's HA partner:


```
lun mapping add-reporting-nodes -vserver vservice_name -path lun_path -igroup igroup_name -local-nodes
```
2. Verify that the LUN has been removed from the existing LUN map:


```
lun mapping show -fields reporting-nodes
```
3. Rescan the host as appropriate for the client operating system to verify that the LUN is mapped only to its owning node and HA partner node.

Related concepts

[Selective LUN Map](#) on page 26

Modifying SLM reporting nodes

If you are moving a LUN or a volume containing LUNs to another HA pair within the same cluster, you should modify the Selective LUN Map (SLM) reporting-nodes list, before initiating the move. This ensures that active/optimized LUN paths are maintained.

Steps

1. Add the destination node and its partner node to the reporting-nodes list of the aggregate or volume:

```
lun mapping add-reporting-nodes -vserver vservice_name -path lun_path -
igroup igroup_name [-destination-aggregate aggregate_name | -
destination-volume volume_name]
```

If you have a consistent naming convention, you can modify multiple LUN mappings at the same time by using `*-igroup` instead of `igroup`.

2. Rescan the host to discover the newly added paths.
3. Add the new paths to your MPIO configuration.
4. Run the command for the needed move operation and wait for completion.
5. Verify that I/O is being serviced through the Active/Optimized path:

```
lun mapping show -fields reporting-nodes
```

6. Remove the previous LUN owner and its partner node from the reporting-nodes list:

```
lun mapping remove-reporting-nodes-remote-nodes -vserver vservice_name -
path lun_path -igroup igroup_name
```

7. Verify that the LUN has been removed from the existing LUN map:

```
lun mapping show -fields reporting-nodes
```

8. Remove stale device entries for the host OS.
9. Rescan the host to verify removal of old paths.

See your host documentation for specific steps to rescan your hosts.

Related concepts

[Selective LUN Map](#) on page 26

Related information


[NetApp Documentation: Host Utilities \(current releases\)](#)

Ways to limit LUN access with port sets and igroups


In addition to using Selective LUN Map (SLM), you can limit access to your LUNs through igroups and port sets.

Port sets can be used with SLM to further restrict access of certain targets to certain initiators. When using SLM with port sets, LUNs will be accessible on the set of LIFs in the port set on the node that owns the LUN and on that node's HA partner.

In the following example, initiator1 does not have a port set. Without a port set, initiator1 can access LUN1 through both LIF1 and LIF2.



You can limit access to LUN1 by using a port set. In the following example, initiator1 can access LUN1 only through LIF1. However, initiator1 cannot access LUN1 through LIF2 because LIF2 is not in port set1.



Related concepts

[Selective LUN Map](#) on page 26

Related tasks

[Creating port sets and binding igroups to port sets](#) on page 24

Considerations for transitioning SAN configurations

In a SAN environment, a disruption in service is required during the transition of a 7-Mode volume to clustered Data ONTAP. You need to shut down your hosts to complete the transition. After transition, you must update your host configurations before you can begin serving data in clustered Data ONTAP.

You need to schedule a maintenance window during which you can shut down your hosts and complete the transition.

For more information about transitioning your 7-Mode SAN volumes to clustered Data ONTAP, see the *7-Mode Transition Tool Copy-Based Transition Guide*.

Capabilities and restrictions of transitioned LUNs

LUNs that have been transitioned from Data ONTAP operating in 7-Mode to clustered Data ONTAP have certain capabilities and restrictions that affect the way the LUNs can be managed.

You can do the following with transitioned LUNs:

- View the LUN using the `lun show` command
- View the inventory of LUNs transitioned from the 7-Mode volume using the `transition 7-mode show` command
- Restore a volume from a 7-Mode Snapshot copy
Restoring the volume transitions all of the LUNs captured in the Snapshot copy

- Restore a single LUN from a 7-Mode Snapshot copy using the `snapshot restore-file` command
- Create a clone of a LUN in a 7-Mode Snapshot copy
- Restore a range of blocks from a LUN captured in a 7-Mode Snapshot copy
- Create a FlexClone of the volume using a 7-Mode Snapshot copy

You cannot do the following with transitioned LUNs:

- Access Snapshot copy-backed LUN clones captured in the volume

Considerations for copying LUNs

There are considerations you should be aware of when copying a LUN.

The cluster administrator can copy a LUN across Storage Virtual Machines (SVMs) within the cluster by using the `lun copy` command. There must be enough space in the source volume for a SIS clone.

LUNs in Snapshot copies can be used as source LUNs for the `lun copy` command. When you copy a LUN using the `lun copy` command, the LUN copy is immediately available for read and write access. The source LUN is unchanged by creation of a LUN copy. Both the source LUN and the LUN copy exist as unique LUNs with different LUN serial numbers. Changes made to the source LUN are not reflected in the LUN copy, and changes made to the LUN copy are not reflected in the source LUN. The LUN mapping of the source LUN is not copied to the new LUN; the LUN copy must be mapped.

Data protection through Snapshot copies occurs at the volume level. Therefore, if you copy a LUN to a volume different from the volume of the source LUN, the destination LUN falls under the data protection scheme of the destination volume. If you do not have Snapshot copies established for the destination volume, Snapshot copies are not created of the LUN copy.

Copying LUNs is a nondisruptive operation.

You cannot copy the following types of LUNs:

- A LUN that has been created from a file
- A LUN that is in NVFail state
- A LUN that is in a load-sharing relationship
- A protocol-endpoint class LUN

Increasing the size of a LUN

You can grow a LUN to approximately 10 times its original size, but not greater than 16 TB.

About this task

For example, if you create a 100 GB LUN, you can grow that LUN to approximately 1,000 GB. However, if you create an 8 TB LUN, you cannot grow it to 80 TB. 16 TB is the approximate maximum LUN size. The actual LUN size might vary slightly based on the OS type of the LUN.

You do not need to take the LUN offline to increase the size. However, after you have increased the size, you must rescan the LUN on the host for the host to recognize the change in size.

See the Command Reference page for the `lun resize` command for more information about resizing a LUN.

Steps

1. Increase the size of the LUN:

```
lun resize -vserver vserver_name -volume volume_name -lun lun_name -size
lun_size
```

2. Verify the increased LUN size:

```
lun show -vserver vserver_name
```

3. Rescan the LUN on the host.
4. Follow your host documentation to make the newly created LUN size visible to the host file system.

Related information

Clustered Data ONTAP 8.3.2 man page: [lun resize](#) - Changes the size of the LUN to the input value size.

Decreasing the size of a LUN

Before you decrease the size of a LUN, the host needs to migrate the blocks containing the LUN data into the boundary of the smaller LUN size. You should use a tool such as SnapDrive for Windows to ensure that the LUN is properly decreased without truncating blocks containing LUN data. Manually decreasing the size of your LUN is not recommended.

About this task

After you decrease the size of your LUN, Data ONTAP automatically notifies the initiator that the LUN size has decreased. However, additional steps might be required on your host for the host to recognize the new LUN size. Check your host documentation for specific information about decreasing the size of the host file structure.

Moving LUNs

You can move a LUN across volumes within a Storage Virtual Machine (SVM), but you cannot move a LUN across SVMs. LUNs moved across volumes within an SVM are moved immediately and without loss of connectivity.

Before you begin

If your LUN is using Selective LUN Map (SLM), the SLM reporting nodes must have been modified to include the destination node and its HA partner.

About this task

Storage efficiency features, such as deduplication, are not preserved during a LUN move. They must be reapplied after the LUN move is completed.

Data protection through Snapshots occurs at the volume level. Therefore, when you move a LUN it fall under the data protection scheme of the destination volume. If you do not have Snapshots established for the destination volume, Snapshots will not be taken of the LUN.

You cannot move a LUN to the following volumes:

- A SnapMirror destination volume
- The SVM root volume

You cannot move the following types of LUNs:

- A LUN that has been created from a file
- A LUN that is in NVFail state
- A LUN that is in a load-sharing relationship
- A protocol-endpoint class LUN

Note: For Solaris `os_type` LUNs that are 1 TB or greater, the host might experience a timeout during the LUN move. For this LUN type you should unmount the LUN before initiating the move.

Steps

1. Run the `lun move start` command to move the LUN.

There will be a very brief period when the LUN is visible on both the origin and destination volume. This is expected and will be resolved upon completion of the move.

2. Run the `lun move show` command to track the status of the move and verify successful completion.

Related concepts

[Selective LUN Map](#) on page 26

Related tasks

[Modifying SLM reporting nodes](#) on page 28

Deleting LUNs

You can delete a LUN from a Storage Virtual Machine (SVM) if you no longer need the LUN. If you create a LUN from a file, you cannot remove the file while the LUN is linked to it. You must first delete the LUN.

Before you begin

The LUN must be unmapped from its igroup before you can delete it.

About this task

If you create a space-reserved LUN from a file, the file becomes space-reserved. If you delete that LUN, the file is no longer space-reserved.

Steps

1. Ensure that the application or host is not using the LUN.
2. Use the `lun mapping delete` to unmap the LUN from the igroup.

Example

```
lun mapping delete -vserver vs5 -volume vo5 -lun lun5 -igroup igr5
```

3. Use the `lun delete` command to delete the LUN.

Example

```
lun delete -vserver vs5 -volume vol15 -lun lun5
```

4. Use the `lun show` command to verify that you deleted the LUN.

Example

```
lun show -vserver vs5
```

Vserver	Path	State	Mapped	Type	Size
vs5	/vol/vol16/lun8	online	mapped	windows	10.00GB

Examining configured and used space of a LUN

Knowing the configured space and actual space used for your LUNs can help you determine the amount of space that can be reclaimed when doing space reclamation, the amount of reserved space that contains data, and the total configured size versus the actual size used for a LUN.

Step

1. Use the `lun show` command to show the configured space versus the actual space used for a LUN.

Example

The following example show the configured space versus the actual space used by the LUNs in the v3 Storage Virtual Machine (SVM):

```
lun show -vserver vs3 -fields path, size, size-used, space-reserve
```

vserver	path	size	space-reserve	size-used
vs3	/vol/vol0/lun1	50.01GB	disabled	25.00GB
vs3	/vol/vol0/lun1_backup	50.01GB	disabled	32.15GB
vs3	/vol/vol0/lun2	75.00GB	disabled	
vs3	/vol/vol0/lun0	5.00GB	enabled	4.50GB

4 entries were displayed.

Commands for managing LUNs

Data ONTAP provides various commands to help you manage your LUNs.

If you want to...	Use this command...
Create a LUN	<code>lun create</code>
Change or add a comment to a LUN	<code>lun modify</code>
Map a LUN to an igroup	<code>lun mapping create</code>
Determine whether a LUN is mapped	<code>lun show</code> or <code>lun mapping show</code>
Move a LUN within the same volume	<code>lun rename</code>
Move a LUN to a different volume within a Storage Virtual Machine (SVM)	<code>lun move start</code>

If you want to...	Use this command...
Move a LUN to a different qtree within the same volume	<code>lun move-in-volume</code>
Rename a LUN	<code>lun rename</code>
Copy a LUN	<code>lun copy</code>
Modify a LUN	<code>lun modify</code>
Enable or disable space reservations for LUNs	<code>lun modify</code>
Take a LUN offline	<code>lun modify</code>
Bring a LUN online	<code>lun modify</code>
Determine the maximum size of a LUN	<code>lun maxsize</code>
Resize a LUN	<code>lun resize</code>
Display LUN configurations	<code>lun show</code>
Display detailed LUN information	<code>lun show -instance</code>
Display LUN paths and LUN IDs	<code>lun mapping show</code>
Unmap a LUN from an igroup	<code>lun mapping delete</code>
Delete a LUN	<code>lun delete</code>
View a man page for a command	<code>man <i>command name</i></code>

See the man page for each command for more information.

Commands for managing port sets

Data ONTAP provides commands to manage your port sets.

See *How to limit LUN access in a virtualized environment* for more information how you can use portsets to limit LUN access.

If you want to....	Use this command...
Create a new port set	<code>lun portset create</code>
Add LIFs to a port set	<code>lun portset add</code>
Display LIFs in a port set	<code>lun portset show</code>
Display igroups that are bound to port sets	<code>lun portset show</code>
Bind an igroup to a port set	<code>lun igroup bind</code>
Unbind an igroup from a port set	<code>lun igroup unbind</code>
Remove a LIF from a port set	<code>lun portset remove</code> Note: If only one LIF is in the port set, you must use the <code>lun igroup unbind</code> command to unbind the port set from the initiator group. After the port set is no longer bound to an initiator group, you can remove the last LIF from the port set.
Delete a port set	<code>lun portset delete</code>


If you want to....	Use this command...
View a man page for a command	<code>man command name</code>

See the man page for each command for more information.

I/O misalignments might occur on properly aligned LUNs

Data ONTAP might report I/O misalignments on properly aligned LUNs. In general, these misalignment warnings can be disregarded as long as you are confident that your LUN is properly provisioned and your partitioning table is correct.

LUNs and hard disks both provide storage as blocks. Because the block size for disks on the host is 512 bytes, LUNs present blocks of that size to the host while actually using larger, 4-KB blocks to store data. The 512-byte data block used by the host is referred to as a logical block. The 4-KB data block used by the LUN to store data is referred to as a physical block. This means that there are eight 512-byte logical blocks in each 4-KB physical block.



The host operating system can begin a read or write I/O operation at any logical block. I/O operations are only considered aligned when they begin at the first logical block in the physical block. If an I/O operation begins at a logical block that is not also the start of a physical block, the I/O is considered misaligned. Data ONTAP automatically detects the misalignment and reports it on the LUN. However, the presence of misaligned I/O does not necessarily mean that the LUN is also misaligned. It is possible for misaligned I/O to be reported on properly aligned LUNs.

If further investigation is required, technical support can run diagnostic commands that show detailed I/O alignment data to confirm the presence or absence of true LUN misalignment.

For more information about tools for correcting alignment problems, see the following documentation:

- *Data ONTAP DSM for Windows MPIO Installation and Administration Guide*

- *Windows Host Utilities Installation and Setup Guide*
- *Virtual Storage Console for VMware vSphere Installation and Administration Guide*

Related concepts

[ESX boot LUNs report as misaligned](#) on page 37

Related information

[NetApp Technical Report 3747: Best Practices for File System Alignment in Virtual Environments](#)
[NetApp Documentation: Host Utilities \(current releases\)](#)

How to achieve I/O alignment using LUN OS types

To achieve I/O alignment with your OS partitioning scheme, you should use the recommended Data ONTAP LUN `ostype` value that most closely matches your operating system.

The partition scheme employed by the host operating system is a major contributing factor to I/O misalignments. Some Data ONTAP LUN `ostype` values use a special offset known as a “prefix” to enable the default partitioning scheme used by the host operating system to be aligned.

Note: In some circumstances, a custom partitioning table might be required to achieve I/O alignment. However, for `ostype` values with a “prefix” value greater than 0, a custom partition might create misaligned I/O.

The LUN `ostype` values in the following table should be used based on your operating system.

LUN <code>ostype</code>	Prefix (bytes)	Prefix (sectors)	Operating system
windows	32,256	63	Windows 2000, 2003 (MBR format)
windows_gpt	17,408	34	Windows 2003 (GPT format)
windows_2008	0	0	Windows 2008 and later
hyper_v	0	0	Windows 2008 Hyper-V and later
linux	0	0	All Linux distributions
xen	0	0	Citrix XenServer
vmware	0	0	VMware ESX
solaris	1MB	2,048	Solaris
solaris_efi	17,408	34	Solaris
hpux	0	0	HP-UX
aix	0	0	AIX

Special I/O alignment considerations for Linux

Linux distributions offer a wide variety of ways to use a LUN including as raw devices for databases, various volume managers, and file systems. It is not necessary to create partitions on a LUN when used as a raw device or as physical volume in a logical volume.

If the LUN will be used without a volume manager, you should partition the LUN to have one partition that begins at an aligned offset, which is a sector that is an even multiple of eight logical blocks.

Special I/O alignment considerations for Solaris LUNs

You need to consider various factors when determining whether you should use the `solaris` `ostype` or the `solaris_efi` `ostype`.

See the *Solaris Host Utilities Installation and Administration Guide* for detailed information.

Related information

[NetApp Documentation: Host Utilities \(current releases\)](#)

ESX boot LUNs report as misaligned

LUNs used as ESX boot LUNs are typically reported by Data ONTAP as misaligned. ESX creates multiple partitions on the boot LUN, making it very difficult to align. Misaligned ESX boot LUNs are not typically a performance problem because the total amount of misaligned I/O is small. Assuming that the LUN was correctly provisioned with the `vmware` `ostype`, no action is needed.

Related concepts

[I/O misalignments might occur on properly aligned LUNs](#) on page 35

Controlling and monitoring I/O performance to LUNs by using Storage QoS

You can control input/output (I/O) performance to LUNs by assigning LUNs to Storage QoS policy groups. You might control I/O performance to ensure that workloads achieve specific performance objectives or to throttle a workload that negatively impacts other workloads.

About this task

Policy groups enforce a maximum throughput limit (for example, 100 MB/s). You can create a policy group without specifying a maximum throughput, which enables you to monitor performance before you control the workload.

You can also assign Storage Virtual Machines (SVMs) with FlexVol volumes and LUNs to policy groups.

Note the following requirements about assigning a LUN to a policy group:

- The LUN must be contained by the SVM to which the policy group belongs. You specify the SVM when you create the policy group.
- If you assign a LUN to a policy group, then you cannot assign the LUN's containing volume or SVM to a policy group.

For more information about how to use Storage QoS, see the *Clustered Data ONTAP System Administration Guide for Cluster Administrators*.

Steps

1. Use the `qos policy-group create` command to create a policy group.
2. Use the `lun create` command or the `lun modify` command with the `-qos-policy-group` parameter to assign a LUN to a policy group.
3. Use the `qos statistics` commands to view performance data.
4. If necessary, use the `qos policy-group modify` command to adjust the policy group's maximum throughput limit.

Tools available to effectively monitor your LUNs

Tools are available to help you effectively monitor your LUNs and avoid running out of space.

- OnCommand Unified Manager is a free tool that enables you to manage all storage across all clusters in your environment.
- System Manager is a graphical user interface built into Data ONTAP that enables you to manually manage storage needs at the cluster level.
- OnCommand Insight presents a single view of your storage infrastructure and enables you to set up automatic monitoring, alerts, and reporting when your LUNs, volumes, and aggregates are running out of storage space.

Ways to address issues when LUNs go offline

When no space is available for writes, LUNs go offline to preserve data integrity. LUNs can run out of space and go offline for various reasons, and there are several ways you can address the issue.

If the...	You can...
Aggregate is full	<ul style="list-style-type: none"> • Add more disks. • Use the <code>volume modify</code> command to shrink a volume that has available space. • If you have space-guarantee volumes that have available space, change the volume space guarantee to none with the <code>volume modify</code> command.

If the...	You can...
Volume is full but there is space available in the containing aggregate	<ul style="list-style-type: none"> • For space guarantee volumes, use the <code>volume modify</code> command to increase the size of your volume. • For thinly provisioned volumes, use the <code>volume modify</code> command to increase the maximum size of your volume. If volume autogrow is not enabled, use <code>volume modify -autogrow-mode</code> to enable it. • Delete Snapshot copies manually with the <code>volume snapshot delete</code> command, or use the <code>volume snapshot autodelete modify</code> command to automatically delete Snapshot copies.

Related information

[Clustered Data ONTAP 8.3 Physical Storage Management Guide](#)

[Clustered Data ONTAP 8.3 Logical Storage Management Guide](#)

[Clustered Data ONTAP 8.3 Data Protection Guide](#)

What host-side space management is

In a thinly provisioned environment, host side space management completes the process of managing space from the storage system that has been freed in the host file system.

A host file system contains metadata to keep track of which blocks are available to store new data and which blocks contain valid data that must not be overwritten. This metadata is stored within the LUN. When a file is deleted in the host file system, the file system metadata is updated to mark that file's blocks as free space. Total file system free space is then recalculated to include the newly freed blocks. To the storage system, these metadata updates appear no different from any other writes being performed by the host. Therefore, the storage system is unaware that any deletions have occurred.

This creates a discrepancy between the amount of free space reported by the host and the amount of free space reported by the underlying storage system. For example, suppose you have a newly provisioned 200-GB LUN assigned to your host by your storage system. Both the host and the storage system report 200 GB of free space. Your host then writes 100 GB of data. At this point, both the host and storage system report 100 GB of used space and 100 GB of unused space.

Then you delete 50 GB of data from your host. At this point, your host will report 50 GB of used space and 150 GB of unused space. However, your storage system will report 100 GB of used space and 100 GB of unused space.

Host-side space management uses various methods to reconcile the space differential between the host and the storage system.

Automatic host-side space management with SCSI thinly provisioned LUNs

If your host supports SCSI thin provisioning, you can enable the `space-allocation` option in Data ONTAP to turn on automatic host-side space management.

Enabling SCSI thin provisioning enables you to do the following.

- **Automatic host-side space management**
When data is deleted on a host that supports SCSI thin provisioning, host-side space management identifies the blocks of deleted data on the host file system and automatically issues one or more SCSI `UNMAP` commands to free corresponding blocks on the storage system.
- **Notify the host when a LUN runs out of space while keeping the LUN online**
On hosts that do not support SCSI thin provisioning, when the volume containing LUN runs out of space and cannot automatically grow, Data ONTAP takes the LUN offline. However, on hosts that support SCSI thin provisioning, Data ONTAP does not take the LUN offline when it runs out of space. The LUN remains online in read-only mode and the host is notified that the LUN can no longer accept writes.

Related concepts

[Tools available to effectively monitor your LUNs](#) on page 38

Related tasks

[Enabling space allocation for SCSI thinly provisioned LUNs](#) on page 40

Related references

[Host support for SCSI thin provisioning](#) on page 41

[Host support for SCSI thin provisioning](#) on page 41

Related information

[NetApp Documentation: Host Utilities \(current releases\)](#)

Enabling space allocation for SCSI thinly provisioned LUNs

If you set the `space-allocation` option to **enabled**, Data ONTAP notifies the host when the volume has run out of space and the LUN in the volume cannot accept writes. This option also enables Data ONTAP to reclaim space automatically when your host deletes data.

About this task

The `space-allocation` option is set to **disabled** by default, and you must take the LUN offline to enable space allocation. After you enable space allocation, you must perform discovery on the host before the host will recognize that space allocation has been enabled.

Steps

1. Take the LUN offline:

```
lun modify -vserver vserver_name -volume volume_name -lun lun_name -state offline
```

2. Set the `-space-allocation` parameter to **enabled**:


```
lun modify -vserver vservice_name -volume volume_name -lun lun_name -
space-allocation enabled
```

3. Verify that space allocation is enabled:

```
lun show -vserver vservice_name -volume volume_name -lun lun_name -fields
space-allocation
```

4. Bring the LUN online:

```
lun modify -vserver vservice_name -volume volume_name -lun lun_name -
state online
```

Example

5. On the host, rescan all disks to ensure that the change to the `-space-allocation` option is correctly discovered.

Related concepts

[Automatic host-side space management with SCSI thinly provisioned LUNs](#) on page 40

[Tools available to effectively monitor your LUNs](#) on page 38

[Using FlexClone LUNs to protect your data](#) on page 52

[Effect of moving or copying a LUN on Snapshot copies](#) on page 50

Related references

[Host support for SCSI thin provisioning](#) on page 41

Host support for SCSI thin provisioning

To leverage the benefits of SCSI thin provisioning, it must be supported by your host. SCSI thin provisioning, uses the Logical Block Provisioning feature as defined in the SCSI SBC-3 standard. Only hosts that support this standard can use SCSI thin provisioning in Data ONTAP.

The following hosts currently support SCSI thin provisioning when you enable space allocation:

- VMware ESX 5.0 and later
- Red Hat Enterprise Linux 6.2 and later
- Microsoft Windows 2012

When you enable the space allocation functionality in Data ONTAP, you turn on the following SCSI thin provisioning features:

- Unmapping and reporting space usage for space reclamation
- Reporting resource exhaustion errors

Related concepts

[Automatic host-side space management with SCSI thinly provisioned LUNs](#) on page 40

Related tasks

[Enabling space allocation for SCSI thinly provisioned LUNs](#) on page 40

Managing igroups

You can manage your initiator groups (igroups) by performing a range of tasks, including creating, destroying, and renaming igroups.

Commands for managing igroups

Data ONTAP provides commands to manage your igroups.

If you want to...	Use this command...
Create a new igroup	<code>lun igroup create</code>
Add an initiator to an igroup	<code>lun igroup add</code>
Bind an igroup to a port set	<code>lun igroup bind</code>
Rename an existing igroup	<code>lun igroup rename</code>
Set the OS type for an igroup	<code>lun igroup modify</code>
Display detailed LUN igroup information	<code>lun igroup show -instance</code>
Remove an initiator from an igroup	<code>lun igroup remove</code>
Unbind an igroup from a port set	<code>lun igroup unbind</code>
Delete an igroup	<code>lun igroup delete</code>
View the man page for a command	<code>man <i>command name</i></code>

See the man page for each command for more information.

Related concepts

[What igroups are](#) on page 42

What igroups are

Initiator groups (igroups) are tables of FC protocol host WWPNs or iSCSI host node names. You can define igroups and map them to LUNs to control which initiators have access to LUNs.

Typically, you want all of the host's initiator ports or software initiators to have access to a LUN. If you are using multipathing software or have clustered hosts, each initiator port or software initiator of each clustered host needs redundant paths to the same LUN.

You can create igroups that specify which initiators have access to the LUNs either before or after you create LUNs, but you must create igroups before you can map a LUN to an igroup.

Initiator groups can have multiple initiators, and multiple igroups can have the same initiator. However, you cannot map a LUN to multiple igroups that have the same initiator. An initiator cannot be a member of igroups of differing otypes.

Related concepts

[Managing igroups](#) on page 42

Example of how igroups give LUN access

You can create multiple igroups to define which LUNs are available to your hosts. For example, if you have a host cluster, you can use igroups to ensure that specific LUNs are visible to only one host in the cluster or to all the hosts in the cluster.

The following table illustrates how four igroups give access to the LUNs for four different hosts that are accessing the storage system. The clustered hosts (Host3 and Host4) are both members of the same igroup (group3) and can access the LUNs mapped to this igroup. The igroup named group4 contains the WWPNs of Host4 to store local information that is not intended to be seen by its partner.

Hosts with HBA WWPNs, IQNs, or EUIs	igroups	WWPNs, IQNs, EUIs added to igroups	LUNs mapped to igroups
Host1, single-path (iSCSI software initiator) iqn.1991-05.com.microsoft:host1	group1	iqn.1991-05.com.microsoft:host1	/vol/vol2/ lun1
Host2, multipath (two HBAs) 10:00:00:00:c9:2b:6b:3c 10:00:00:00:c9:2b:02:3c	group2	10:00:00:00:c9:2b:6b:3c 10:00:00:00:c9:2b:02:3c	/vol/vol2/ lun2
Host3, multipath, clustered with host 4 10:00:00:00:c9:2b:32:1b 10:00:00:00:c9:2b:41:02	group3	10:00:00:00:c9:2b:32:1b 10:00:00:00:c9:2b:41:02 10:00:00:00:c9:2b:51:2c 10:00:00:00:c9:2b:47:a2	/vol/vol2/ qtreen1/ lun3
Host4, multipath, clustered (connected to Host3) 10:00:00:00:c9:2b:51:2c 10:00:00:00:c9:2b:47:a2	group4	10:00:00:00:c9:2b:51:2c 10:00:00:00:c9:2b:47:a2	/vol/vol2/ qtreen2/ lun4 /vol/vol2/ qtreen1/ lun5

How to specify initiator WWPNs and node names for an igroup

You can specify the node names and WWPNs of the initiators when you create an igroup or you can add them later. If you choose to specify the initiator node names and WWPNs when you create the LUN, they can be later removed, if needed.

Follow instructions in your Host Utilities documentation to obtain WWPNs and to find the node names associated with a specific host. For hosts running the latest ESX software, Virtual Storage Console (also known as OnCommand Plug-in for VMware) has replaced the Host Utilities.

Related concepts

[How FC nodes are identified](#) on page 101

[What Host Utilities are](#) on page 97

Related information

[NetApp Documentation: Host Utilities \(current releases\)](#)

Why Data ONTAP uses ALUA

Data ONTAP uses Asymmetric Logical Unit Access (ALUA) for iSCSI and FC to identify optimized and unoptimized paths. Host support for ALUA is required for cluster failover and data mobility operations to work correctly.

ALUA is an industry standard protocol for identifying optimized paths between a storage system and a host. ALUA enables the initiator to query the target about path attributes, such as primary path and secondary path. It also allows the target to communicate events back to the initiator. It is beneficial because multipathing software can be developed to support any storage array. Proprietary SCSI commands are no longer required to determine primary and secondary paths.

Related information

[NetApp Interoperability](#)

MPIO and ALUA

Clustered Data ONTAP uses multipath I/O (MPIO) and ALUA to provide high availability for iSCSI and FC.

MPIO provides more than one physical path between the controller and the initiator. ALUA determines the optimized and non-optimized paths between the initiator and the LUN. The optimized paths are the most direct route from your initiator to the LUN. The non-optimized paths provide additional routes to the LUN if the optimized paths are not available.

Your initiator must support MPIO and ALUA in cluster SAN environment.

Managing LIFs

LIFs can be removed from port sets, can be moved to different nodes within a Storage Virtual Machine (SVM), and can be deleted.

For more information on configuring LIFs, see the *Clustered Data ONTAP Network Management Guide*.

Related concepts

[Considerations for LIFs in cluster SAN environments](#) on page 92

Considerations for SAN LIF migration

You only need to perform a LIF migration if you are changing the contents of your cluster, for example, adding nodes to the cluster or deleting nodes from the cluster. If you perform a LIF migration, you do not have to re-zone your FC fabric or create new iSCSI sessions between the attached hosts of your cluster and the new target interface.

You cannot migrate a SAN LIF using the `network interface migrate` command. SAN LIF migrations must be performed by taking the LIF offline, moving the LIF to a different home node or port, and then bringing it back online in its new location. ALUA provides redundant paths and automatic path selection as part of any clustered Data ONTAP SAN solution. Therefore, there is no I/O interruption when the LIF is taken offline for the migration. The host simply retries and then moves I/O to another LIF.

Using LIF migration, you can nondisruptively do the following:

- Replace one HA pair of a cluster with an upgraded HA pair in a way that is transparent to hosts accessing LUN data
- Upgrade a target interface card
- Shift the resources of a Storage Virtual Machine (SVM) from one set of nodes in the cluster to another set of nodes in the cluster

Removing a SAN LIF from a port set

If the LIF you want to delete or move is in a port set, you must remove the LIF from the port set before you can delete or move the LIF.

About this task

You need to do Step 1 in the following procedure only if one LIF is in the port set. You cannot remove the last LIF in a port set if the port set is bound to an initiator group. Otherwise, you can start with Step 2 if multiple LIFs are in the port set.

Steps

1. If only one LIF is in the port set, use the `lun igroup unbind` command to unbind the port set from the initiator group.

Note: When you unbind an initiator group from a port set, all of the initiators in the initiator group have access to all target LUNs mapped to the initiator group on all network interfaces.

Example

```
cluster1::>lun igroup unbind -vserver vs1 -igroup ig1
```

2. Use the `lun portset remove` command to remove the LIF from the port set.

```
cluster1::> port set remove -vserver vs1 -portset ps1 -port-name lif1
```

Moving SAN LIFs

If a node needs to be taken offline, you can move a SAN LIF to preserve its configuration information, such as its WWPN, and avoid rezoning the switch fabric. Because a SAN LIF must be taken offline before it is moved, host traffic must rely on host multipathing software to ensure nondisruptive access to the LUN. You can move SAN LIFs to any node in a cluster, but you cannot move the SAN LIFs between Storage Virtual Machines (SVMs).

Before you begin

If the LIF is a member of a port set, the LIF must have been removed from the port set before the LIF can be moved to a different node.

About this task

The destination node and physical port for a LIF that you want to move must be on the same FC fabric or Ethernet network. If you move a LIF to a different fabric that has not been properly zoned, or if you move a LIF to an Ethernet network that does not have connectivity between iSCSI initiator and target, the LIF will be inaccessible when you bring it back online.

Steps

1. View the administrative and operational status of the LIF:

```
network interface show -vserver vserver_name
```

2. Change the status of the LIF to **down** (offline):

```
network interface modify -vserver vserver_name -lif LIF_name -status-admin down
```

3. Assign the LIF a new node and port:

```
network interface modify -vserver vserver_name -lif LIF_name -home-node node_name -home-port port_name
```

4. Change the status of the LIF to **up** (online):

```
network interface modify -vserver vserver_name -lif LIF_name -status-admin up
```

5. Verify your changes:

```
network interface show -vserver vserver_name
```

Related tasks

[Removing a SAN LIF from a port set](#) on page 45

Deleting a LIF in a SAN environment

When you delete a LIF, you should ensure that the host connected to the LIF can access the LUNs through another path.

Before you begin

If the LIF you want to delete is a member of a port set, you must first remove the LIF from the port set before you can delete the LIF.

Steps

1. Use the `network interface delete` command to delete the LIF.

Example

```
cluster1::> network interface delete -vserver vs1 -lif lif1
```

2. Use the `network interface show` command to verify that you deleted the LIF.

Example

```
cluster1::> network interface show -vserver vs1
```

Logical Vserver	Status Interface	Network Admin/Oper	Address/Mask	Current Node	Current Port	Is Home
vs1	lif2	up/up	192.168.2.72/24	node-01	e0b	true
	lif3	up/up	192.168.2.73/24	node-01	e0b	true

Related tasks

[Removing a SAN LIF from a port set](#) on page 45

FC and FCoE LIFs on the same port need to be in separate zones

When using Cisco FC and FCoE switches, a single fabric zone must not contain more than one target LIF for the same physical port. If multiple LIFs on the same port are in the same zone, then the LIF ports might fail to recover from a connection loss.

Multiple LIFs for the FC and FCoE protocols can share physical ports on a node as long as they are in different zones. Cisco FC and FCoE switches require each LIF on a given port to be in a separate zone from the other LIFs on that port.

A single zone can have both FC and FCoE LIFs. A zone can contain a LIF from every target port in the cluster, but be careful to not exceed the host's path limits.

LIFs on different physical ports can be in the same zone.

While this is a requirement for Cisco switches, separating LIFs is a good idea for all switches.

Considerations for adding nodes to a cluster

You need to be aware of certain considerations when adding nodes to a cluster.

- You must create LIFs on the new nodes as appropriate before you create LUNs on those new nodes.
- You must discover those LIFs from the hosts as dictated by the host stack and protocol.

Data protection methods in SAN environments

You can protect your data by making copies of it so that it is available for restoration in the event of accidental deletion, application crashes, data corruption, or disaster. Depending on your data protection and backup needs, Data ONTAP offers a variety of methods that enable you to protect your data.

Snapshot copy

Enables you to manually or automatically create, schedule, and maintain multiple backups of your LUNs. Snapshot copies use only a minimal amount of additional volume space and do not have a performance cost. If your LUN data is accidentally modified or deleted, that data can easily and quickly be restored from one of the latest Snapshot copies.

FlexClone LUNs (FlexClone license required)

Point-in-time, writable copies of another LUN in an active volume or in a Snapshot copy. A clone and its parent can be modified independently without affecting each other.

SnapRestore (license required)

Enables you to perform fast, space-efficient, on-request data recovery from Snapshot copies on an entire volume. You can use SnapRestore to restore a LUN to an earlier preserved state without rebooting the storage system.

Data protection mirror copies (SnapMirror license required)

Provide asynchronous disaster recovery by enabling you to periodically create Snapshot copies of data on your volume; copy those Snapshot copies over a local or wide area network to a partner volume, usually on another cluster; and retain those Snapshot copies. The mirror copy on the partner volume provides quick availability and restoration of data from the time of the last Snapshot copy, if the data on the source volume is corrupted or lost.

SnapVault backups (SnapVault license required)

Provides storage efficient and long-term retention of backups. SnapVault relationships enable you to back up selected Snapshot copies of volumes to a destination volume and retain the backups.

If you conduct tape backups and archival operations, you can perform them on the data that is already backed up on the SnapVault secondary volume.

SnapDrive for Windows or UNIX (SnapDrive licence required)

Configures access to LUNs, manages LUNs, and manages storage system Snapshot copies directly from a Windows or UNIX hosts.

Native tape backup and recovery

Support for most existing tape drives is included in Data ONTAP, as well as a method for tape vendors to dynamically add support for new devices. Data ONTAP also supports the Remote Magnetic Tape (RMT) protocol, enabling backup and recovery to any capable system.

Related information

[NetApp Documentation: SnapDrive for UNIX](#)

[NetApp Documentation: SnapDrive for Windows \(current releases\)](#)

Clustered Data ONTAP 8.3 Data Protection Guide

Clustered Data ONTAP 8.3 Data Protection Tape Backup and Recovery Guide

Effect of moving or copying a LUN on Snapshot copies

Snapshot copies are created of the volume. Therefore, if you copy or move a LUN to a different volume, the moved LUN or LUN copy will fall under the data protection scheme of the destination volume. If you do not have Snapshot copies established for the destination volume, Snapshot copies will not be created of the LUN or LUN copy on that volume.

Related information

Clustered Data ONTAP 8.3 Data Protection Guide

Restoring a single LUN from a Snapshot copy

You can restore a single LUN from a Snapshot copy without restoring the entire volume that contains the single LUN. You can restore the LUN in place or to a new path in the volume. The operation restores only the single LUN without impacting other files or LUNs in the volume.

Before you begin

- You must have enough space on your volume to complete the restore:
 - If you are restoring a space-reserved LUN where the fractional reserve is 0%, you require 1 times the size of the restored LUN.
 - If you are restoring a space-reserved LUN where the fractional reserve is 100%, you require 2 times the size of the restored LUN.
 - If you are restoring a non-space-reserved LUN, you only require the actual space used for the restored LUN.
- A Snapshot copy of the destination LUN must have been created.
If the restore operation fails, the destination LUN might be truncated. In such cases, you can use the Snapshot copy to prevent data loss.
- A Snapshot copy of the source LUN must have been created.
In rare cases, the LUN restore can fail, leaving the source LUN unusable. If this occurs, you can use the Snapshot copy to return the LUN to the state just before the restore attempt.
- The destination LUN and source LUN must have the same OS type.
If your destination LUN has a different OS type from your source LUN, your host can lose data access to the destination LUN after the restore operation.

Steps

1. From the host, stop all host access to the LUN.
2. Unmount the LUN on its host to ensure that the host does not access the LUN.
3. Unmap the LUN:


```
lun mapping delete -vserver vservice_name -volume volume_name -lun
lun_name -igroup igroup_name
```
4. Determine the Snapshot copy you want to restore your LUN to:


```
volume snapshot show -vserver vservice_name -volume volume_name
```

5. Create a Snapshot copy of the LUN prior to restoring the LUN:

```
volume snapshot create -vserver vserver_name -volume volume_name -
snapshot snapshot_name
```

6. Restore the specified LUN in a volume:

```
volume snapshot restore-file -vserver vserver_name -volume volume_name -
snapshot snapshot_name -path lun_path
```

7. Follow the steps on the screen.

8. If necessary, bring the LUN online:

```
lun modify -vserver vserver_name -path lun_path -state online
```

9. If necessary, remap the LUN:

```
lun mapping create -vserver vserver_name -volume volume_name -lun
lun_name -igroup igroup_name
```

10. From the host, remount the LUN.

11. From the host, restart access to the LUN.

Restoring all LUNs in a volume from a Snapshot copy

You can use `volume snapshot restore` command to restore all the LUNs in a specified volume from a Snapshot copy.

Steps

1. From the host, stop all host access to the LUNs.

Using SnapRestore without stopping all host access to LUNs in the volume can cause data corruption and system errors.

2. Unmount the LUNs on that host so that the host cannot access the LUNs.

3. Unmap your LUNs:

```
lun mapping delete -vserver vserver_name -volume volume_name -lun
lun_name -igroup igroup_name
```

4. Determine the Snapshot copy to which you want to restore your volume:

```
volume snapshot show -vserver vserver_name -volume volume_name
```

5. Change your privilege setting to advanced:

```
set -privilege advanced
```

6. Restore your data:

```
volume snapshot restore -vserver vserver_name -volume volume_name -
snapshot snapshot_name
```

7. Follow the instructions on the screen.

8. Remap your LUNs:

```
lun mapping create -vserver vserver_name -volume volume_name -lun
lun_name -igroup igroup_name
```

9. Verify that your LUNs are online:

```
lun show -vserver vserver_name -path lun_path -fields state
```

10. If your LUNs are not online, bring them online:

```
lun modify -vserver vserver_name -path lun_path -state online
```

11. Change your privilege setting to admin:

```
set -privilege admin
```

12. From the host, remount your LUNs.
13. From the host, restart access to your LUNs.

Deleting one or more existing Snapshot copies in a volume

You can manually delete one or more existing Snapshot copies in the volume. You might want to do this if you need more space on your volume.

Steps

1. Use the `volume snapshot show` command to verify that these are the Snapshot copies that you want to delete.

Example

```
volume snapshot show -vserver vs3 -volume vol3
```

Vserver	Volume	Snapshot	Size	---Blocks---	
				Total%	Used%
vs3	vol3	weekly.2013-05-01_0015	100KB	0%	38%
		weekly.2013-05-08_0015	76KB	0%	32%
		daily.2013-05-09_0010	76KB	0%	32%
		daily.2013-05-10_0010	76KB	0%	32%
		hourly.2013-05-10_1005	72KB	0%	31%
		hourly.2013-05-10_1105	72KB	0%	31%
		hourly.2013-05-10_1205	72KB	0%	31%
		hourly.2013-05-10_1305	72KB	0%	31%
		hourly.2013-05-10_1405	72KB	0%	31%
		hourly.2013-05-10_1505	72KB	0%	31%

10 entries were displayed.

2. Use the `volume snapshot delete` command to delete all the Snapshot copies.

Example

```
volume snapshot delete -vserver vs3 -volume vol3 *
```

```
10 entries were acted on.
```

Using FlexClone LUNs to protect your data

A FlexClone LUN is a point-in-time, writeable copy of another LUN in an active volume or in a Snapshot copy. The clone and its parent can be modified independently without affecting each other.

A FlexClone LUN shares space initially with its parent LUN. By default, the FlexClone LUN inherits the space-reserved attribute of the parent LUN. For example, if the parent LUN is non-space-reserved, the FlexClone LUN is also non-space-reserved by default. However, you can create a non-space-reserved FlexClone LUN from a parent that is a space-reserved LUN.

When you clone a LUN, block sharing occurs in the background and you cannot create a volume Snapshot copy until the block sharing is finished.

You must configure the volume to enable the FlexClone LUN automatic deletion function with the `volume snapshot autodelete modify` command. Otherwise, if you want FlexClone LUNs to be deleted automatically but the volume is not configured for FlexClone auto delete, none of the FlexClone LUNs are deleted.

When you create a FlexClone LUN, the FlexClone LUN automatic deletion function is disabled by default. You must manually enable it on every FlexClone LUN before that FlexClone LUN can be automatically deleted. If you are using semi-thick volume provisioning and you want the “best effort” write guarantee provided by this option, you must make *all* FlexClone LUNs available for automatic deletion.

Note: When you create a FlexClone LUN from a Snapshot copy, the LUN is automatically split from the Snapshot copy by using a space-efficient background process so that the LUN does not continue to depend on the Snapshot copy or consume any additional space. If this background split has not been completed and this Snapshot copy is automatically deleted, that FlexClone LUN is deleted even if you have disabled the FlexClone auto delete function for that FlexClone LUN. After the background split is complete, the FlexClone LUN is not deleted even if that Snapshot copy is deleted.

Related information

[Clustered Data ONTAP 8.3 Logical Storage Management Guide](#)

Reasons for using FlexClone LUNs

You can use FlexClone LUNs to create multiple read/write copies of a LUN.

You might want to do this for the following reasons:

- You need to create a temporary copy of a LUN for testing purposes.
- You need to make a copy of your data available to additional users without giving them access to the production data.
- You want to create a clone of a database for manipulation and projection operations, while preserving the original data in an unaltered form.
- You want to access a specific subset of a LUN's data (a specific logical volume or file system in a volume group, or a specific file or set of files in a file system) and copy it to the original LUN, without restoring the rest of the data in the original LUN. This works on operating systems that support mounting a LUN and a clone of the LUN at the same time. SnapDrive for UNIX supports this with the `snap connect` command.
- You need multiple SAN boot hosts with the same operating system.

How a FlexVol volume can reclaim free space from FlexClone files and FlexClone LUNs

You can configure the autodelete settings of a FlexVol volume to automatically delete FlexClone files and FlexClone LUNs that have autodelete enabled when a volume is nearly full to reclaim a target amount of free space in the volume.

You can configure a volume to automatically start deleting FlexClone files and FlexClone LUNs when the free space in the volume decreases below a particular threshold value and automatically stop deleting clones when a target amount of free space in the volume is reclaimed. Although you cannot specify the threshold value that starts the automatic deletion of clones, you can specify whether a clone is eligible for deletion, and you can specify the target amount of free space for a volume to reclaim by deleting clones.

A volume automatically deletes FlexClone files and FlexClone LUNs when the free space in the volume decreases below a particular threshold and when *both* the following requirements are met:

- The autodelete capability is enabled for the volume that contains the FlexClone files and FlexClone LUNs.
You can enable the autodelete capability for a FlexVol volume by using the `volume snapshot autodelete modify` command. You must set the `-trigger` parameter to **volume** or **snap_reserve** for a volume to automatically delete FlexClone files and FlexClone LUNs.
- The autodelete capability is enabled for the FlexClone files and FlexClone LUNs.
You can enable autodelete for a FlexClone file or FlexClone LUN by using the `file clone create` command with the `-autodelete` parameter. As a result, you can preserve certain FlexClone files and FlexClone LUNs by disabling autodelete for the clones and ensuring that other volume settings do not override the clone setting.

You can specify the autodelete setting for FlexClone LUNs created using Data ONTAP 8.2 and later, and you can specify the autodelete setting for FlexClone files created using Data ONTAP 8.3 and later. After upgrading to Data ONTAP 8.3, FlexClone files created using Data ONTAP versions earlier than 8.3 have autodelete disabled. You can enable autodelete for the FlexClone files by using the `volume file clone autodelete` command.

Configuring a FlexVol volume to automatically delete FlexClone files and FlexClone LUNs

You can enable a FlexVol volume to automatically delete FlexClone files and FlexClone LUNs with autodelete enabled when the free space in the volume decreases below a particular threshold.

Before you begin

- The FlexVol volume must contain FlexClone files and FlexClone LUNs and be online.
- The FlexVol volume must not be a read-only volume.

Steps

1. Enable automatic deletion of FlexClone files and FlexClone LUNs in the FlexVol volume by using the `volume snapshot autodelete modify` command.
 - For the `-trigger` parameter, you can specify **volume** or **snap_reserve**.
 - For the `-destroy-list` parameter, you must always specify **lun_clone,file_clone** regardless of whether you want to delete only one type of clone.

Example

The following example shows how you can enable volume `vol1` to trigger the automatic deletion of FlexClone files and FlexClone LUNs for space reclamation until 25% of the volume consists of free space:

```
cluster1::> volume snapshot autodelete modify -vserver vs1 -volume
vol1 -enabled true -commitment disrupt -trigger volume -target-free-
space 25 -destroy-list lun_clone,file_clone

Volume modify successful on volume:vol1
```

Note: While enabling FlexVol volumes for automatic deletion, if you set the value of the `-commitment` parameter to **destroy**, all the FlexClone files and FlexClone LUNs with the `-autodelete` parameter set to **true** might be deleted when the free space in the volume decreases below the specified threshold value. However, FlexClone files and FlexClone LUNs with the `-autodelete` parameter set to **false** will not be deleted.

2. Verify that automatic deletion of FlexClone files and FlexClone LUNs is enabled in the FlexVol volume by using the `volume snapshot autodelete show` command.

Example

The following example shows that volume `vol1` is enabled for automatic deletion of FlexClone files and FlexClone LUNs:

```
cluster1::> volume snapshot autodelete show -vserver vs1 -volume vol1

Vserver Name: vs1
Volume Name: vol1
Enabled: true
Commitment: disrupt
Defer Delete: user_created
Delete Order: oldest_first
Defer Delete Prefix: (not specified)
Target Free Space: 25%
Trigger: volume
Destroy List: lun_clone,file_clone
Is Constituent Volume: false
```

3. Ensure that autodelete is enabled for the FlexClone files and FlexClone LUNs in the volume that you want to delete by performing the following steps:
 - a. Enable automatic deletion of a particular FlexClone file or FlexClone LUN by using the `volume file clone autodelete` command.

You can force a specific FlexClone file or FlexClone LUN to be automatically deleted by using the `volume file clone autodelete` command with the `-force` parameter.

Example

The following example shows that automatic deletion of the FlexClone LUN `lun1_clone` contained in volume `vol1` is enabled:

```
cluster1::> volume file clone autodelete -vserver vs1 -clone-
path /vol/vol1/lun1_clone -enabled true
```

You can enable autodelete when you create FlexClone files and FlexClone LUNs.

- b. Verify that the FlexClone file or FlexClone LUN is enabled for automatic deletion by using the `volume file clone show-autodelete` command.

Example

The following example shows that the FlexClone LUN `lun1_clone` is enabled for automatic deletion:

```
cluster1::> volume file clone show-autodelete -vserver vs1 -clone-
path vol/vol1/lun1_clone

Vserver Name: vs1
Path: vol/vol1/lun1_clone
Autodelete Enabled: true
Clone
```

For more information about using the commands, see the respective man pages.

Cloning LUNs from an active volume

You can create copies of your LUNs by cloning the LUNs in the active volume. These FlexClone LUNs are readable and writeable copies of the original LUNs in the active volume.

Before you begin

A FlexClone license must be installed.

About this task

Note: A space-reserved FlexClone LUN requires as much space as the space-reserved parent LUN. If the FlexClone LUN is not space-reserved, you must ensure that the volume has enough space to accommodate changes to the FlexClone LUN.

Steps

1. Use the `lun show` command to verify that the LUN exists.

Example

```
lun show -vserver vs1
```

Vserver	Path	State	Mapped	Type	Size
vs1	/vol/vol1/lun1	online	unmapped	windows	47.07MB

2. Use the `volume file clone create` command to create the FlexClone LUN.

Example

```
volume file clone create -vserver vs1 -volume vol1 -source-path lun1 -
destination-path /lun1_clone
```

If you need the FlexClone LUN to be available for automatic deletion, you include `-autodelete true`. If you are creating this FlexClone LUN in a volume using semi-thick provisioning, you must enable automatic deletion for all FlexClone LUNs.

3. Use the `lun show` command to verify that you created a LUN.

Example

```
lun show -vserver vs1
```

Vserver	Path	State	Mapped	Type	Size
vs1	/vol/volX/lun1	online	unmapped	windows	47.07MB
vs1	/vol/volX/lun1_clone	online	unmapped	windows	47.07MB

Related concepts

[Reasons for using FlexClone LUNs](#) on page 53

[Using FlexClone LUNs to protect your data](#) on page 52

Creating FlexClone LUNs from a Snapshot copy in a volume

You can use a Snapshot copy in your volume to create FlexClone copies of your LUNs. FlexClone copies of LUNs are both readable and writeable.

Before you begin

A FlexClone license must be installed.

About this task

The FlexClone LUN inherits the space reservations attribute of the parent LUN. A space-reserved FlexClone LUN requires as much space as the space-reserved parent LUN. If the FlexClone LUN is not space-reserved, the volume must have enough space to accommodate changes to the clone.

Steps

1. Create a Snapshot copy of the volume that contains the LUNs:

```
volume snapshot create -vserver vserver_name -volume volume_name -
snapshot snapshot_name
```

You must create a Snapshot copy (the backing Snapshot copy) of the LUN you want to clone.

2. Create the FlexClone LUN from the Snapshot copy:

```
file clone create -vserver vserver_name -volume volume_name -source-path
source_path -snapshot-name snapshot_name -destination-path
destination_path
```

If you need the FlexClone LUN to be available for automatic deletion, you include `-autodelete true`. If you are creating this FlexClone LUN in a volume using semi-thick provisioning, you must enable automatic deletion for all FlexClone LUNs.

3. Verify that the FlexClone LUN is correct:

```
lun show -vserver vserver_name
```

Vserver	Path	State	Mapped	Type	Size
vs1	/vol/vol1/lun1_clone	online	unmapped	windows	47.07MB
vs1	/vol/vol1/lun1_snap_clone	online	unmapped	windows	47.07MB

Related concepts

[Reasons for using FlexClone LUNs](#) on page 53

[Using FlexClone LUNs to protect your data](#) on page 52

Preventing a specific FlexClone file or FlexClone LUN from being automatically deleted

If you configure a FlexVol volume to automatically delete FlexClone files and FlexClone LUNs, any clone that fits the criteria you specify might be deleted. If you have specific FlexClone files or FlexClone LUNs that you want to preserve, you can exclude them from the automatic FlexClone deletion process.

Before you begin

A FlexClone license must be installed.

About this task

Starting with Data ONTAP 8.3, when you create a FlexClone file or FlexClone LUN, by default the autodelete setting for the clone is disabled. FlexClone files and FlexClone LUNs with autodelete disabled are preserved when you configure a FlexVol volume to automatically delete clones to reclaim space on the volume.

Attention: If you set the `commitment` level on the volume to **try** or **disrupt**, you can individually preserve specific FlexClone files or FlexClone LUNs by disabling autodelete for those clones. However, if you set the `commitment` level on the volume to **destroy** and the destroy lists include **lun_clone**, **file_clone**, the volume setting overrides the clone setting, and all FlexClone files and FlexClone LUNs can be deleted regardless of the autodelete setting for the clones.

Steps

1. Prevent a specific FlexClone file or FlexClone LUN from being automatically deleted by using the `volume file clone autodelete` command.

Example

The following example shows how you can disable autodelete for FlexClone LUN `lun1_clone` contained in `vol1`:

```
cluster1::> volume file clone autodelete -vserver vs1 -volume vol1 -
clone-path lun1_clone -enable false
```

A FlexClone file or FlexClone LUN with autodelete disabled cannot be deleted automatically to reclaim space on the volume.

2. Verify that autodelete is disabled for the FlexClone file or FlexClone LUN by using the `volume file clone show-autodelete` command.

Example

The following example shows that autodelete is false for the FlexClone LUN `lun1_clone`:

```
cluster1::> volume file clone show-autodelete -vserver vs1 -clone-
path vol/vol1/lun1_clone
Name: vs1
Path: vol/vol1/lun1_clone
Autodelete Enabled: false
```

Configuring and using SnapVault backups in a SAN environment

SnapVault configuration and use in a SAN environment is very similar to configuration and use in a NAS environment, but restoring LUNs in a SAN environment requires some special procedures.

SnapVault backups contain a set of read-only copies of a source volume. In a SAN environment you always back up entire volumes to the SnapVault secondary volume, not individual LUNs.

The procedure for creating and initializing the SnapVault relationship between a primary volume containing LUNs and a secondary volume acting as a SnapVault backup is identical to the procedure used with FlexVol volumes used for file protocols. This procedure is described in detail in the *Clustered Data ONTAP Data Protection Guide*.

It is important to ensure that LUNs being backed up are in a consistent state before the Snapshot copies are created and copied to the SnapVault secondary volume. Automating the Snapshot copy creation with a product like SnapManager for Microsoft SQL Server ensures that backed up LUNs are complete and usable by the original application.

There are three basic choices for restoring LUNs from a SnapVault secondary volume:

- You can map a LUN directly from the SnapVault secondary volume and connect a host to the LUN to access the contents of the LUN.
The LUN is read-only and you can map only from the most recent Snapshot copy in the SnapVault backup. Persistent reservations and other LUN metadata are lost. If desired, you can use a copy program on the host to copy the LUN contents back to the original LUN if it is still accessible.
The LUN has a different serial number from the source LUN.
- You can clone any Snapshot copy in the SnapVault secondary volume to a new read-write volume.
You can then map any of the LUNs in the volume and connect a host to the LUN to access the contents of the LUN. If desired, you can use a copy program on the host to copy the LUN contents back to the original LUN if it is still accessible.
- You can restore the entire volume containing the LUN from any Snapshot copy in the SnapVault secondary volume.
Restoring the entire volume replaces all of the LUNs, and any files, in the volume. Any new LUNs created since the Snapshot copy was created are lost.
The LUNs retain their mapping, serial numbers, UUIDs, and persistent reservations.

Accessing a read-only LUN copy from a SnapVault backup

You can access a read-only copy of a LUN from the latest Snapshot copy in a SnapVault backup. The LUN ID, path, and serial number are different from the source LUN and must first be mapped. Persistent reservations, LUN mappings, and igroups are not replicated to the SnapVault secondary volume.

Before you begin

- The SnapVault relationship must be initialized and the latest Snapshot copy in the SnapVault secondary volume must contain the desired LUN.
- The Storage Virtual Machine (SVM) containing the SnapVault backup must have one or more LIFs with the desired SAN protocol accessible from the host used to access the LUN copy.
- If you plan to access LUN copies directly from the SnapVault secondary volume, you must create your igroups on the SnapVault SVM in advance.
You can access a LUN directly from the SnapVault secondary volume without having to first restore or clone the volume containing the LUN.

About this task

If a new Snapshot copy is added to the SnapVault secondary volume while you have a LUN mapped from a previous Snapshot copy, the contents of the mapped LUN changes. The LUN is still mapped with the same identifiers, but the data is taken from the new Snapshot copy. If the LUN size changes, some hosts automatically detect the size change; Windows hosts require a disk rescan to pick up any size change.

Steps

1. Run the `lun show` command to list the available LUNs in the SnapVault secondary volume.

Example

In this example, you can see both the original LUNs in the primary volume srcvolA and the copies in the SnapVault secondary volume dstvolB:

```
cluster::> lun show

Vserver      Path                               State  Mapped  Type           Size
-----
vserverA     /vol/srcvolA/lun_A                online mapped   windows_2008  300.0GB
vserverA     /vol/srcvolA/lun_B                online mapped   windows_2008  300.0GB
vserverA     /vol/srcvolA/lun_C                online mapped   windows_2008  300.0GB
vserverB     /vol/dstvolB/lun_A                online unmapped windows_2008  300.0GB
vserverB     /vol/dstvolB/lun_B                online unmapped windows_2008  300.0GB
vserverB     /vol/dstvolB/lun_C                online unmapped windows_2008  300.0GB

6 entries were displayed.
```

2. If the igroup for the desired host does not already exist on the SVM containing the SnapVault secondary volume, run the `igroup create` command to create an igroup.

Example

This command creates an igroup for a Windows host that uses the iSCSI protocol:

```
cluster::> igroup create -vserver vserverB -igroup temp_igroup
               -protocol iscsi -ostype windows
               -initiator iqn.1991-05.com.microsoft:hostA
```

3. Run the `lun mapping create` command to map the desired LUN copy to the igroup.

Example

```
cluster::> lun mapping create -vserver vserverB -path /vol/dstvolB/
lun_A
               -igroup temp_igroup
```

4. Connect the host to the LUN and access the contents of the LUN as desired.

Restoring a single LUN from a SnapVault backup

You can restore a single LUN to a new location or to the original location. You can restore from any Snapshot copy in the SnapVault secondary volume. To restore the LUN to the original location, you first restore it to a new location and then copy it.

Before you begin

The SnapVault relationship must be initialized and the SnapVault secondary volume must contain an appropriate Snapshot copy to restore.

The Storage Virtual Machine (SVM) containing the SnapVault secondary volume must have one or more LIFs with the desired SAN protocol that are accessible from the host used to access the LUN copy.

The igroups must already exist on the SnapVault SVM in advance.

About this task

The process includes creating a read-write volume clone from a Snapshot copy in the SnapVault secondary volume. You can use the LUN directly from the clone, or you can optionally copy the LUN contents back to the original LUN location.

The LUN in the clone has a different path and serial number from the original LUN. Persistent reservations are not retained.

Steps

1. Run the `snapmirror show` command to verify the secondary volume that contains the SnapVault backup.

Example

```
cluster::> snapmirror show
```

Source Path	Type	Dest Path	Mirror State	Relation Status	Total Progress	Healthy	Last Updated
vserverA:srcvolA	XDP	vserverB:dstvolB	Snapmirrored	Idle	-	true	-

2. Run the `volume snapshot show` command to identify the Snapshot copy that you want to restore the LUN from.

Example

```
cluster::> volume snapshot show
```

Vserver	Volume	Snapshot	State	Size	Total%	Used%
vserverB	dstvolB	daily.2013-02-10_0010	valid	124KB	0%	0%
		weekly.2013-02-10_0015	valid	112KB	0%	0%
		daily.2013-02-11_0010	valid	164KB	0%	0%

3. Run the `volume clone create` command to create a read-write clone from the desired Snapshot copy.

The volume clone is created in the same aggregate as the SnapVault backup. There must be enough space in the aggregate to store the clone.

Example

```
cluster::> volume clone create -vserver vserverB
  -flexclone dstvolB_clone -type RW -parent-volume dstvolB
  -parent-snapshot daily.2013-02-10_0010
[Job 108] Job succeeded: Successful
```

4. Run the `lun show` command to list the LUNs in the volume clone.

```
cluster::> lun show -vserver vserverB -volume dstvolB_clone
```

Vserver	Path	State	Mapped	Type
vserverB	/vol/dstvolB_clone/lun_A	online	unmapped	windows_2008
vserverB	/vol/dstvolB_clone/lun_B	online	unmapped	windows_2008
vserverB	/vol/dstvolB_clone/lun_C	online	unmapped	windows_2008

3 entries were displayed.


```
vserverA:srcvolA
      XDP vserverB:dstvolB
          Snapmirrored
              Idle          -          true  -
```

3. Run the `volume snapshot show` command to identify the Snapshot copy that you want to restore from.

Example

```
cluster::> volume snapshot show

Vserver  Volume  Snapshot                               State Size  Total% Used%
-----  -
vserverB  dstvolB
          daily.2013-02-10_0010  valid  124KB   0%    0%
          weekly.2013-02-10_0015  valid  112KB   0%    0%
          daily.2013-02-11_0010  valid  164KB   0%    0%
```

4. Run the `snapmirror restore` command and specify the `-source-snapshot` option to specify the Snapshot copy to use.

The destination you specify for the restore is the original volume you are restoring to.

Example

```
cluster::> snapmirror restore -destination-path vserverA:srcvolA
      -source-path vserverB:dstvolB -source-snapshot daily.2013-02-10_0010

Warning: All data newer than Snapshot copy hourly.2013-02-11_1205 on
volume vserverA:src_volA will be deleted.
Do you want to continue? {y|n}: y
[Job 98] Job is queued: snapmirror restore from source
"vserverB:dstvolB" for the snapshot daily.2013-02-10_0010.
```

5. If you are sharing LUNs across a host cluster, restore the persistent reservations on the LUNs from the affected hosts.

Restoring a volume from a SnapVault backup

In the following example, the LUN named `lun_D` was added to the volume after the Snapshot copy was made. After restoring the entire volume from the Snapshot copy, `lun_D` no longer appears.

In the `lun show` command output, you can see the LUNs in the primary volume `srcvolA` and the read-only copies of those LUNs in the SnapVault secondary volume `dstvolB`. There is no copy of `lun_D` in the SnapVault backup.

```
cluster::> lun show
Vserver  Path                               State  Mapped  Type          Size
-----  -
vserverA  /vol/srcvolA/lun_A                 online mapped   windows_2008  300.0GB
vserverA  /vol/srcvolA/lun_B                 online mapped   windows_2008  300.0GB
vserverA  /vol/srcvolA/lun_C                 online mapped   windows_2008  300.0GB
vserverA  /vol/srcvolA/lun_D                 online mapped   windows_2008  250.0GB
vserverB  /vol/dstvolB/lun_A                 online unmapped windows_2008  300.0GB
vserverB  /vol/dstvolB/lun_B                 online unmapped windows_2008  300.0GB
vserverB  /vol/dstvolB/lun_C                 online unmapped windows_2008  300.0GB
```

7 entries were displayed.

```
cluster::> snapmirror restore -destination-path vserverA:srcvolA
      -source-path vserverB:dstvolB
```

```

-source-snapshot daily.2013-02-10_0010

Warning: All data newer than Snapshot copy hourly.2013-02-11_1205
on volume vserverA:src_volA will be deleted.
Do you want to continue? {y|n}: y
[Job 98] Job is queued: snapmirror restore from source
"vserverB:dstvolB" for the snapshot daily.2013-02-10_0010.

cluster::> lun show
Vserver  Path                               State  Mapped  Type           Size
-----  -
vserverA /vol/srcvolA/lun_A                 online mapped  windows_2008  300.0GB
vserverA /vol/srcvolA/lun_B                 online mapped  windows_2008  300.0GB
vserverA /vol/srcvolA/lun_C                 online mapped  windows_2008  300.0GB
vserverB /vol/dstvolB/lun_A                 online unmapped windows_2008  300.0GB
vserverB /vol/dstvolB/lun_B                 online unmapped windows_2008  300.0GB
vserverB /vol/dstvolB/lun_C                 online unmapped windows_2008  300.0GB

6 entries were displayed.


```

After the volume is restored from the SnapVault secondary volume, the source volume no longer contains lun_D. You do not need to remap the LUNs in the source volume after the restore; they are still mapped.

How you can connect a host backup system to the primary storage system

You can back up SAN systems to tape through a separate backup host to avoid performance degradation on the application host.

It is imperative that you keep SAN and NAS data separated for backup purposes. The figure below shows the recommended physical configuration for a host backup system to the primary storage system. You must configure volumes as SAN-only, and you must also configure qtrees within a single volume as SAN-only. LUNs can be confined to a single volume or qtree. The LUNs also can be spread across multiple volumes, qtrees, or storage systems.



Volumes on a host can consist of a single LUN mapped from the storage system or multiple LUNs using a volume manager, such as VxVM on HP-UX systems.

Related tasks

[Backing up a LUN through a host backup system](#) on page 65

Related information

[Clustered Data ONTAP 8.3 Data Protection Guide](#)

Backing up a LUN through a host backup system

You can use a cloned LUN from a Snapshot copy as source data for the host backup system.

Before you begin

A production LUN must exist and be mapped to an igroup that includes the WWPN or initiator node name of the application server. The LUN must also be formatted and accessible to the host

Steps

1. Save the contents of the host file system buffers to disk.

You can use the command provided by your host operating system, or you can use SnapDrive for Windows or SnapDrive for UNIX. You can also opt to make this step part of your SAN backup pre-processing script.

2. Use the `volume snapshot create` command to create a Snapshot copy of the production LUN.

Example

```
volume snapshot create -vserver vs0 -volume vol3 -snapshot vol3_snapshot
-comment "Single snapshot" -foreground false
```

3. Use the `volume file clone create` command to create a clone of the production LUN.

Example

```
volume file clone create -vserver vs3 -volume vol3 -source-path lun1 -
snapshot-name snap_vol3 -destination-path lun1_backup
```

4. Use the `lun igroup create` command to create an igroup that includes the WWPN of the backup server.

Example

```
lun igroup create -vserver vs3 -igroup igroup3 -protocol fc -ostype
windows -initiator 10:00:00:00:c9:73:5b:91
```

5. Use the `lun mapping create` command to map the LUN clone you created in Step 3 to the backup host.

Example

```
lun mapping create -vserver vs3 -volume vol3 -lun lun1_backup -igroup
igroup3
```

You can opt to make this step part of your SAN backup application's post-processing script.

6. From the host, discover the new LUN and make the file system available to the host.
You can opt to make this step part of your SAN backup application's post-processing script.
7. Back up the data in the LUN clone from the backup host to tape by using your SAN backup application.
8. Use the `lun modify` command to take the LUN clone offline.

Example

```
lun modify -vserver vs3 -path /vol/vol3/lun1_backup -state offline
```

9. Use the `lun delete` to remove the LUN clone.

Example

```
lun delete -vserver vs3 -volume vol3 -lun lun1_backup
```

10. Use the `volume snapshot delete` command to remove the Snapshot copy.

Example

```
volume snapshot delete -vserver vs3 -volume vol3 -snapshot vol3_snapshot
```

Related concepts

[How you can connect a host backup system to the primary storage system](#) on page 64

Ways to implement SVM disaster recovery in SAN environments

Storage Virtual Machine (SVM) disaster recovery provides support for data recovery at the SVM level if an SVM becomes inaccessible.

The *primary* SVM is the SVM requiring disaster recovery support. The *secondary* SVM is an SVM in another cluster that is peered with the primary SVM to provide disaster recovery support. The cluster containing the secondary SVM does not have to have the same number of nodes or the same number of FC and Ethernet ports as the primary SVM.

Volume, configuration, and metadata is replicated to the secondary SVM at regular scheduled intervals. If the primary SVM becomes inaccessible, the secondary SVM can be brought online. Persistent reservations are not copied to the secondary SVM. This means that certain hosts must either be rebooted or have persistent reservations reset after the secondary SVM becomes active. The following hosts require a reboot or a persistent reservation reset:

- AIX
- Solaris
- Veritas

SVM disaster recovery can be implemented in two ways.

Identity preserve SVM disaster recovery

With identity preserve SVM disaster recovery, volumes, LIFs, and LUNs in the secondary SVM are not visible to hosts until the entire disaster recovery operation has been successfully completed.

The volumes, LIFs, and LUNs in the secondary SVM have the same identity as the corresponding volumes, LIFs, and LUNs in the primary SVM. Therefore, the primary SVM and secondary SVM cannot be visible to the host simultaneously.

Identity preserve SVM disaster recovery does not require FC port zoning.

Identity discard SVM disaster recovery

With identity discard SVM disaster recovery, the volumes, LIFs, and LUNs in the primary SVM remain visible to the host in read-only mode for the duration of the disaster recovery event.

The volumes, LIFs, and LUNs in the secondary SVM do not have the same identity as the corresponding volumes, LIFs, and LUNs in the primary SVM. Therefore, the primary SVM and secondary SVM can be visible to the host simultaneously.

To use identity discard SVM disaster recovery, the primary SVM FC LIFs must use WWPN zoning.

Related information

[Data ONTAP 8.3 SVM Disaster Recovery Express Guide](#)

[Data ONTAP 8.3 SVM Disaster Recovery Preparation Express Guide](#)

Managing your iSCSI service

You can perform various task for managing your iSCSI service, such as deleting iSCSI service from a particular SVM, enabling error recovery levels, and defining authentication methods.

Commands for managing iSCSI services

You can use iSCSI commands to manage iSCSI services on your Storage Virtual Machine (SVM).

- [Commands for managing iSCSI services](#) on page 68
- [Commands for managing storage system interfaces on iSCSI](#) on page 69
- [Commands for managing iSCSI interface access management](#) on page 69
- [Commands for managing iSCSI initiator security](#) on page 69
- [Commands for managing iSCSI sessions](#) on page 70
- [Commands for managing iSCSI connections](#) on page 70
- [Commands for managing initiators](#) on page 70

Commands for managing iSCSI services

If you want to....	Use this command....
Verify that the iSCSI service is running	<code>vserver iscsi show</code>
Display the iSCSI license	<code>license show</code>
Enable the iSCSI license	<code>license add</code>
Create an iSCSI service	<code>vserver iscsi create</code>
Start an iSCSI service	<code>vserver iscsi start</code>
Add a LIF	<code>network interface create</code>
Modify a LIF	<code>network interface modify</code>
Delete a LIF	<code>network interface delete</code> Note: If the LIF is in a port set, you must remove the LIF from the port set before you can delete the LIF. Use the <code>lun portset remove</code> command to remove the LIF from the port set.
Display a node name	<code>vserver iscsi show</code>
Display the target alias	<code>vserver iscsi show</code>

If you want to....	Use this command....
Change the target node name	<code>vserver iscsi modify</code> Note: You need to change the privilege to advanced mode to change the target node name. Note: You must stop the iSCSI service before you can change the node name.
Add or change the target alias	<code>vserver iscsi modify</code>
Disable iSCSI license	<code>license delete</code>
Stop an iSCSI service	<code>vserver iscsi stop</code>
Delete an iSCSI service	<code>vserver iscsi delete</code>
View a man page for a command	<code>man command</code>

Commands for managing storage system interfaces on iSCSI

If you want to....	Use this command....
Display the iSCSI logical interfaces for an SVM	<code>vserver iscsi interface show</code>
Enable a logical interface	<code>vserver iscsi interface enable</code>
Disable a logical interface	<code>vserver iscsi interface disable</code>
View a man page for a command	<code>man command</code>

Commands for managing iSCSI interface access management

If you want to....	Use this command....
Add an iSCSI LIF to an access list	<code>vserver iscsi interface accesslist add</code>
Display an access list for an initiator	<code>vserver iscsi interface accesslist show</code>
Remove an iSCSI LIF from an access list	<code>vserver iscsi interface accesslist remove</code>
View a man page for a command	<code>man command</code>

Commands for managing iSCSI initiator security

If you want to....	Use this command....
Configure a security policy method for an initiator	<code>vserver iscsi security create</code>
Generate a hexadecimal secret password	<code>vserver iscsi security generate</code>
Display default and manually created security policy method information	<code>vserver iscsi security show</code>

If you want to....	Use this command....
Modify an existing security policy method for an initiator	<code>vserver iscsi security modify</code>
Define the default security policy method	<code>vserver iscsi security default</code>
Delete the security policy for an initiator	<code>vserver iscsi security delete</code>
View a man page for a command	<code>man <i>command</i></code>

Commands for managing iSCSI sessions

If you want to....	Use this command....
Display the status of active iSCSI commands	<code>vserver iscsi command show</code>
Display iSCSI session information	<code>vserver iscsi session show</code>
Display session parameter information	<code>vserver iscsi session parameter show</code>
Shut down all connections in a session	<code>vserver iscsi session shutdown</code> Note: You need to change the privilege to advanced mode to shut down all connections in a session.
View a man page for a command	<code>man <i>command</i></code>

Commands for managing iSCSI connections

If you want to....	Use this command....
Display iSCSI connection information	<code>vserver iscsi connection show</code>
Shut down a connection in a session	<code>vserver iscsi connection shutdown</code> Note: You need to change the privilege to advanced mode to shut down a connection in a session.
View a man page for a command	<code>man <i>command</i></code>

Commands for managing initiators

If you want to....	Use this command....
Display a list of active initiators	<code>vserver iscsi initiator show</code>
View a man page for a command	<code>man <i>command</i></code>

Configuring your network for best performance

Ethernet networks vary greatly in performance. You can maximize the performance of the network used for iSCSI by selecting specific configuration values.

Steps

1. Connect the host and storage ports to the same network.
It is best to connect to the same switches. Routing should never be used.
2. Select the highest speed ports available, and dedicate them to iSCSI.
10 GbE ports are best. 1 GbE ports are the minimum.
3. Disable Ethernet flow control for all ports.
4. Enable jumbo frames (typically MTU of 9000).
All devices in the data path, including initiators, targets, and switches, must support jumbo frames. Otherwise, enabling jumbo frames actually reduces network performance substantially.

Defining a security policy method for an initiator

You can define a list of initiators and their authentication methods. You can also modify the default authentication method that applies to initiators that do not have a user-defined authentication method.

About this task

You can generate unique passwords using security policy algorithms in the product or you can manually specify the passwords that you want to use.

Note: Not all initiators support hexadecimal CHAP secret passwords.

Steps

1. Use the `vserver iscsi security create` command to create a security policy method for an initiator.

Example

```
vserver iscsi security create -vserver vs2 -initiator iqn.1991-05.com.microsoft:host1 -auth-type CHAP -user-name bob1 -outbound-user-name bob2
```

2. Follow the screen commands to add the passwords.
Creates a security policy method for initiator `iqn.1991-05.com.microsoft:host1` with inbound and outbound CHAP user names and passwords.

Related concepts

[How iSCSI authentication works](#) on page 73

[Guidelines for using CHAP authentication](#) on page 74

[What CHAP authentication is](#) on page 74

Deleting an iSCSI service for an SVM

You can delete an iSCSI service for a Storage Virtual Machine (SVM) if it is no longer required.

Before you begin

The administration status of the iSCSI service must be in the “down” state before you can delete an iSCSI service. You can move the administration status to down with the `vserver iscsi modify` command.

Steps

1. Use the `vserver iscsi modify` command to stop the I/O to the LUN.

Example

```
vserver iscsi modify -vserver vs1 -status-admin down
```

2. Use the `vserver iscsi delete` command to remove the iscsi service from the SVM.

Example

```
vserver iscsi delete -vserver vs_1
```

3. Use the `vserver iscsi show` command to verify that you deleted the iSCSI service from the SVM.

```
vserver iscsi show -vserver vs1
```

Getting more details in iSCSI session error recoveries

Increasing the iSCSI session error recovery level enables you to receive more detailed information about iSCSI error recoveries. Using a higher error recovery level might cause a minor degrade in iSCSI session performance.

About this task

By default, Data ONTAP is configured to use error recovery level 0 for iSCSI sessions. If you are using an initiator that has been qualified for error recovery level 1 or 2, you can choose to increase the error recovery level. The modified session error recovery level affects only the newly created sessions and does not affect existing sessions.

Steps

1. Enter advanced mode:


```
set -privilege advanced
```
2. Verify the current setting by using the `iscsi show` command.

Example

```
iscsi show -vserver vs3 -fields max-error-recovery-level
```

```
vserver max-error-recovery-level
-----
vs3      0
```


3. Change the error recovery level by using the `iscsi modify` command.

Example

```
iscsi modify -vserver vs3 -max-error-recovery-level 2
```

iSCSI service management

You can manage the availability of the iSCSI service on the iSCSI logical interfaces of the Storage Virtual Machine (SVM) by using the `vserver iscsi interface enable` or `vserver iscsi interface disable` commands.

By default, the iSCSI service is enabled on all iSCSI logical interfaces.

How iSCSI is implemented on the host

iSCSI can be implemented on the host using hardware or software.

You can implement iSCSI in one of the following ways:

- Using Initiator software that uses the host's standard Ethernet interfaces.
- Through an iSCSI host bus adapter (HBA): An iSCSI HBA appears to the host operating system as a SCSI disk adapter with local disks.
- Using a TCP Offload Engine (TOE) adapter that offloads TCP/IP processing. The iSCSI protocol processing is still performed by host software.

How iSCSI authentication works

During the initial stage of an iSCSI session, the initiator sends a login request to the storage system to begin an iSCSI session. The storage system will then either permit or deny the login request, or determine that a login is not required.

iSCSI authentication methods are:

- Challenge Handshake Authentication Protocol (CHAP)—The initiator logs in using a CHAP user name and password. You can specify a CHAP password or generate a hexadecimal secret password. There are two types of CHAP user names and passwords:
 - Inbound—The storage system authenticates the initiator. Inbound settings are required if you are using CHAP authentication.
 - Outbound—This is an optional setting to enable the initiator to authenticate the storage system. You can use outbound settings only if you defined an inbound user name and password on the storage system.
- deny—The initiator is denied access to the storage system.
- none—The storage system does not require authentication for the initiator.

You can define a list of initiators and their authentication methods. You can also define a default authentication method that applies to initiators that are not on this list.

The default iSCSI authentication method is `none`, which means any initiator not in the authentication list can log in to the storage system without authentication. However, you can change the default method to `deny` or `CHAP`.

Related concepts

[How Data ONTAP implements an iSCSI network](#) on page 98

Related information

[Data ONTAP documentation on the NetApp support website-media.netapp.com/documents/tr-3441.pdf](#)

iSCSI initiator security management

Data ONTAP provides a number of features for managing security for iSCSI initiators. You can define a list of iSCSI initiators and the authentication method for each, display the initiators and their associated authentication methods in the authentication list, add and remove initiators from the authentication list, and define the default iSCSI initiator authentication method for initiators not in the list.

What CHAP authentication is

The Challenge Handshake Authentication Protocol (CHAP) enables authenticated communication between iSCSI initiators and targets. When you use CHAP authentication, you define CHAP user names and passwords on both the initiator and the storage system.

During the initial stage of an iSCSI session, the initiator sends a login request to the storage system to begin the session. The login request includes the initiator's CHAP user name and CHAP algorithm. The storage system responds with a CHAP challenge. The initiator provides a CHAP response. The storage system verifies the response and authenticates the initiator. The CHAP password is used to compute the response.

Guidelines for using CHAP authentication

You should follow certain guidelines when using CHAP authentication.

- If you define an inbound user name and password on the storage system, you must use the same user name and password for outbound CHAP settings on the initiator. If you also define an outbound user name and password on the storage system to enable bidirectional authentication, you must use the same user name and password for inbound CHAP settings on the initiator.
- You cannot use the same user name and password for inbound and outbound settings on the storage system.
- CHAP user names can be 1 to 128 bytes.
A null user name is not allowed.
- CHAP passwords (secrets) can be 1 to 512 bytes.
Passwords can be hexadecimal values or strings. For hexadecimal values, you should enter the value with a prefix of "0x" or "0X". A null password is not allowed.
- For additional restrictions, you should see the initiator's documentation.
For example, the Microsoft iSCSI software initiator requires both the initiator and target CHAP passwords to be at least 12 bytes if IPsec encryption is not being used. The maximum password length is 16 bytes regardless of whether IPsec is used.

Related concepts

[What CHAP authentication is](#) on page 74

iSCSI interface access lists to limit initiator interfaces

iSCSI interface access lists can be used to limit the number of interfaces an initiator can access, thereby increasing performance and security.

When an initiator begins a discovery session using an iSCSI `SendTargets` command, it receives the IP addresses associated with network interfaces on its access list. By default, all initiators have access to all iSCSI interfaces. You can use the access list to restrict the number of interfaces an initiator has access to.

iSNS server registration requirement

If you decide to use an iSNS service, you must ensure that your Storage Virtual Machines (SVMs) are properly registered with an Internet Storage Name Service (iSNS) server.

What iSNS is

The Internet Storage Name Service (iSNS) is a protocol that enables automated discovery and management of iSCSI devices on a TCP/IP storage network. An iSNS server maintains information about active iSCSI devices on the network, including their IP addresses, iSCSI node names, and portal groups.

You can obtain an iSNS server from a third-party vendor. If you have an iSNS server on your network configured and enabled for use by the initiator and target, you can use the management LIF for a Storage Virtual Machine (SVM) to register all the iSCSI LIFs for that SVM on the iSNS server. After the registration is complete, the iSCSI initiator can query the iSNS server to discover all the LIFs for that particular SVM.

If you do not have an iSNS server on your network, you must manually configure each target to be visible to the host.

What an iSNS server does

An iSNS server uses the Internet Storage Name Service protocol to maintain information about active iSCSI devices on the network, including their IP addresses, iSCSI node names, and portal groups.

The iSNS protocol enables automated discovery and management of iSCSI devices on an IP storage network. An iSCSI initiator can query the iSNS server to discover iSCSI target devices.

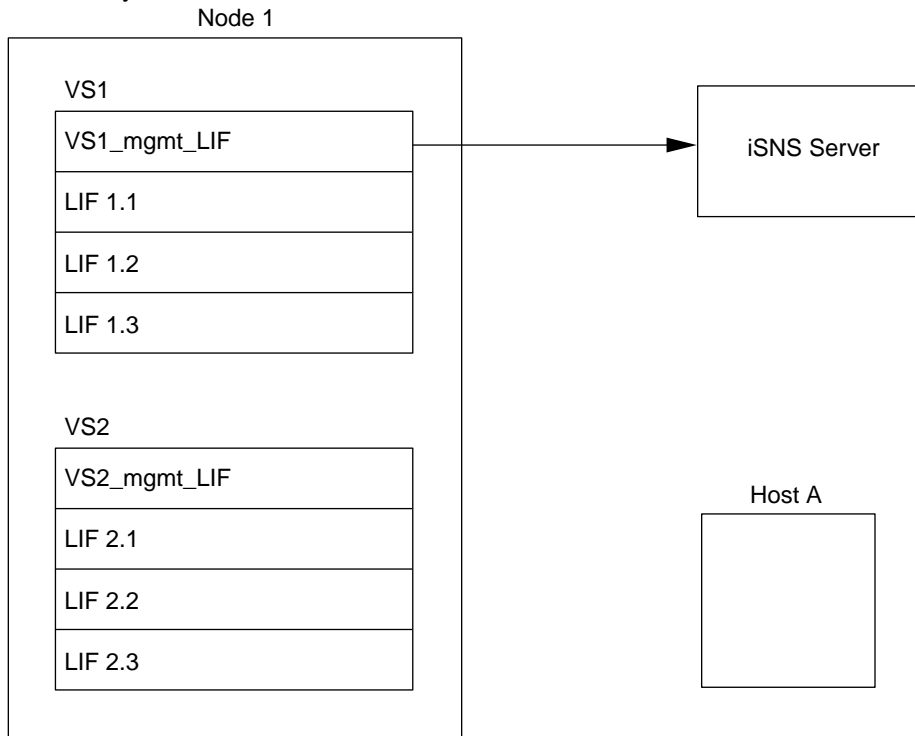
NetApp does not supply or resell iSNS servers. You obtain these servers from a vendor supported by NetApp. Be sure to check the NetApp iSCSI Support Matrix to see which iSNS servers are currently supported.

How SVMs interact with an iSNS server

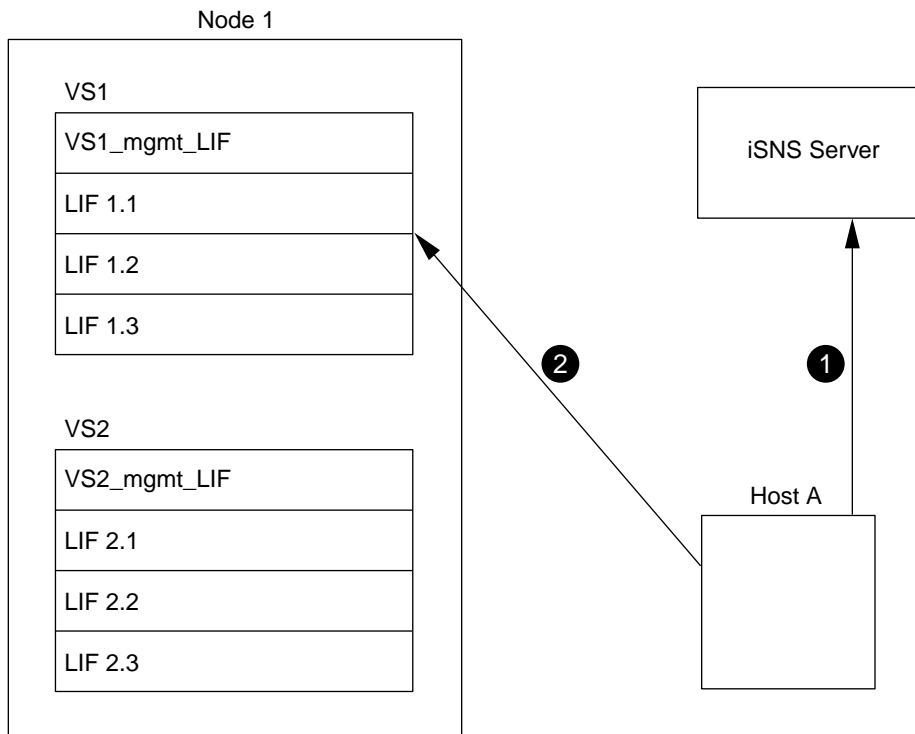
The iSNS server communicates with each Storage Virtual Machine (SVM) through the SVM management LIF. The management LIF registers all iSCSI target node name, alias, and portal information with the iSNS service for a specific SVM.

In the following example, SVM VS1 uses the SVM management LIF `VS1_mgmt_LIF` to register with the iSNS server. During iSNS registration, an SVM sends all the iSCSI LIFs through the SVM management LIF to the iSNS Server. After the iSNS registration is complete, the iSNS server has a list of all the LIFs serving iSCSI in VS1. If a node contains multiple SVMs, each SVM must register

individually with the iSNS server to use the iSNS service.



In the next example after the iSNS server completes the registration with the target, Host A can discover all the LIFs for VS1 through the iSNS server as indicated in step 1. After Host A completes the discovery of the LIFs for VS1, Host A can establish a connection with any of the LIFs in VS1 as shown in step 2. Host A is not aware of any of the LIFs in VS2 until the management LIF VS2_mgmt_LIF for VS2 registers with the iSNS server.



However, if you define the interface access lists, the host can only use the defined LIFs in the interface access list to access the target.

After iSNS is initially configured, Data ONTAP automatically updates the iSNS server any time the SVM configuration settings change.

A delay of a few minutes can occur between the time you make the configuration changes and when Data ONTAP sends the iSNS server the updates. You can use the `vserver iscsi isns update` command to force an immediate update of the iSNS information on the iSNS server.

Related tasks

[Registering the SVM with an iSNS server](#) on page 77

Related references

[Commands for managing iSNS](#) on page 78

About iSNS service version incompatibility

In Data ONTAP 8.1.1, the default iSNS version is draft 22. You cannot use previous iSNS versions. This draft is also used by Microsoft iSNS server 3.0.

Registering the SVM with an iSNS server

You can use the `vserver iscsi isns` command to configure the Storage Virtual Machine (SVM) to register with an iSNS server.

About this task

The `vserver iscsi isns create` command configures the SVM to register with the iSNS server. The SVM does not provide commands that enable you to configure or manage the iSNS server. To manage the iSNS server, you can use the server administration tools or the interface provided by the vendor for the iSNS server.

Steps

1. On your iSNS server, ensure that your iSNS service is up and available for service.
2. Create the SVM management LIF on a data port:


```
network interface create -vserver SVM_name -lif LIF_name -role data -
data-protocol none -home-node home_node_name -home-port home_port -
address IP_address -netmask network_mask
```
3. Create an iSCSI service on your SVM if one does not already exist:


```
vserver iscsi create -vserver SVM_name
```
4. Verify that the iSCSI service was created successfully:


```
iscsi show -vserver SVM_name
```
5. Verify that a default route exists for the SVM:


```
network route show -vserver SVM_name
```
6. If a default route does not exist for the SVM, create a default route:


```
network route create -vserver SVM_name -routing-group routing_group -
destination destination -gateway gateway
```
7. Configure the SVM to register with the iSNS service:


```
vserver iscsi isns create -vserver SVM_name -address IP_address
```

Both IPv4 and IPv6 address families are supported. The address family of the iSNS server must be the same as that of the SVM management LIF.

For example, you cannot connect an SVM management LIF with an IPv4 address to an iSNS server with an IPv6 address.

8. Verify that the iSNS service is running:

```
vserver iscsi isns show -vserver SVM_name
```

9. If the iSNS service is not running, start it:

```
vserver iscsi isns start -vserver SVM_name
```

Related concepts

[How SVMs interact with an iSNS server](#) on page 75

[What iSNS is](#) on page 75

Related references

[Commands for managing iSNS](#) on page 78

Commands for managing iSNS

Data ONTAP provides commands to manage your iSNS service.

If you want to...	Use this command...
Configure an iSNS service	<code>vserver iscsi isns create</code>
Start an iSNS service	<code>vserver iscsi isns start</code>
Modify an iSNS service	<code>vserver iscsi isns modify</code>
Display iSNS service configuration	<code>vserver iscsi isns show</code>
Force an update of registered iSNS information	<code>vserver iscsi isns update</code>
Stop an iSNS service	<code>vserver iscsi isns stop</code>
Remove an iSNS service	<code>vserver iscsi isns delete</code>
View the man page for a command	<code>man <i>command name</i></code>

See the man page for each command for more information.

Related concepts

[How SVMs interact with an iSNS server](#) on page 75

[What iSNS is](#) on page 75

Related tasks

[Registering the SVM with an iSNS server](#) on page 77

iSCSI troubleshooting tips

You can troubleshoot common problems that occur with iSCSI networks.

Troubleshooting iSCSI LUNs not visible on the host

The iSCSI LUNs appear as local disks to the host. If the storage system LUNs are not available as disks on the host, you should verify the configuration settings.

Configuration setting	What to do
Cabling	Verify that the cables between the host and the storage system are properly connected.
Network connectivity	Verify that there is TCP/IP connectivity between the host and the storage system. <ul style="list-style-type: none"> From the storage system command line, ping the host interfaces that are being used for iSCSI. From the host command line, ping the storage system interfaces that are being used for iSCSI.
System requirements	Verify that the components of your configuration are qualified. Also, verify that you have the correct host operating system (OS) service pack level, initiator version, Data ONTAP version, and other system requirements. You can check the most up-to-date system requirements at Interoperability Matrix.
Jumbo frames	If you are using jumbo frames in your configuration, verify that jumbo frames are enabled on all devices in the network path: the host Ethernet NIC, the storage system, and any switches.
iSCSI service status	Verify that the iSCSI service is licensed and started on the storage system.
Initiator login	Verify that the initiator is logged in to the storage system. If the command output shows no initiators are logged in, check the initiator configuration on the host. Also verify that the storage system is configured as a target of the initiator.
iSCSI node names	Verify that you are using the correct initiator node names in the igroup configuration. On the host, you can use the initiator tools and commands to display the initiator node name. The initiator node names configured in the igroup and on the host must match.
LUN mappings	Verify that the LUNs are mapped to an igroup. On the storage system console, you can use one of the following commands: <ul style="list-style-type: none"> <code>lun mapping show</code> displays all LUNs and the igroups to which they are mapped. <code>lun mapping show -igroup</code> displays the LUNs mapped to a specific igroup.
iSCSI LIFs enable	Verify that the iSCSI logical interfaces are enabled.

Related information

[NetApp Interoperability Matrix Tool](#)

Resolving iSCSI error messages on the storage system

There are a number of common iSCSI-related error messages that you can view with the `event log show` command. You need to know what these messages mean and what you can do to resolve the issues they identify.

The following table contains the most common error messages, and instructions for resolving them:

Message	Explanation	What to do
ISCSI: network interface <i>identifier</i> disabled for use; incoming connection discarded	The iSCSI service is not enabled on the interface.	You can use the <code>iscsi interface enable</code> command to enable the iSCSI service on the interface. For example: <code>iscsi interface enable -vserever vs1 -lif lif1</code>
ISCSI: Authentication failed for initiator <i>nodename</i>	CHAP is not configured correctly for the specified initiator.	Check CHAP settings; you cannot use the same user name and password for inbound and outbound settings on the storage system: <ul style="list-style-type: none"> • Inbound credentials on the storage system must match outbound credentials on the initiator. • Outbound credentials on the storage system must match inbound credentials on the initiator.

Related concepts

[Guidelines for using CHAP authentication](#) on page 74

Managing your FC service

You can start, verify, stop, or delete an FC service. You can also add LIFs, create WWPN aliases, and display FC logical interface information.

Commands for managing FC protocols

You can use FC commands to manage FC protocols on your Storage Virtual Machine (SVM).

- [Commands for managing SVM FC services](#) on page 81
- [Commands for displaying active FC initiators](#) on page 81
- [Commands for managing FC logical interfaces](#) on page 82
- [Commands for managing FC initiator WWPN aliases](#) on page 82

Commands for managing SVM FC services

If you want to....	Use this command...
Verify that the FC service is running	<code>vserver fcp status</code> or <code>vserver fcp show</code>
Verify the FC license	<code>license show</code>
Enable the FC license	<code>license add</code>
Start an FC service	<code>vserver fcp start</code>
Enable an FC service	<code>vserver fcp create</code>
Modify an FC service	<code>vserver fcp modify</code>
Add a LIF	<code>network interface create</code>
Modify a LIF	<code>network interface modify</code>
Delete a LIF	<code>network interface delete</code>
Modify a target name	<code>vserver fcp modify</code>
Disable the FC license	<code>license delete</code>
Stop an FC service	<code>vserver fcp stop</code>
Delete an FC service	<code>vserver fcp delete</code>
View a man page for a command	<code>man command _name</code>

Commands for displaying active FC initiators

If you want to....	Use this command...
Display information about FC initiators	<code>vserver fcp initiator show</code>
View a man page for a command	<code>man command _name</code>

Commands for managing FC logical interfaces

If you want to....	Use this command...
Display FC logical interface information	<code>vserver fcp interface show</code>
Assign a new WWPN to a logical interface	<code>vserver fcp portname set</code>
Display the WWPN used by the FC logical interfaces	<code>vserver fcp portname show</code>
View a man page for a command	<code>man <i>command_name</i></code>

Commands for managing FC initiator WWPN aliases

If you want to....	Use this command...
Create a WWPN alias name	<code>vserver fcp wwpn-alias set</code>
Modify a WWPN alias name	<code>vserver fcp wwpn-alias set</code>
Display WWPN alias information	<code>vserver fcp wwpn-alias show</code>
Remove a WWPN alias name	<code>vserver fcp wwpn-alias remove</code>
View a man page for a command	<code>man <i>command_name</i></code>

Changing the WWPN for an FC logical interface

Data ONTAP automatically assigns the World Wide Port Numbers (WWPNs) for all FC logical interfaces when they are created. However, there are some circumstances in which you might need to change the WWPN assignments on your FC logical interfaces.

About this task

For instance, if you are replacing an existing storage system, you might want to reuse the existing WWPNs to minimize the impact on the existing configuration.

Steps

1. Use the `set -privilege advanced` command to change the privilege to advanced.
2. Use the `network interface modify` command to take the logical interfaces offline.

Example

```
network interface modify -vserver vs3 -lif lif1 -status-admin down
```

3. Use the `vserver fcp portname set` command to change the WWPN of your logical interface.

Example

```
vserver fcp portname set -vserver vs3 -lif lif1 -wwpn
20:09:00:a0:98:27:db:a3
```

4. Use the `network interface modify` command to bring the logical interfaces online.

Example

```
network interface modify -vserver vs3 -lif lif1 -status-admin up
```

Related concepts

[How worldwide name assignments work](#) on page 83

[How WWPNs are used](#) on page 101

Deleting an FC service for an SVM

You can delete an FC service for a Storage Virtual Machine (SVM) if it is no longer required.

Before you begin

The administration status must be “down” before you can delete a FC service for an SVM. You can set the administration status to down with either the `vserver fcp modify` command or the `vserver fcp stop` command.

Steps

1. Use the `vserver fcp stop` command to stop the I/O to the LUN.

Example

```
vserver fcp stop -vserver vs_1
```

2. Use the `vserver fcp delete` command to remove the service from the SVM.

Example

```
vserver fcp delete -vserver vs_1
```

3. Use the `vserver fcp show` to verify that you deleted the FC service from your SVM:

```
vserver fcp show -vserver vs_1
```

How worldwide name assignments work

Worldwide names are created sequentially in Data ONTAP. However, because of the way Data ONTAP assigns them, they might appear to be assigned in a non-sequential order.

Each adapter has a pre-configured WWPN and WWNN, but Data ONTAP does not use these pre-configured values. Instead, Data ONTAP assigns its own WWPNs or WWNNs, based on the MAC addresses of the onboard Ethernet ports.

You can change this value only after creation of the FC logical interface (with the `vserver fcp portname set` command).


The worldwide names might appear to be non-sequential when assigned for the following reasons:

- Worldwide names are assigned across all the nodes and Storage Virtual Machines (SVMs) in the cluster.
- Freed worldwide names are recycled and added back to the pool of available names.

How FC switches are identified

Fibre Channel switches have one worldwide node name (WWNN) for the device itself, and one worldwide port name (WWPN) for each of its ports.

For example, the following diagram shows how the WWPNs are assigned to each of the ports on a 16-port Brocade switch. For details about how the ports are numbered for a particular switch, see the vendor-supplied documentation for that switch.



Port **0**, WWPN 20:**00**:00:60:69:51:06:b4

Port **1**, WWPN 20:**01**:00:60:69:51:06:b4

Port **14**, WWPN 20:**0e**:00:60:69:51:06:b4

Port **15**, WWPN 20:**0f**:00:60:69:51:06:b4

Related concepts

[How WWPNs are used](#) on page 101

Recommended MTU configurations for FCoE jumbo frames

For Fibre Channel over Ethernet (FCoE), jumbo frames for the Ethernet adapter portion of the CNA should be configured at 9000 MTU. Jumbo frames for the FCoE adapter portion of the CNA should be configured at greater than 1500 MTU. Only configure jumbo frames if the initiator, target, and all intervening switches support and are configured for jumbo frames.

Managing systems with FC adapters

There are commands available to manage onboard FC adapters and FC adapter cards. These commands can be used to configure the adapter mode, display adapter information, and change the speed.

Most systems have onboard FC adapters that you can configure as initiators or targets. You can also use FC adapter cards configured as initiators or targets. Initiators connect to back-end disk shelves, tape libraries, and possibly foreign storage arrays (FlexArray). Targets connect to FC switches or other storage controllers.

Related information

[Clustered Data ONTAP 8.3 SAN Configuration Guide](#)

Commands for managing FC adapters

You can use FC commands to manage FC target adapters, FC initiator adapters, and onboard FC adapters for your storage controller.

FC initiator adapter commands work only at the node level. You must use the `run -node node_name` command before you can use the FC initiator adapter commands.

Commands for managing FC target adapters

If you want to...	Use this command...
Display FC adapter information on a node	<code>network fcp adapter show</code>
Modify FC target adapter parameters	<code>network fcp adapter modify</code>
Display FC protocol traffic information	<code>run -node <i>node_name</i> sysstat -f</code>
Display how long the FC protocol has been running	<code>run -node <i>node_name</i> uptime</code>
Display adapter configuration and status	<code>run -node <i>node_name</i> sysconfig -v <i>adapter</i></code>
Verify which expansion cards are installed and whether there are any configuration errors	<code>run -node <i>node_name</i> sysconfig -ac</code>
View a man page for a command	<code>man <i>command_name</i></code>

Commands for managing FC initiator adapters

If you want to...	Use this command...
Display information for all initiators and their adapters in a node	<code>run -node <i>node_name</i> storage show adapter</code>
Display adapter configuration and status	<code>run -node <i>node_name</i> sysconfig -v <i>adapter</i></code>
Verify which expansion cards are installed and whether there are any configuration errors	<code>run -node <i>node_name</i> sysconfig -ac</code>

Commands for managing onboard FC adapters

If you want to...	Use this command...
Display the status of the onboard FC ports	<code>run -node <i>node_name</i> system hardware unified-connect show</code>

Configuring FC adapters for initiator mode

You can configure individual FC ports of onboard adapters and certain FC adapter cards for initiator mode. Initiator mode is used to connect the ports to back-end disk shelves and tapes.

Before you begin

LIFs on the adapter must be removed from any portsets of which they are members.

About this task

Each onboard FC port can be individually configured as an initiator or a target. Ports on certain FC adapters can also be individually configured as either a target port or an initiator port, just like the onboard FC ports. A list of adapters that can be configured for target mode is available in the *Hardware Universe*.

Steps

1. Remove all of the LIFs from the adapter:

```
network interface delete -vserver SVM_name -lif LIF_name,LIF_name
```

2. Take your adapter offline:

```
network fcp adapter modify -node node_name -adapter adapter_port -state down
```

If the adapter does not go offline, you can also remove the cable from the appropriate adapter port on the system.

3. Change the adapter from target to initiator:

```
system hardware unified-connect modify -t initiator adapter_port
```

4. Reboot the node hosting the adapter you changed.

5. Verify that the FC ports are configured in the correct state for your configuration:

```
system hardware unified-connect show
```

6. Bring the adapter back online:

```
node run -node node_name storage enable adapter -e adapter_port
```

Related tasks

[Deleting a LIF in a SAN environment](#) on page 47

Related information

[NetApp Hardware Universe](#)

Configuring FC adapters for target mode

You can configure individual FC ports of onboard adapters and certain FC adapter cards for target mode. Target mode is used to connect the ports to FC initiators.

About this task

Each onboard FC port can be individually configured as an initiator or a target. Ports on certain FC adapters can also be individually configured as either a target port or an initiator port, just like the onboard FC ports. A list of adapters that can be configured for target mode is available in the *Hardware Universe*.

Steps

1. Take the adapter offline:

```
node run -node node_name storage disable adapter -d adapter_port
```

If the adapter does not go offline, you can also remove the cable from the appropriate adapter port on the system.

2. Change the adapter from initiator to target:

```
system hardware unified-connect modify -t target adapter_port
```

3. Reboot the node hosting the adapter you changed.

4. Verify that the target port has the correct configuration:

```
network fcp adapter show -node node_name
```

5. Bring your adapter online:

```
network fcp adapter modify -node node_name -adapter adapter_port -state up
```

Related information

[NetApp Hardware Universe](#)

Displaying information about an FC target adapter

You can use the `network fcp adapter show` command to display system configuration and adapter information for any FC adapter in the system.

Step

1. Display information about the FC adapter by using the `network fcp adapter show` command.

Example

The output displays system configuration information and adapter information for each slot that is used.

```
network fcp adapter show -instance -node node1 -adapter 0a
```

Changing the FC adapter speed

You should set your adapter target port speed to match the speed of the device to which it connects, instead of using autonegotiation. A port that is set to autonegotiation can take longer to reconnect after a takeover/giveback or other interruption.

Before you begin

All LIFs that use this adapter as their home port must be offline.

About this task

Because this task encompasses all Storage Virtual Machines (SVMs) and all LIFs in a cluster, you must use the `-home-port` and `-home-lif` parameters to limit the scope of this operation. If you do not use these parameters, the operation applies to all LIFs in the cluster, which might not be desirable.

Steps

1. Take all of the LIFs on this adapter offline:


```
network interface modify -vserver * -lif * { -home-node node1 -home-port e0c } -status-admin down
```
2. Take the adapter offline:


```
network fcp adapter modify -node node1 -adapter 0c -state down
```

If the adapter does not go offline, you can also remove the cable from the appropriate adapter port on the system.
3. Determine the maximum speed for the port adapter:


```
fcp adapter show -instance
```

You cannot modify the adapter speed beyond the maximum speed.
4. Change the adapter speed:


```
network fcp adapter modify -node node1 -adapter 0a -speed 8
```
5. Bring the adapter online:


```
network fcp adapter modify -node node1 -adapter 0c -state up
```

- Bring all of the LIFs on the adapter online:

```
network interface modify -vserver * -lif * { -home-node node1 -home-port
e0c } -status-admin up
```

Supported port configurations for X1143A-R6 adapters

FC target mode is the default configuration for X1143A-R6 adapter ports. However, ports on this adapter can be configured as either 10-Gb Ethernet and FCoE ports or as 16-Gb FC ports.

When configured for Ethernet and FCoE, X1143A-R6 adapters support concurrent NIC and FCoE target traffic on the same 10-GBE port. When configured for FC, each 2-port pair that shares the same ASIC can be individually configured for FC target or FC initiator mode. This means that a single X1143A-R6 adapter can support FC target mode on one 2-port pair and FC initiator mode on another 2-port pair.

Related information

[NetApp Hardware Universe](#)

[Clustered Data ONTAP 8.3 SAN Configuration Guide](#)

Configuring your ports

To configure the unified target adapter (X1143A-R6), you must configure your two adjacent ports on the same chip in the same personality mode.

Steps

- Configure your ports as needed for Fibre Channel (FC) or Converged Network Adapter (CNA) using the `system hardware unified-connect modify` command.
- Attach the appropriate cables for FC or 10 Gb Ethernet.

Changing the CNA/UTA2 target adapter optical modules

You should change the optical modules on the unified target adapter (CNA/UTA2) to support the personality mode you have selected for the adapter.

Steps

- Remove the current optical modules from the X1143A-R6 adapter.
- Insert the correct modules for your preferred personality mode (FC or CNA) optics.
- Verify that you have the correct SFP+ installed:

```
network fcp adapter show -instance -node -adapter
```

Supported SFP+ modules and Cisco-branded Copper (Twinax) cables are listed in the *Hardware Universe*.

Related information

[NetApp Hardware Universe](#)

Viewing adapter settings

To view the settings for your unified target adapter (X1143A-R6), you must run the `system hardware unified-connect show` command to display all modules on your controller.

Steps

1. Boot your controller without the cables attached.
2. Run the `system hardware unified-connect show` command to see the port configuration and modules.
3. View the port information before configuring the CNA and ports.

How to prevent loss of connectivity when using the X1133A-R6 adapter

The X1133A-R6 HBA is a 4-port, 16-Gb, target-only FC adapter consisting of two, 2-port pairs. Each 2-port pair is supported by a single ASIC. If an error occurs with the ASIC supporting a pair, both ports in the pair will go offline.

To prevent loss of connectivity in the event of port failure, it is recommended that you configure your system with redundant paths to separate X1133A-R6 HBAs, or with redundant paths to ports supported by different ASICs on the HBA.

Storage virtualization with VMware and Microsoft copy offload

VMware and Microsoft support copy offload operations to increase performance and network throughput. You must configure your system to meet the requirements of the VMware and Windows operating system environments to use their respective copy offload functions.

When using VMware and Microsoft copy offload in virtualized environments, your LUNs must be aligned. Unaligned LUNs can degrade performance.

Related concepts

[I/O misalignments might occur on properly aligned LUNs](#) on page 35

[How to achieve I/O alignment using LUN OS types](#) on page 36

[Special I/O alignment considerations for Linux](#) on page 37

[Special I/O alignment considerations for Solaris LUNs](#) on page 37

[ESX boot LUNs report as misaligned](#) on page 37

Advantages of using a virtualized SAN environment

Creating a virtualized environment by using Storage Virtual Machines (SVMs) and LIFs enables you to expand your SAN environment to all of the nodes in your cluster.

- Distributed management
You can log in to any node in the SVM to administer all of the nodes in a cluster.
- Increased data access
With MPIO and ALUA, you have access to your data through any active iSCSI or FC LIFs for the SVM.
- Controlled LUN access
If you use SLM and portsets, you can limit which LIFs an initiator can use to access LUNs.

Related concepts

[Why Data ONTAP uses ALUA](#) on page 44

[MPIO and ALUA](#) on page 44

[Ways to limit LUN access with port sets and igroups](#) on page 28

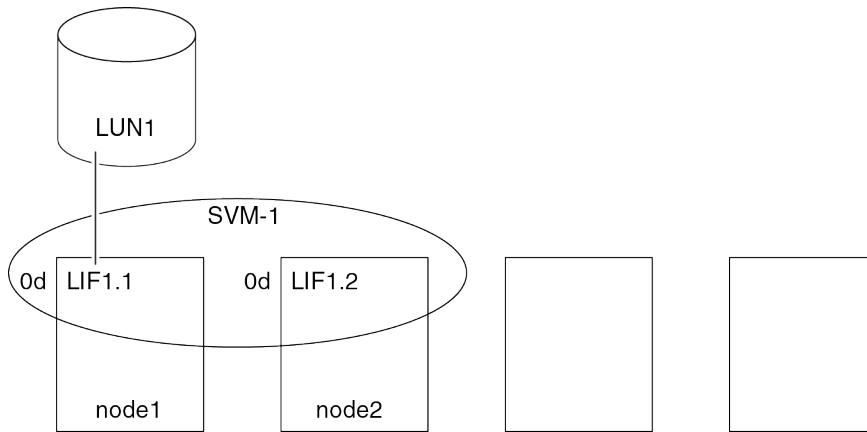
How LUN access works in a virtualized environment

In a virtualized environment, LIFs enable Storage Virtual Machines (SVMs) to access LUNs through optimized and unoptimized paths. You can have a single SVM per node or you can have multiple SVMs on a single node.

A LIF is a logical interface that connects the SVM to a physical port. Each node that you add a LIF to becomes a member of the SVM. Although multiple SVMs can have multiple LIFs on the same ports, a LIF belongs to one SVM. LUNs can then be accessed through the node LIFs


Example of LUN access with a single SVM on a node

In the following example, LIF1.1 connects SVM SVM-1 to node1:0d. LIF1.1 and LIF 1.2 belong only to SVM-1. If a new LUN is created, it can connect SVM-1 to node2:0d.



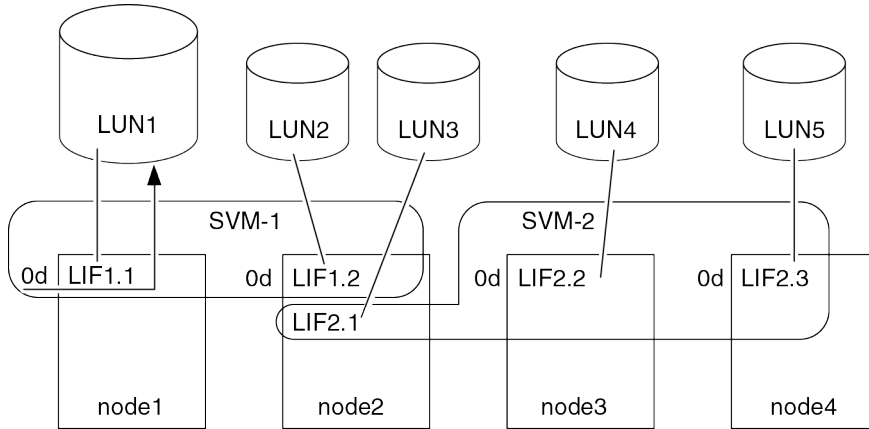
Example of LUN access with multiple SVMs on a node

A physical port can have multiple LIFs serving different SVMs. Because LIFs are associated with a particular SVM, the cluster node can send the incoming data traffic to the correct SVM. In the following example, node2 has two LIFs: LIF1.2 and LIF2.1 for two SVMs on port 0d. SVM SVM-1 has LIF1.2 and SVM SVM-2 has LIF2.1, both of which are on node2:0d. SVM SVM-1 connects LIF1.2 to LUN2 and SVM SVM-2 connects LIF2.1 to LUN3.




Example of an optimized path to a LUN

In an optimized path, the data traffic does not travel over the cluster network; it travels the most direct route to the LUN. The optimized path to LUN1 is through LIF1.1 in node1:0d, as shown in the following example:



Example of an unoptimized path to a LUN

In an unoptimized path, the data travels over the cluster network. The following example illustrates two unoptimized paths. One unoptimized path to LUN1 is through the cluster network for node2. The data traffic destined for LUN1 enters through node2:0d and travels through the cluster network to node1 to reach LUN1. Another unoptimized path to LUN1 is through node3:0d. The data for LUN1 enters through node3:0d and travels through the cluster network to node1 to reach LUN1.



Considerations for LIFs in cluster SAN environments

You need to be aware of certain LIF considerations in a SAN environment.

- Initiators must use MPIO and ALUA for failover capability for clusters in a SAN iSCSI or FC environment because SAN does not support automatic failover for LIFs.
- There must be at least one SAN LIF of the appropriate protocol configured on each node that hosts a mapped LUN and its HA partner.
- Some options are not applicable for iSCSI or FC.
For example, you cannot use IP addresses with FC.

Improving VMware VAAI performance for ESX hosts

Data ONTAP 8.2 and later supports certain VMware vStorage APIs for Array Integration (VAAI) features when the ESX host is running ESX 4.1 or later. These features help offload operations from the ESX host to the storage system and increase the network throughput. The ESX host enables the features automatically in the correct environment.

The VAAI feature supports the following SCSI commands:

- EXTENDED_COPY

This feature enables the host to initiate the transfer of data between the LUNs or within a LUN without involving the host in the data transfer. This results in saving ESX CPU cycles and increasing the network throughput. The extended copy feature, also known as "copy offload," is used in scenarios such as cloning a virtual machine. When invoked by the ESX host, the copy offload feature copies the data within the NetApp storage system rather than going through the host network. Copy offload transfers data in the following ways:

- Within a LUN
- Between LUNs within a volume
- Between LUNs on different volumes within a Storage Virtual Machine (SVM)
- Between LUNs on different SVMs within a cluster

If this feature cannot be invoked, the ESX host automatically uses the standard READ and WRITE commands for the copy operation.

- WRITE_SAME

This feature offloads the work of writing a repeated pattern, such as all zeros, to a storage array. The ESX host uses this feature in operations such as zero-filling a file.

- COMPARE_AND_WRITE

This feature bypasses certain file access concurrency limits, which speeds up operations such as booting up virtual machines.

Requirements for using the VAAI environment

The VAAI features are part of the ESX operating system and are automatically invoked by the ESX host when you have set up the correct environment.

The environment requirements are as follows:

- The ESX host must be running ESX 4.1 or later.
- The NetApp storage system that is hosting the VMware datastore must be running Data ONTAP 8.2 or later.
- (Copy offload only) The source and the destination of the VMware copy operation must be hosted on the same storage system within the same cluster.

Note: The copy offload feature currently does not support copying data between VMware datastores that are hosted on different storage systems.

How to determine if VAAI features are supported by ESX

To confirm whether the ESX operating system supports the VAAI features, you can check the vSphere Client or use any other means of accessing the host. Data ONTAP 8.2 and later supports the SCSI commands by default.

You can check your ESX host advanced settings to determine whether VAAI features are enabled. The table indicates which SCSI commands correspond to ESX control names.

SCSI command	ESX control name (VAAI feature)
EXTENDED_COPY	HardwareAcceleratedMove
WRITE_SAME	HardwareAcceleratedInit
COMPARE_AND_WRITE	HardwareAcceleratedLocking

Microsoft Offloaded Data Transfer (ODX)

Microsoft Offloaded Data Transfer (ODX), also known as *copy offload*, enables direct data transfers within a storage device or between compatible storage devices without transferring the data through the host computer.

Data ONTAP supports ODX for both the CIFS and SAN protocols.

In non-ODX file transfers, the data is read from the source and is transferred across the network to the host. The host transfers the data back over the network to the destination. In ODX file transfer, the data is copied directly from the source to the destination without passing through the host.

Because ODX offloaded copies are performed directly between the source and destination, significant performance benefits are realized, including faster copy time, reduced utilization of CPU and memory on the client, and reduced network I/O bandwidth utilization.

For SAN environments, ODX is only available when it is supported by both the host and the storage system. Client computers that support ODX and have ODX enabled automatically and transparently use offloaded file transfer when moving or copying files. ODX is used regardless of whether you drag-and-drop files through Windows Explorer or use command-line file copy commands, or whether a client application initiates file copy requests.

Requirements for using ODX

If you plan to use ODX for copy offloads, you need to be familiar with volume support considerations, system requirements, and software capability requirements.

ODX is only supported on Storage Virtual Machines (SVMs) with FlexVol volumes for intra-cluster copies. ODX cannot be used to copy files or folders to a volume in another cluster. ODX also cannot be used to copy data to or from volumes in SVMs with Infinite Volume.

To use ODX, your system must have the following:

- Clustered Data ONTAP 8.2 or later
ODX is automatically enabled in supported versions of Data ONTAP.
- Minimum source volume of 2 GB
For optimal performance, the source volume should be greater than 260 GB.
- Deduplication
ODX uses deduplication as part of the copy process. If you do not want deduplication on your SVM, you should disable ODX on that SVM.
- ODX support on the Windows client

ODX is supported in Windows Server 2012 or later and in Windows 8 or later. For the latest information about supported Windows clients, see the Interoperability Matrix.

- Copy application support for ODX
The application that performs the data transfer must support ODX. Application operations that support ODX include the following:
 - Hyper-V management operations, such as creating and converting virtual hard disks (VHDs), managing Snapshot copies, and copying files between virtual machines
 - Windows Explorer operations
 - Windows PowerShell copy commands
 - Windows command prompt copy commands

For more information about supported ODX applications on Windows servers and clients, consult the Microsoft TechNet Library.

ODX does not work with the following volume types:

- Source volumes with capacities of less than 2 GB
- Compressed volumes
- Sparse volumes
- Read-only volumes
- Semi-thick provisioned volumes

Use cases for ODX

You should be aware of the use cases for using ODX on SVMs with FlexVol volumes so that you can determine under what circumstances this feature provides you with performance benefits.

Windows servers and clients that support ODX use copy offload as the default way of copying data across remote servers. If the Windows server or client does not support ODX or the ODX copy offload fails at any point, the copy or move operation falls back to traditional reads and writes for the copy or move operation.

The following use cases support using ODX copies and moves:

- Intra-volume
The source and destination files or LUNs are within the same volume.
- Inter-volume, same node, same SVM
The source and destination files or LUNs are on different volumes that are located on the same node. The data is owned by the same SVM.
- Inter-volume, different nodes, same SVM
The source and destination files or LUNs are on different volumes that are located on different nodes. The data is owned by the same SVM.
- Inter-SVM, same node
The source and destination file or LUNs are on different volumes that are located on the same node. The data is owned by different SVMs.
- Inter-SVM, different nodes
The source and destination file or LUNs are on different volumes that are located on different nodes. The data is owned by different SVMs.

There are some additional special use cases:

- With the Data ONTAP ODX implementation, you can use ODX to copy files between SMB shares and FC or iSCSI attached virtual drives.
You can use Windows Explorer, the Windows CLI or PowerShell, Hyper-V, or other applications that support ODX to copy or move files seamlessly using ODX copy offload between SMB shares and connected LUNs, provided the SMB shares and LUNs are on the same cluster.
- Hyper-V provides some additional use cases for ODX copy offload:
 - You can use ODX copy offload pass-through with Hyper-V to copy data within or across virtual hard disk (VHD) files or to copy data between mapped SMB shares and connected iSCSI LUNs within the same cluster.
This allows copies from guest operating systems to pass through to the underlying storage.
 - When creating fixed-sized VHDs, ODX is used for initializing the disk with zeros, using a well-known zeroed token.
 - ODX copy offload is used for virtual machine storage migration if the source and destination storage is on the same cluster.

Note: To take advantage of the use cases for ODX copy offload pass-through with Hyper-V, the guest operating system must support ODX and the guest operating system's disks must be SCSI disks backed by storage (either SMB or SAN) that supports ODX. IDE disks on the guest operating system do not support ODX pass-through.

Special system file requirements

When using the ODX feature, there are ODX system files that exist in every volume of the system. You must not remove or modify these files unless technical support tells you to do so.

These files enable point-in-time representation of data used during the ODX transfer. The following system files are in the root level of each volume that contains LUNs or files to which data was offloaded:

- .copy-offload (directory)
- .tokens (file under .copy_offload directory)

Basic iSCSI and FC concepts

In iSCSI networks and FC fabrics, storage systems are targets that have storage target devices, which are referred to as LUN (logical units). Using the Data ONTAP operating system, you configure the storage by creating LUNs. The LUNs are accessed by hosts, which are initiators in the storage network.

Protocols that hosts can use to connect to SAN storage systems

Hosts can connect to SAN storage systems using Internet Small Computer Systems Interface (iSCSI) or Fibre Channel (FC) protocol networks.

To connect to iSCSI networks, hosts can use standard Ethernet network adapters (NICs), TCP offload engine (TOE) cards with software initiators, converged network adapters (CNAs), or dedicated iSCSI host bus adapters (HBAs).

To connect to FC networks, hosts require FC HBAs or CNAs.

What Host Utilities are

Host Utilities includes support software and documentation for connecting a supported host to an iSCSI or FC network.

The support software includes programs that display information about storage and programs to collect information that technical support personnel need to diagnose problems. It also includes software to help tune and optimize the host settings for use in a NetApp storage infrastructure.

Separate host utilities are offered for each supported host operating system. In some cases, different versions of the Host Utilities are available for different versions of the host operating system.

The documentation included with the host utilities describes how to install and use the host utilities software. It includes instructions for using the commands and features specific to your host operating system.

You must use the Host Utilities documentation along with this guide to set up and manage your iSCSI or FC network.

Related information

[NetApp Interoperability Matrix Tool](#)

[NetApp Documentation: Host Utilities \(current releases\)](#)

Simplified host management with SnapDrive

You can use SnapDrive software to simplify some of the management and data protection tasks associated with iSCSI and FC storage. SnapDrive is an optional management package for Windows and UNIX hosts.

You can use SnapDrive for Windows to easily create virtual disks from pools of storage that can be distributed among several storage systems.

You can use SnapDrive for UNIX to automate storage provisioning tasks and simplify the process of creating Snapshot copies and clones from Snapshot copies consistent with host data.

See NetApp product documentation for more information on SnapDrive.

Related information

NetApp Support: File Upload Utility

How Data ONTAP implements an iSCSI network

You should be aware of important concepts that are required to understand how Data ONTAP implements an iSCSI network.

Related concepts

iSCSI service management on page 73

What iSCSI is

The iSCSI protocol is a licensed service on the storage system that enables you to transfer block data to hosts using the SCSI protocol over TCP/IP. The iSCSI protocol standard is defined by RFC 3270.

In an iSCSI network, storage systems are targets that have storage target devices, which are referred to as LUNs (logical units). A host with an iSCSI host bus adapter (HBA), or running iSCSI initiator software, uses the iSCSI protocol to access LUNs on a storage system. The iSCSI protocol is implemented over the storage system's standard Ethernet interfaces using a software driver.

The connection between the initiator and target uses a standard TCP/IP network. No special network configuration is needed to support iSCSI traffic. The network can be a dedicated TCP/IP network, or it can be your regular public network. The storage system listens for iSCSI connections on TCP port 3260.

Related information

RFC 3270: www.ietf.org/rfc/rfc3270.txt

What iSCSI nodes are

In an iSCSI network, there are two types of nodes: targets and initiators. Targets are storage systems, and initiators are hosts. Switches, routers, and ports are TCP/IP devices only, and are not iSCSI nodes.

How iSCSI target nodes connect to the network

You can implement iSCSI on the storage system using several different software solutions.

Target nodes can connect to the network in the following ways:

- Over Ethernet interfaces using software that is integrated into Data ONTAP.
 - Over multiple system interfaces, with an interface used for iSCSI that can also transmit traffic for other protocols, such as CIFS and NFS.
- Using a unified target adapter (UTA) or a converged network adapter (CNA).

How iSCSI nodes are identified

Every iSCSI node must have a node name.

The two formats, or type designators, for iSCSI node names are *iqn* and *eui*. The storage system always uses the iqn-type designator. The initiator can use either the iqn-type or eui-type designator.

iqn-type designator

The iqn-type designator is a logical name that is not linked to an IP address.

It is based on the following components:

- The type designator, such as `iqn`
- A node name, which can contain alphabetic characters (a to z), numbers (0 to 9), and three special characters:
 - Period (“.”)
 - Hyphen (“-”)
 - Colon (“:”)
- The date when the naming authority acquired the domain name, followed by a period
- The name of the naming authority, optionally followed by a colon (:)
- A unique device name

Note: Some initiators might provide variations on the preceding format. Also, even though some hosts do support underscores in the host name, they are not supported on NetApp systems. For detailed information about the default initiator-supplied node name, see the documentation provided with your iSCSI Host Utilities.

An example format is as follows:

```
iqn.yyyymm.backward_naming_authority:unique_device_name
```

`yyyy-mm` is the month and year in which the naming authority acquired the domain name.

`backward_naming_authority` is the reverse domain name of the entity responsible for naming this device. An example reverse domain name is `com.microsoft`.

`unique-device-name` is a free-format unique name for this device assigned by the naming authority.

The following example shows the iSCSI node name for an initiator that is an application server:

```
iqn.1991-05.com.microsoft:example
```

eui-type designator

The eui-type designator is based on the type designator, `eui`, followed by a period, followed by sixteen hexadecimal digits.

A format example is as follows:

```
eui.0123456789abcdef
```

Storage system node name

Each storage system has a default node name based on a reverse domain name and a unique encoding number.

The node name is displayed in the following format:

```
iqn.1992-08.com.netapp:sn.unique-encoding-number
```

The following example shows the default node name for a storage system with a unique encoding number:

```
iqn.1992-08.com.netapp:sn.812921059e6c11e097b3123478563412:vs.6
```

How the storage system checks initiator node names

The storage system checks the format of the initiator node name at session login time. If the initiator node name does not comply with storage system node name requirements, the storage system rejects the session.

Default port for iSCSI

The iSCSI protocol is configured in Data ONTAP to use TCP port number 3260.

Data ONTAP does not support changing the port number for iSCSI. Port number 3260 is registered as part of the iSCSI specification and cannot be used by any other application or service.

How iSCSI communication sessions work

During an iSCSI session, the initiator and the target communicate over their standard Ethernet interfaces, unless the host has an iSCSI HBA or a CNA.

The Storage Virtual Machine (SVM) appears as a single iSCSI target node with one iSCSI node name.

On the storage system, the interface can be an Ethernet port, interface group (ifgrp), UTA, or a virtual LAN (VLAN) interface.

Each interface on the target belongs to its own portal group by default. This enables an initiator port to conduct simultaneous iSCSI sessions on the target, with one session for each portal group. The maximum number of sessions for a storage system depends on the number of nodes in a cluster and the memory capacity of that storage system. To determine whether your host's initiator software or HBA can have multiple sessions with one storage system, see your host OS or initiator documentation.

Each session has an Initiator Session ID (ISID), a number that is determined by the initiator.

How high availability is maintained in an iSCSI SVM

An iSCSI Storage Virtual Machine (SVM) provides high availability through MPIO and ALUA. When a node fails, availability is maintained by rerouting traffic to the other nodes within the SVM. When the path to a primary node fails, availability is maintained by rerouting traffic through indirect node paths.

How Data ONTAP implements an FC SAN

You should be aware of the important concepts that are required to understand how Data ONTAP implements an FC SAN.

What FC is

FC is a licensed service on the storage system that enables you to export LUNs and transfer block data to hosts using the SCSI protocol over a Fibre Channel fabric.

What FC nodes are

In an FC network, nodes include targets, initiators, and switches.

Targets are storage systems, and initiators are hosts. Nodes register with the Fabric Name Server when they are connected to an FC switch. Each Storage Virtual Machine (SVM) that has a FCP service is a different FC target node.

How FC target nodes connect to the network

Storage systems and hosts have adapters, so they can be directly connected to each other or to FC switches with cables. For switch or storage system management, they might be connected to each other or to TCP/IP switches with Ethernet cables.

When a node is connected to the FC SAN, it registers each of its ports with the switch's Fabric Name Server service, using a unique identifier.

How FC nodes are identified

Each FC node is identified by a worldwide node name (WWNN).

How WWPNs are used

WWPNs identify each port in an FC node.

- Creating an initiator group

The WWPNs of the host's HBAs are used to create an initiator group (igroup). An igroup is used to control host access to specific LUNs. You can create an igroup by specifying a collection of WWPNs of initiators in an FC network. When you map a LUN on a storage system to an igroup, you can grant all the initiators in that group access to that LUN. If a host's WWPN is not in an igroup that is mapped to a LUN, that host does not have access to the LUN. This means that the LUNs do not appear as disks on that host.

You can also create port sets to make a LUN visible only on specific target ports. A port set consists of a group of FC target ports. You can bind an igroup to a port set. Any host in the igroup can access the LUNs only by connecting to the target ports in the port set.

- Uniquely identifying FC LIFs

WWPNs uniquely identify each FC logical interface. The host operating system uses the combination of the WWNN and WWPN to identify Storage Virtual Machines (SVMs) and FC LIFs. Some operating systems require persistent binding to ensure that the LUN appears at the same target ID on the host.

Related concepts

[Ways to limit LUN access with port sets and igroups](#) on page 28

Copyright information

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

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

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

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

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

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

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

Trademark information

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

- access lists
 - commands [69](#)
 - iSNS [75](#)
- active initiators
 - available commands [70](#)
 - displaying [70](#)
 - showing [70](#)
- active volumes
 - cloning LUNs from [56](#)
- adapter ports
 - configurations supported for X1143A-R6 [88](#)
- adapters
 - active initiators [81](#)
 - changing WWPNs for FC logical interfaces [82](#)
 - commands for managing FC [84](#)
 - displaying
 - FC logical interfaces [82](#)
 - displaying information about FC target [87](#)
 - FC, configuring for initiator mode [85](#)
 - how to prevent loss of connectivity when using X1133A-R6 [89](#)
 - introduction to managing systems with FC [84](#)
- adapters, FC
 - changing speed [87](#)
- adapters, unified target
 - changing optical modules [88](#)
- adding initiators to igroups
 - command for [42](#)
- adding nodes to a cluster
 - considerations for [48](#)
- aggregates
 - overcommitment, defined [8](#)
- ALUA
 - considerations for LIFs in cluster SAN environments [92](#)
 - how Data ONTAP uses to identify optimized and unoptimized paths [44](#)
- authentication
 - displaying security information [69](#)
 - iSCSI [73](#)
- autogrow
 - SAN volume configuration option [14](#)
- automatically deleting
 - disabling autodelete for FlexClone files and LUNs [57](#)
 - FlexClone files and FlexClone LUNs [54](#)

B

- backups
 - about configuring and using SnapVault, in SAN environments [58](#)
 - accessing read-only LUN copies from SnapVault [59](#)
 - backing up LUNs through host system [65](#)
 - restoring all LUNs in a volume from SnapVault [62](#)
 - restoring single LUNs from SnapVault [60](#)
- binding igroups

- command for [42](#)
- block access
 - enabling for a specific host [23](#)
 - enabling for SVMs with FC [21](#)
 - enabling for SVMs with iSCSI [21](#)

C

- CHAP
 - defined [74](#)
 - guidelines [74](#)
 - iSCSI authentication [73](#)
- Cisco switches
 - FC and FCoE zoning requirement [47](#)
- cluster SAN environments
 - considerations for LIFs in [92](#)
- clusters
 - considerations for adding nodes to [48](#)
- CNA/UTA2 target adapters
 - changing optical modules [88](#)
- commands
 - FC
 - active initiators [81](#)
 - logical interfaces [82](#)
 - service [81](#)
 - WWPN alias [82](#)
 - for managing igroups [42](#)
 - for managing LUNs [33](#)
 - initiator [70](#)
 - iSCSI
 - connections [70](#)
 - initiator security [69](#)
 - interface access [69](#)
 - security [71](#)
 - service [68](#)
 - sessions [70](#)
 - iSNS service [78](#)
 - port sets [34](#)
 - storage system interfaces [69](#)
- comments
 - how to send feedback about documentation [104](#)
- communication
 - iSNS [75](#)
- COMPARE WRITE feature
 - VAAI feature [93](#)
- compression
 - impact with zero fractional reserve [15](#)
- configuration limits
 - volume and SVM requirements [14](#)
- configuration options
 - SAN volume [14](#)
- configuration settings
 - for non-space-reserved LUNs with thin volume provisioning [11](#)
 - for space-reserved LUNs with semi-thick volume provisioning [11](#)
 - for space-reserved LUNs with thick-provisioned volumes [10](#)
- configured space

- examining LUN [33](#)
- configuring
 - FC adapters and onboard ports for initiator mode [85](#)
 - FC adapters for target mode [86](#)
 - switches for FCoE [19](#)
 - unified target adapter ports [88](#)
- connectivity loss
 - how to prevent when using X1133A-R6 adapters [89](#)
- considerations
 - for adding nodes to a cluster [48](#)
 - for LIFs in cluster SAN environments [92](#)
 - LUNs and files [25](#)
- copies, Snapshot
 - effect of moving or copying a LUN on [50](#)
- copy offload
 - use cases for [95](#)
 - using VMware and Microsoft for storage virtualization [90](#)
 - See also* ODX
- copy offloads
 - requirements for using ODX [94](#)
- copying
 - considerations for LUNs [30](#)
- creating
 - igroups, command for [42](#)
 - iSCSI initiator security method [69](#)
 - LUNs, command for [33](#)
 - port sets [34](#)

D

- data growth
 - calculating rate for LUNs [16](#)
- data protection
 - FlexClone LUNs, how you use for [52](#)
 - methods of, in SAN environments [49](#)
- decreasing
 - LUN size [31](#)
- deduplication
 - impact with zero fractional reserve [15](#)
- default security authentication
 - defining [69](#)
- defining
 - default security authentication [69](#)
 - security password [69](#)
- deleting
 - disabling autodelete for FlexClone files and LUNs [57](#)
 - FlexClone files and LUNs automatically [54](#)
 - igroups, command for [42](#)
 - LUNs [32](#)
 - LUNs, command for [33](#)
 - port sets [34](#)
- destroying
 - LUNs, command for [33](#)
 - port sets [34](#)
- disaster recovery, SVM
 - available types in SAN [67](#)
- displaying
 - authentication security information [69](#)
 - FC adapter information, commands for [84](#)
 - igroup [34](#)
 - igroup information, command for [42](#)

- LUN mapping information, command for [33](#)
- port sets [34](#)
- documentation
 - how to receive automatic notification of changes to [104](#)
 - how to send feedback about [104](#)

E

- environments, SAN cluster
 - considerations for LIFs in [92](#)
- error messages
 - resolving iSCSI storage system [80](#)
- error recovery levels
 - enabling iSCSI levels 1 and 2 [72](#)
- ESX boot LUNs
 - reported as misaligned [37](#)
- Ethernet [73](#), [98](#)
- eui type designator [99](#)
- examples
 - of how igroups give LUN access [43](#)
- extended copy feature
 - environment [93](#)
 - invoked automatically [93](#)
 - VAAI feature [93](#)
 - when the standard copy operation is used [93](#)

F

- fabrics
 - changing FC adapter speed to match [87](#)
- FC
 - active initiators
 - commands [81](#)
 - displaying [81](#)
 - adding [81](#), [82](#)
 - adding target alias [81](#)
 - assigning [82](#)
 - deleting [82](#)
 - displaying [82](#)
 - displaying target alias [81](#)
 - enabling block access for SVMs with [21](#)
 - license [20](#)
 - LIF [81](#)
 - logical interfaces [82](#)
 - WWPN alias [82](#)
- FC adapter speeds
 - changing [87](#)
- FC adapters
 - commands for managing [84](#)
 - configuring for initiator mode [85](#)
 - configuring for target mode [86](#)
 - introduction to managing systems with [84](#)
- FC commands [81](#), [82](#)
- FC license
 - adding [81](#)
 - disabling [81](#)
 - enabling [81](#)
 - verifying [81](#)
- FC LIF
 - zoning restrictions for Cisco switches [47](#)
- FC logical interfaces

- available commands [82](#)
 - changing WWPNs for [82](#)
 - FC onboard ports
 - configuring for initiator mode [85](#)
 - FC service
 - commands [81](#)
 - deleting [81](#), [83](#)
 - modifying [81](#)
 - starting [81](#)
 - stopping [81](#)
 - FC target
 - creating [81](#)
 - FC target adapters
 - displaying information about [87](#)
 - target adapters
 - displaying [87](#)
 - FC target alias
 - adding [81](#)
 - FC WWPN alias
 - available commands [82](#)
 - creating [82](#)
 - deleting [82](#)
 - destroying [82](#)
 - displaying [82](#)
 - modifying [82](#)
 - FCoE
 - configuring switches for [19](#)
 - recommended MTU configurations for jumbo frames [84](#)
 - FCoE LIF
 - zoning restrictions for Cisco switches [47](#)
 - FCP
 - defined [100](#)
 - node connection [101](#)
 - node identification [101](#)
 - nodes defined [100](#)
 - switch nodes [84](#)
 - FCP service
 - destroying [81](#)
 - feedback
 - how to send comments about documentation [104](#)
 - FlexClone files and FlexClone LUNs
 - disabling automatic deletion of [57](#)
 - enabling automatic deletion of [54](#)
 - how a volume reclaims free space from [53](#)
 - FlexClone LUNs
 - data protection, how you use for [52](#)
 - impact with zero fractional reserve [15](#)
 - reasons for using [53](#)
 - using a Snapshot copy to create [57](#)
 - FlexVol volumes
 - configuration settings for non-space-reserved LUNs with thin volume provisioning [11](#)
 - configuration settings for space-reserved LUNs with semi-thick volume provisioning [11](#)
 - configuration settings for thick-provisioned with space-reserved LUNs [10](#)
 - configuring automatic deletion of FlexClone files and LUNs [54](#)
 - considerations for setting fractional reserve [15](#)
 - determining the correct configuration for your environment [12](#)
 - how they reclaim free space from FlexClone files and FlexClone LUNs [53](#)
 - provisioning options [8](#)
 - recommended configuration combinations with LUNs [9](#)
 - required for SAN [14](#)
 - thin-provisioned, defined [7](#)
 - what it means to overcommit volumes that supply storage to [8](#)
 - fractional reserve
 - considerations for setting for FlexVol volumes [15](#)
 - free space
 - how FlexVol volumes reclaim from FlexClone files and LUNs [53](#)
- ## G
- growth, data
 - calculating rate for LUNs [16](#)
 - guidelines
 - CHAP authentication [74](#)
 - for mapping LUNs to igroups [25](#)
- ## H
- HBA [73](#), [98](#)
 - head swaps
 - changing WWPNs for FC logical interfaces [82](#)
 - high availability [44](#)
 - host
 - iSCSI implementation [73](#)
 - iSNS [75](#)
 - host backup systems
 - how you can connect to the primary storage system [64](#)
 - host configurations
 - SAN considerations when transitioning from 7-Mode to Data ONTAP [29](#)
 - host management
 - for Windows and UNIX using SnapDrive [97](#)
 - host utilities
 - guidelines for assigning LUN IDs [25](#)
 - Host Utilities
 - defined [97](#)
 - host-side space management
 - automatically enabled with SCSI thin provisioned LUNs [40](#)
 - defined [39](#)
 - hosts
 - backing up a LUN through [65](#)
 - protocols they can use to connect to SAN storage systems [97](#)
 - support for SCSI thin provisioning [41](#)
 - troubleshooting when iSCSI LUNs not visible on [79](#)
 - hosts, SAN
 - enabling block access for [23](#)
- ## I
- I/O
 - alignment considerations for Linux [37](#)
 - alignment considerations for Solaris LUNs [37](#)

- I/O alignment
 - how to ensure proper, using OS types [36](#)
- I/O misalignment warnings
 - might be reported on properly aligned LUNs [35](#)
- I/O performance
 - controlling and monitoring to LUNs using Storage QoS [37](#)
- Identity discard SVM disaster recovery
 - explained [67](#)
- identity preserve SVM disaster recovery
 - explained [67](#)
- IDs
 - guidelines for assigning LUN [25](#)
- igroup
 - WWPN [101](#)
- igroups
 - binding [34](#)
 - binding to port sets [24](#)
 - commands for managing [42](#)
 - defined [42](#)
 - example of how LUN access is given with [43](#)
 - guidelines for mapping LUNs to [25](#)
 - how to specify initiator WWPNs and node names for [43](#)
 - mapping to LUNs [22](#)
 - unbinding [34](#)
 - ways to limit LUN access in a virtualized environment with [28](#)
- Infinite Volumes
 - not supported for SAN [14](#)
- information
 - how to send feedback about improving documentation [104](#)
- initiator
 - node name
 - login [100](#)
- initiator mode
 - configuring FC adapters and onboard ports for [85](#)
- initiator security list
 - removing [69](#)
- initiators
 - how to specify WWPNs and node names for an igroup [43](#)
- interface access list
 - for iSCSI initiator interface limits [75](#)
- interfaces
 - changing WWPNs for FC logical [82](#)
- iqn type designator [98](#)
- iSCSI
 - how high availability is maintained in SVMs [100](#)
- iSCSI
 - adding
 - LIFs [68](#)
 - target alias [68](#)
 - available commands [68](#), [69](#)
 - commands [69](#)
 - configuring network for best performance [71](#)
 - connections
 - available commands [70](#)
 - deleting [70](#)
 - shutdown [70](#)
 - creating
 - targets [68](#)
 - default TCP port [100](#)
 - disabling [69](#)
 - displaying
 - connections [70](#)
 - node names [68](#)
 - target alias [68](#)
 - enabling [69](#)
 - enabling block access for SVMs with [21](#)
 - enabling error recovery levels 1 and 2 [72](#)
 - explained [98](#)
 - host implementation [73](#)
 - how communication sessions work [100](#)
 - how nodes are identified [98](#)
 - implementation on the storage system [98](#)
 - initiator security
 - method [69](#)
 - initiator security method [69](#)
 - interface access [69](#)
 - interface access lists to limit initiator interfaces [75](#)
 - iSNS [75](#)
 - license
 - adding [68](#)
 - verifying [68](#)
 - modifying
 - service [68](#)
 - target alias [68](#)
 - target node name [68](#)
 - modifying initiator security method [69](#)
 - new service [68](#)
 - nodes defined [98](#)
 - resolving storage system error messages [80](#)
 - security [73](#)
 - security policy [71](#)
 - service
 - commands [68](#)
 - deleting [68](#)
 - destroying [68](#)
 - starting [68](#)
 - stopping [68](#)
 - verifying [68](#)
 - showing [69](#)
 - stopping [69](#)
 - storage system interfaces [69](#)
 - troubleshooting [79](#)
- iSCSI
 - service
 - creating [68](#)
- iSCSI displaying
 - interface access [69](#)
- iSCSI error recovery levels
 - increasing Data ONTAP default [72](#)
- iSCSI initiator security commands [69](#)
- iSCSI service
 - deleting [72](#)
 - introduction to managing [68](#)
- iSCSI sessions
 - available commands [70](#)
 - deleting [70](#)
 - displaying [70](#)
 - showing [70](#)
 - shutdown [70](#)
- iSNS
 - access lists [75](#)

- communication [75](#)
 - defined [75](#)
 - discovery [75](#)
 - host [75](#)
 - iSCSI service [75](#)
 - LIF [75](#)
 - registration [75](#)
 - server versions [77](#)
 - service
 - adding [78](#)
 - configuration [78](#)
 - configuring [78](#)
 - creating [78](#)
 - deleting [78](#)
 - displaying [78](#)
 - modifying [78](#)
 - removing [78](#)
 - showing [78](#)
 - stopping [78](#)
 - updating [78](#)
 - iSNS server
 - registration requirement [75](#)
 - iSNS servers
 - registering SVMs with [77](#)
 - iSNS service
 - available commands [78](#)
- J**
- jumbo frames
 - recommended MTU configurations for FCoE [84](#)
- L**
- license
 - FC [100](#)
 - iSCSI [68](#)
 - verifying
 - FC [20](#)
 - iSCSI [20](#)
 - LIF
 - port set
 - removal [45](#)
 - removing IIF
 - from port sets [45](#)
 - LIFs
 - adding [68](#)
 - adding FC [81](#)
 - adding iSCSI [68](#)
 - commands
 - FC [81](#)
 - iSCSI [68](#)
 - considerations for, in cluster SAN environments [92](#)
 - creating as part of enabling block access for SVMs with FC [21](#)
 - creating when enabling block access for SVMs with iSCSI [21](#)
 - deleting [47, 68, 81](#)
 - FC [81](#)
 - FC and FCoE zoning restrictions for Cisco switches [47](#)
 - iSCSI [68](#)
 - managing [45](#)
 - migration considerations for SAN [45](#)
 - modifying [68, 81](#)
 - moving SAN [46](#)
 - Linux
 - I/O alignment considerations [37](#)
 - logical interface [73](#)
 - logical interfaces
 - changing WWPNs for FC [82](#)
 - login
 - initiator
 - checks [100](#)
 - LUN mapping
 - using Selective LUN Map [26](#)
 - LUN maps
 - how to determine whether SLM is enabled on [27](#)
 - LUN paths
 - decreasing using SLM [27](#)
 - LUNs
 - about restoring in SnapVault backup SAN environments [58](#)
 - accessing read-only copies from SnapVault backups [59](#)
 - adding, command for [33](#)
 - available commands for managing [33](#)
 - backing up through host backup systems [65](#)
 - bringing online, command for [33](#)
 - calculating rate of data change for [16](#)
 - capabilities and restrictions of transitioned [29](#)
 - cloning from active volumes [56](#)
 - command for adding [33](#)
 - command for changing serial number [33](#)
 - commands for managing [33](#)
 - configuration settings for non-space-reserved with thin volume provisioning [11](#)
 - configuration settings for space reserved with semi-thick volume provisioning [11](#)
 - configuration settings for space-reserved with thick-provisioned volumes [10](#)
 - considerations for copying [30](#)
 - controlling and monitoring I/O performance using Storage Qos [37](#)
 - creating and mapping igroups [22](#)
 - creating FlexClone, from Snapshot copies [57](#)
 - creating from files [25](#)
 - creating, command for [33](#)
 - decreasing the size of [31](#)
 - deleting [32](#)
 - deleting, command for [33](#)
 - determining the correct configuration for your environment [12](#)
 - displaying configurations, command for [33](#)
 - displaying information about, command for [33](#)
 - displaying mapping information, command for [33](#)
 - displaying serial numbers, command for [33](#)
 - displaying statistics, command for [33](#)
 - enabling space allocation for SCSI thinly provisioned [40](#)
 - examining configured and used space [33](#)
 - example of how access to is given with igroups [43](#)
 - guidelines for assigning IDs [25](#)
 - guidelines for mapping to igroups [25](#)
 - how access works in a virtualized environment [90](#)

- how moving or copying effects Snapshot copies [50](#)
- I/O alignment considerations for Solaris [37](#)
- I/O alignment for Linux [37](#)
- I/O misaligned warnings might be reported on properly aligned [35](#)
- increasing the size of [30](#)
- introduction to guidelines for SAN environments [25](#)
- introduction to managing [26](#)
- introduction to setting up [19](#)
- introduction to using space management capabilities for [7](#)
- maximum size, command for displaying [33](#)
- modifying, command for [33](#)
- moving across volumes [31](#)
- prerequisites for setting up [20](#)
- recommended configuration combinations with volumes [9](#)
- relocating, command for [33](#)
- renaming, command for [33](#)
- requirement for moving volumes that contain, in SAN environments [15](#)
- resizing, command for [33](#)
- restoring all in a volume from a Snapshot copy [51](#)
- restoring all in a volume, from SnapVault backup [62](#)
- restoring single LUN from a Snapshot copy [50](#)
- restoring single, from SnapVault backup [60](#)
- SCSI thin provisioned automatic host-side space management [40](#)
- setup workflow [18](#)
- sizing, command for [33](#)
- steps for setting up [19](#)
- taking offline, command for [33](#)
- thin provisioned, defined [7](#)
- tools available to effectively monitor [38](#)
- unmapping from igroups, command for [33](#)
- volume provisioning options for containing [8](#)
- ways to address offline issues [38](#)
- ways to limit access in a virtualized environment with port sets and igroups [28](#)
- what it means to overcommit aggregates that supply storage to [8](#)

LUNs, FlexClone

- data protection, how you use for [52](#)

LUNs, iSCSI

- troubleshooting when not visible on host [79](#)

M

managing host

- simplified using SnapDrive [97](#)

mapping

- LUNs to igroups, command for [33](#)

methods

- data protection, in SAN environments [49](#)

Microsoft copy offload

- using for storage virtualization [90](#)

Microsoft Offloaded Data Transfer

- See* ODX

migrations

- considerations for SAN LIF [45](#)

modifying

- FC adapters, commands for [84](#)
- iSCSI initiator security method [69](#)

- LUN comments, command for [33](#)

- LUN serial numbers, command for [33](#)

- security authentication method [69](#)

modules, optical

- changing for CNA/UTA2 target adapters [88](#)

monitoring

- configured and used space of LUNs [33](#)

moving

- LUNs across volumes [31](#)

- LUNs, command for [33](#)

- volumes, requirement in SAN environments [15](#)

MPIO

- considerations for LIFs in cluster SAN environments [92](#)

MTU configurations

- recommended for FCoE jumbo frames [84](#)

multiple paths

- reducing number with Selective LUN Map [26](#)

N

networks

- configuring for best iSCSI performance [71](#)

new

- FC service [81](#)

- iSCSI initiator security [69](#)

- port sets [34](#)

node name

- storage system [99](#)

node names

- how to specify for an igroup [43](#)

node type designator

- eui [99](#)

- iqn [98](#)

nodes

- considerations for adding to a cluster [48](#)

- FCP [100](#)

- iSCSI [98](#)

non-space-reserved LUNs

- configuration settings for thin volume provisioning [11](#)

O

ODX

- improving remote copy performance with [94](#)

- requirements [96](#)

- requirements for using [94](#)

- use cases for [95](#)

Offloaded Data Transfer

- See* ODX

onboard FC ports

- configuring for initiator mode [85](#)

optical modules

- changing for CNA/UTA2 target adapters [88](#)

optimized paths

- how Data ONTAP uses ALUA to identify [44](#)

options, configuration

- SAN volume [14](#)

OS types

- how to ensure I/O alignment using [36](#)

out of space errors

- possible with zero fractional reserve [15](#)
- overcommitment
 - defined [8](#)
- overviews
 - host-side space management [39](#)

P

- paths
 - how Data ONTAP uses ALUA to identify optimized and unoptimized [44](#)
- performance
 - configuring iSCSI networks for best [71](#)
 - improving remote copy performance with ODX [94](#)
- port configurations
 - supported for X1143A-R6 adapters [88](#)
- port set
 - LIF [45](#)
- port sets
 - adding [34](#)
 - commands [34](#)
 - creating [34](#)
 - creating and binding to igroups [24](#)
 - deleting [34](#)
 - displaying [34](#)
 - removing [34](#)
 - ways to limit LUN access in a virtualized environment with [28](#)
- ports
 - displaying [34](#)
- prerequisites
 - LUN setup [20](#)
- protection, data
 - methods o, in SAN environments [49](#)
- protocols
 - used by hosts to connect to SAN storage systems [97](#)
- provisioning options
 - volume [8](#)

R

- rate of data growth
 - calculating for LUNs [16](#)
- reducing LUN paths
 - using SLM [27](#)
- removing
 - initiator security list [69](#)
 - LUNs, command for [33](#)
 - port set [34](#)
- removing initiators from igroups
 - command for [42](#)
- renaming igroups
 - command for [42](#)
- requirements
 - for using ODX [94](#)
- restoring
 - all LUNs in a volume from SnapVault backup [62](#)
 - single LUN from SnapVault backup [60](#)

S

SAN

- considerations for transitioning from 7-Mode to clustered Data ONTAP [29](#)
- introduction to LUN guidelines for [25](#)
- SVM disaster recovery types [67](#)
- SAN configuration limits
 - volume and SVM requirements [14](#)
- SAN environments
 - about configuring and using SnapVault backups in [58](#)
 - advantages of virtualized [90](#)
 - methods of data protection in [49](#)
- SAN environments, cluster
 - considerations for LIFs in [92](#)
- SAN hosts
 - enabling block access for [23](#)
- SAN LIF migrations
 - considerations for [45](#)
- SAN LIFs
 - moving [46](#)
- SAN storage systems
 - protocols that hosts use to connect to [97](#)
- SAN systems
 - backing up LUNs through host backup [65](#)
- SAN volumes
 - configuration options [14](#)
- SCSI
 - thin provisioning, host support for [41](#)
- SCSI thin provisioning
 - defined [7](#)
 - host support for [41](#)
- SCSI thinly provisioned LUNs
 - automatic host-side space management enabled with [40](#)
 - enabling space allocation for [40](#)
- security authentication method
 - modifying [69](#)
- security method
 - defining [71](#)
 - iSCSI [71](#)
- security password
 - defining [69](#)
- Selective LUN Map
 - decreasing mapped LUN paths using [27](#)
 - defined [26](#)
 - how to determine whether it is enabled on LUN maps [27](#)
 - modifying reporting-nodes lists [28](#)
- semi-thick volume provisioning
 - configuration settings for space-reserved LUNs with [11](#)
- session
 - checks [100](#)
- single points of failure
 - requirement for moving volumes in SAN environments to avoid [15](#)
- SLM
 - decreasing mapped LUN paths using [27](#)
 - how to determine whether it is enabled on LUN maps [27](#)
 - modifying reporting-nodes list [28](#)
- Snapshot autodelete
 - SAN volume configuration option [14](#)
- Snapshot copies

- deleting [52](#)
 - effect of moving or copying a LUN on [50](#)
 - restoring a single LUN from [50](#)
 - restoring all LUNs in a volume from [51](#)
 - using to create FlexClone LUNs [57](#)
 - Snapshot reserve
 - SAN volume configuration option [14](#)
 - SnapVault backups
 - about configuring and using in SAN environments [58](#)
 - accessing read-only LUN copies from [59](#)
 - restoring all LUNs in a volume from [62](#)
 - restoring single LUNs from [60](#)
 - Solaris LUNs
 - I/O alignment considerations [37](#)
 - space
 - examining LUN configured and used [33](#)
 - space allocation
 - enabling for SCSI thinly provisioned LUNs [40](#)
 - space management
 - introduction to using capabilities for LUNs [7](#)
 - space-reserved LUNs
 - configuration settings for semi-thick volume provisioning [11](#)
 - configuration settings for thick-provisioned volumes [10](#)
 - speeds, FC adapter
 - changing [87](#)
 - Storage QoS
 - using for controlling and monitoring I/O performance to LUNs [37](#)
 - storage system node name
 - defined [99](#)
 - storage systems
 - resolving iSCSI error messages on [80](#)
 - storage systems, primary
 - how you can connect host backup systems to [64](#)
 - storage virtualization
 - using VMware and Microsoft copy offload [90](#)
 - suggestions
 - how to send feedback about documentation [104](#)
 - SVM
 - commands
 - FC [81](#)
 - iSCSI [68](#)
 - deleting [68, 72, 83](#)
 - FC
 - service [83](#)
 - FC traffic [100](#)
 - iSCSI
 - service [72](#)
 - iSCSI service [68](#)
 - iSNS [75](#)
 - SVM disaster recovery
 - available types in SAN [67](#)
 - SVMs
 - how high availability of iSCSI traffic is maintained [100](#)
 - registering with iSNS server [77](#)
 - restrictions for SAN [14](#)
 - with FC, enabling block access for [21](#)
 - with iSCSI, enabling block access [21](#)
 - switch
 - FC and FCoE zoning requirement [47](#)
 - switches
 - configuring for FCoE [19](#)
- ## T
- target adapters
 - changing WWPNs for FC logical interfaces [82](#)
 - displaying information about FC [87](#)
 - target adapters, unified
 - changing optical modules [88](#)
 - target mode
 - configuring FC adapters for [86](#)
 - TCP port
 - default for iSCSI [100](#)
 - thin provisioning
 - for LUNs, defined [7](#)
 - introduction to how Data ONTAP provides the ability to use [7](#)
 - SCSI, host support for [41](#)
 - volumes, explained for [7](#)
 - tools
 - available to monitor your LUNs [38](#)
 - transition considerations
 - SAN 7-Mode to clustered Data ONTAP [29](#)
 - transitioned LUNs
 - capabilities and restrictions [29](#)
 - troubleshooting
 - iSCSI LUNs not visible on the host [79](#)
 - iSCSI storage system error messages [80](#)
 - troubleshooting iSCSI [79](#)
 - Twitter
 - how to receive automatic notification of documentation changes [104](#)
- ## U
- unbinding
 - igroups [34](#)
 - unbinding igroups
 - command for [42](#)
 - unified target adapters
 - changing optical modules [88](#)
 - configuring ports [88](#)
 - viewing settings [89](#)
 - UNIX hosts
 - managing with SnapDrive [97](#)
 - unoptimized paths
 - how Data ONTAP uses ALUA to identify [44](#)
 - use cases
 - for ODX [95](#)
 - used space
 - examining LUN [33](#)
- ## V
- VAAI features
 - copy offload [93](#)
 - extended copy feature [93](#)
 - VERIFY AND WRITE feature [93](#)
 - WRITE SAME feature [93](#)
 - VERIFY AND WRITE feature

- environment [93](#)
- invoked automatically [93](#)
- see COMPARE WRITE feature [93](#)
- verifying
 - FC license [20](#)
 - iSCSI license [20](#)
- virtualized environments
 - how LUN access works in [90](#)
 - ways to limit LUN access with port sets and igroups [28](#)
- virtualized SAN environments
 - advantages of [90](#)
- VMware copy offload
 - using for storage virtualization [90](#)
- volume
 - deleting
 - Snapshot copies [52](#)
- volumes
 - cloning LUNs from active [56](#)
 - considerations for setting fractional reserve for FlexVol [15](#)
 - determining the correct configuration for your environment [12](#)
 - FlexVol, configuration settings for space-reserved LUNs with semi-thick volume provisioning [11](#)
 - FlexVol, configuration settings for thick provisioned with space-reserved LUNs [10](#)
 - FlexVol, configuration settings for thin-provisioned LUNs with thin volume provisioning [11](#)
 - move requirement in SAN environments [15](#)
 - moving LUNs across [31](#)
 - overcommitment, defined [8](#)
 - provisioning options for [8](#)
 - recommended configuration combinations with LUNs [9](#)
 - restoring all LUNs from a Snapshot copy [51](#)
 - restrictions for SAN [14](#)
 - SAN configuration options [14](#)
 - thin-provisioned, defined [7](#)

- what it means to overcommit aggregates that supply storage to FlexVol [8](#)

W

- Windows hosts
 - managing with SnapDrive [97](#)
- workflows
 - LUN setup [18](#)
- worldwide name assignments
 - how they can appear to be non-sequential [83](#)
- WRITE SAME feature
 - environment [93](#)
 - invoked automatically [93](#)
 - VAAI feature [93](#)
- WWNNs
 - how name assignments can appear to be non-sequential [83](#)
- WWPN
 - assigning [82](#)
 - assignment [84](#)
 - displaying [82](#)
 - usage [101](#)
- WWPNs
 - changing for FC logical interfaces [82](#)
 - how name assignments can appear to be non-sequential [83](#)
 - how to specify for an igroup [43](#)

X

- X1133A-R6 adapters
 - how to prevent loss of connectivity when using [89](#)
- X1143A-R6 adapters
 - supported port configurations [88](#)

Z

- zoning
 - requirements for Cisco switches [47](#)