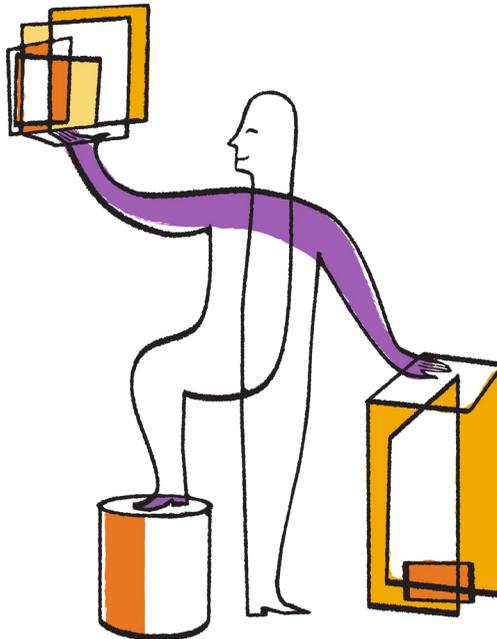




Updated for 8.2.1

Clustered Data ONTAP® 8.2

SAN Administration Guide



NetApp, Inc.
495 East Java Drive
Sunnyvale, CA 94089
U.S.

Telephone: +1 (408) 822-6000
Fax: +1 (408) 822-4501
Support telephone: +1 (888) 463-8277
Web: www.netapp.com
Feedback: doccomments@netapp.com

Part number: 215-08513_B0
February 2014

Contents

Preparing for LUN setup workflow	9
LUN setup workflow	10
Deciding which LUN type to use in SAN environments	11
How the rate of change in your data determines LUN type	11
Calculating Rate of Change	12
How much room do you need for Snapshot copies	13
How LUNs consume space on a volume	14
Strategies for keeping thinly provisioned LUNs online	14
What LUN thin provisioning is	15
When to use thinly provisioned LUNs	16
Why thinly provisioned LUNs go offline	16
How to keep your thinly provisioned LUNs online	17
What to do if your LUNs go offline	26
Additional considerations for space-reserved LUNs	27
What space-reserved LUNs are	27
When to use space-reserved LUNs without Snapshot reserve	27
When to use space-reserved LUNs with Snapshot reserve	28
Space-reserved LUNs in a thinly provisioned volume	28
Understanding space usage of LUNs and Snapshot copies	29
Showing disk space usage for LUNs without Snapshot copies	29
Showing disk space usage for LUNs with Snapshot copies	32
SAN configurations always use FlexVol volumes	38
Volume configuration options	39
Considerations for setting fractional reserve for FlexVol volumes	39
What Snapshot autodelete is	40
What Snapshot reserve is	41
How Data ONTAP can automatically provide more space for full FlexVol volumes	41
Selecting the first method to increase space for full FlexVol volumes	42
Configuring volumes in a SAN environment	43
Configuring volumes for thinly provisioned LUNs	43
Configuring volumes for space-reserved LUNs without Snapshot reserve	46

Configuring volumes for space-reserved LUNs with Snapshot reserve	50
LUN setup	54
LUN guidelines for SAN environments	54
Recommended and supported SAN configurations	55
Considerations when creating a LUN from an existing file	55
Considerations about LUN identification numbers	55
Guidelines for provisioning storage in a SAN environment	56
Guidelines for mapping LUNs to igroups	56
Guidelines for LUN layout and space allocation	57
Understanding LUN functions	57
How file and LUN reservations work for FlexVol volumes	57
How the host operating system affects usable space for a LUN	58
LUN availability on specific target LIFs	58
How Data ONTAP avoids igroup mapping conflicts	59
SnapMirror destinations and read-only LUNs	59
Introduction to LUNs in a virtualized SAN environment	60
How LUN access works in a virtualized environment	60
How to limit LUN access in a virtualized environment	63
Considerations for LIFs in cluster SAN environments	63
Setting up LUNs	63
Setting up LUNs manually with iSCSI block access	64
Setting up LUNs manually with FC block access	71
Setting up LUNs using the SVM setup wizard	79
Managing data protection in your SAN environment	81
Data protection methods	81
Copying all LUNs on a volume with the volume copy command	83
Using Snapshot copies to protect your data	84
Required changes to Snapshot copy default settings	85
Restoring a single LUN from a Snapshot copy	85
Restoring all the LUNs in a volume from a Snapshot copy	87
Deleting one or more existing Snapshot copies in a volume	89
Using FlexClone LUNs to protect your data	90
Reasons for using FlexClone LUNs	90
What FlexClone LUN autodelete is	91
How a FlexVol volume can reclaim free space from FlexClone LUNs	91
Configuring a FlexVol volume to automatically delete FlexClone LUNs	92

Cloning LUNs from an active volume	94
Cloning LUNs from a Snapshot copy in a volume	95
Preventing a specific FlexClone LUN from being autodeleted	96
Configuring and using SnapVault backups in a SAN environment	98
Accessing a read-only LUN copy from a SnapVault backup	99
Restoring a single LUN from a SnapVault backup	100
Restoring all LUNs in a volume from a SnapVault backup	102
How you can connect a host backup system to the primary storage system	105
Backing up a LUN through a host backup system	106
Managing LUNs	109
Misaligned I/O can occur on properly aligned LUNs	109
The effect of ostype on I/O alignment	110
Special I/O alignment considerations for Linux	111
Special I/O alignment considerations for Solaris	111
Controlling and monitoring I/O performance to LUNs by using Storage QoS	112
Increasing the size of a LUN	113
Decreasing the size of a LUN	114
Deleting LUNs	114
Examining configured and used space of a LUN	115
Commands for managing LUNs	116
Commands for managing port sets	117
Managing LIFs	118
LIF migration in a SAN environment	118
Removing a SAN LIF from a port set	119
Moving SAN LIFs	119
Deleting a LIF in a SAN environment	121
FC and FCoE LIFs on the same port need to be in separate zones	122
Understanding igroup functions	123
What igroups are	123
igroup example	124
About iSCSI initiator node names	125
FC protocol initiator WWPN	125
What ALUA is	125
MPIO and ALUA	126
Managing igroups	127
Commands for managing igroups	127

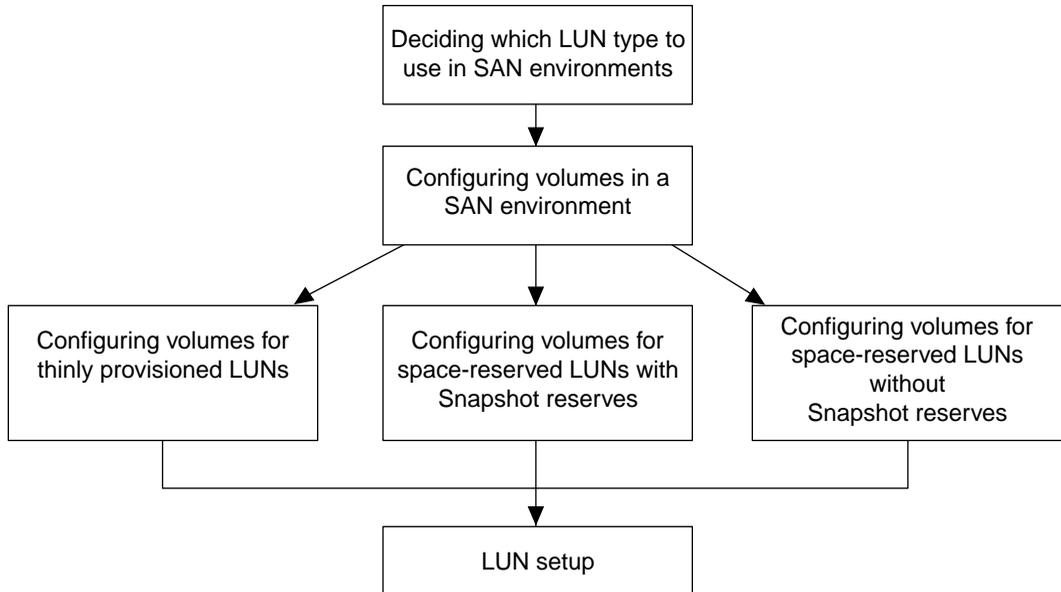
Configuring your iSCSI network for best performance	128
Managing your iSCSI service	129
Defining a security policy method for an initiator	129
Enabling iSCSI session error recovery levels 1 and 2	130
Deleting an iSCSI service for an SVM	131
Commands for managing iSCSI services	131
iSCSI service management	135
How iSCSI is implemented on the host	135
How iSCSI authentication works	135
iSCSI initiator security management	136
What CHAP authentication is	136
Guidelines for using CHAP authentication	136
iSCSI interface access management	137
iSNS server registration	138
What iSNS is	138
What an iSNS server does	138
How SVMs interact with an iSNS server	138
About iSNS service version incompatibility	141
Registering an SVM with an iSNS server	141
Commands for managing iSNS	143
iSCSI troubleshooting tips	144
LUNs not visible on the host	144
Resolving iSCSI error messages on the storage system	145
Understanding FC and FCoE SAN functions	147
How FC switches are identified	147
How worldwide name assignments work	148
Fibre Channel over Ethernet overview	148
Unified Ethernet network management	149
Data center bridging	149
Support for iSCSI DCB	150
Verifying that DCB is configured correctly	150
Managing systems with Fibre Channel adapters	152
Configuring FC adapters for initiator mode	152
Configuring FC adapters for target mode	154
Displaying the status of onboard FC adapter	156
Displaying information about an FC target adapter	156

Changing the FC adapter speed	158
Commands for managing FC adapters	159
Managing systems with the X1143A-R6 adapter	161
Function of the X1143A-R6 adapter	161
Managing the X1143A-R6 adapter	162
Default configuration for the X1143A-R6	163
Configuration rules and supported settings	163
Managing your FC service	165
Changing the WWPN for an FC logical interface	165
Deleting an FC service for an SVM	166
Commands for managing FC protocols	166
Volume management	169
Before performing a volume move in SAN	170
Virtualization of storage	171
LUN alignment in virtual environments	171
Improving VMware VAAI performance for ESX hosts	172
Requirements for using the VAAI environment	172
How to determine if VAAI features are supported by ESX	173
Microsoft Offloaded Data Transfer (ODX)	173
Requirements for using ODX	174
Use cases for ODX	175
Special system file requirements	176
Basic iSCSI and FC concepts	177
How hosts connect to storage systems	177
What Host Utilities are	177
About SnapDrive for Windows and UNIX	178
How Data ONTAP implements an iSCSI network	178
What iSCSI is	178
What iSCSI nodes are	179
How iSCSI target nodes connect to the network	179
How iSCSI nodes are identified	179
Storage system node name	180
How the storage system checks initiator node names	180
Default port for iSCSI	180
How iSCSI communication sessions work	181
How iSCSI works in a Storage Virtual Machine (SVM)	181

How Data ONTAP implements an FC SAN	181
What FC is	181
What FC nodes are	182
How FC target nodes connect to the network	182
How FC nodes are identified	182
How WWPNs are used	182
Copyright information	183
Trademark information	184
How to send your comments	185
Index	186

Preparing for LUN setup workflow

A prerequisite for creating LUNs is that you have an aggregate, a Storage Virtual Machine (SVM), and a volume. Also before you can begin LUN setup, you must decide what type of LUN that you need for your SAN environment and configure your volume for that LUN type.

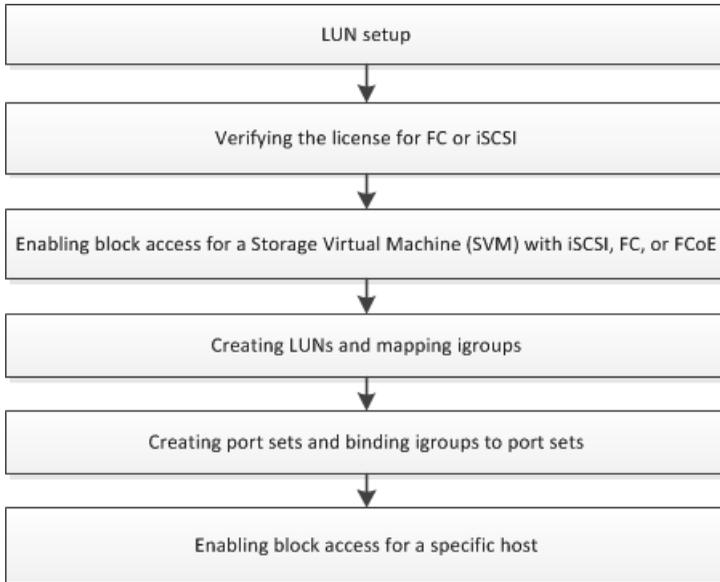


Related concepts

[Deciding which LUN type to use in SAN environments](#) on page 11

LUN setup workflow

After you have configured your volume for a specific LUN type, you can begin configuring (setting up) your LUNs. Configuration involves several steps.



Related concepts

[LUN setup](#) on page 54

Deciding which LUN type to use in SAN environments

You should decide how you want to allocate space for LUNs and Snapshot copies before you configure your volume or set up your LUNs. Do you want to reserve space ahead of time (space-reserved LUNs), or do you want to allocate space as needed (thinly provisioned LUNs)?

You can reserve space up front or add space as needed for LUNs and Snapshot copies in your volume. You should answer the following questions to determine the types of LUNs and Snapshot copies that work best in your environment:

- Do you want to allocate space on your volume as needed for your LUNs and Snapshot copies?
- Do you want to reserve space on your volume for your LUNs and Snapshot copies?
- Do you want to reserve space on your volume for your LUNs but allocate space as needed for your Snapshot copies?
- How closely do you need to monitor your environment?
- Will the amount of data in your LUNs grow quickly?

How you answer these questions determines which of the three common usage scenarios for allocating space in your volume for your LUNs and Snapshot copies works best for your environment. The three common usage scenarios are as follows:

- Thinly provisioned LUNs
- Space-reserved LUNs without Snapshot reserve
- Space-reserved LUNs with Snapshot reserve

Related concepts

[Strategies for keeping thinly provisioned LUNs online](#) on page 14

[Additional considerations for space-reserved LUNs](#) on page 27

How the rate of change in your data determines LUN type

The rate of change in your data helps you determine what type of LUN best suits your environment, space-reserved LUN or thinly provisioned LUN.

Rate of deletion or change of data	Notes
Lots of deletes (high rate of change)	<ul style="list-style-type: none"> • Use space-reserved LUNs • Need extra room for Snapshot copies • Set fractional reserve to 100%

Rate of deletion or change of data	Notes
Very low rate of change (low rate of change)	<ul style="list-style-type: none"> • Use thinly provisioned LUNs
Steady growth (low rate of change)	<ul style="list-style-type: none"> • Use thinly provisioned LUNs • Use volume autogrow
Inconsistent growth	<ul style="list-style-type: none"> • Use thinly provisioned LUNs • Use volume autogrow to ensure room for LUNs • Use Snapshot autodelete • Use space reclamation if possible

Calculating Rate of Change

You will need to know the rate at which your data is changing over time to determine whether you should use space-reserved LUNs or thinly provisioned LUNs.

About this task

If you have a consistently high rate of data change, then space-reserved LUNs might be a better option for you. If you have a low rate of data change, then you should consider leveraging the advantages of thin provisioning. You will need to observe your data over a set period of time to determine your rate of change as accurately as possible.

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.
3. Each day, record in GB how much your data changes.
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 change.

You need a 200 GB LUN and are trying to determine if it should be a space-reserved LUN or a thinly provisioned 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 change is $(20+18+17+20+20+23+22) / 7 = 20$ GB per day.

How much room do you need for Snapshot copies

The longer you need to keep Snapshot copies, the more space you need to set aside for your Snapshot copies.

You can use volume autosize to automatically grow your volume to ensure you have enough space for your LUN and Snapshot copies. You can also use Snapshot autodelete to remove Snapshot copies.

For more information about Snapshot autodelete, see the *Clustered Data ONTAP Data Protection Guide*. For more information about volume autosize, see the *Clustered Data ONTAP Logical Storage Management Guide*.

How long do you need a Snapshot copy for?	Notes
Lots of Snapshot copies for a long duration	<ul style="list-style-type: none"> • Leave a large percentage of room for Snapshot copies • Use Snapshot autodelete aggressively • Higher the rate of change the more room you need
Some Snapshot copies for a long duration	<ul style="list-style-type: none"> • Use Snapshot reserve • Use volume autogrow before snapshot autodelete
Some Snapshot copies for a short duration	<ul style="list-style-type: none"> • Use Snapshot without reserve • Leave a small percentage of room for Snapshot copies. • Ensure snapshot autodelete does not delete the SnapMirror copies. • Use volume autogrow before snapshot autodelete

How LUNs consume space on a volume

Thinly provisioned and space-reserved LUNs consume space on a volume differently.

When you create a thinly provisioned LUN, this LUN consumes almost no space in its containing volume. Data written to the thinly provisioned LUN is allocated from the volume's free space, and an equal amount of space is added to the LUN's used size.

In contrast, when you create a space-reserved LUN, the amount of space equal to the LUN's configured size is removed from the free space of the volume immediately upon creation.

The used size of a LUN cannot be larger than the configured size.

Related tasks

[Enabling space allocation for thinly provisioned LUNs](#) on page 23

Strategies for keeping thinly provisioned LUNs online

There are various options available in Data ONTAP that help minimize the risk of your thinly provisioned LUNs running out of space and going offline.

These options can give a thinly provisioned LUN the space it needs to grow by automatically making more space available on the volume containing the LUN.

- volume autogrow
- Snapshot auto delete
- FlexClone LUN automatic deletion
- automatic space reclamation

For more information about these strategies, see the *Clustered Data ONTAP Logical Storage Management Guide*.

Related concepts

[How Data ONTAP can automatically provide more space for full FlexVol volumes](#) on page 41

[What Snapshot autodelete is](#) on page 40

[What FlexClone LUN autodelete is](#) on page 91

What LUN thin provisioning is

Thin provisioning enables storage administrators to provision more storage on a LUN than is currently available on the volume. Users often do not consume all the space they request, which reduces storage efficiency if space-reserved LUNs are used.

By over-provisioning the volume, storage administrators can increase the capacity utilization of that volume. When a new thinly provisioned LUN is created, it consumes almost no space from the containing volume. As blocks are written to the LUN and space within the LUN is consumed, an equal amount of space within the containing volume is consumed.

With thin provisioning, you can present more storage space to the hosts connecting to the Storage Virtual Machine (SVM) than is actually available on the SVM. Storage provisioning with thinly provisioned LUNs enables storage administrators to provide users with the storage they need at any given time.

The advantages of thin provisioning are as follows:

- Provides better storage efficiency.
- Allows free space to be shared between LUNs.
- Enables LUNs to consume only the space they actually use.

Example of a volume with thinly provisioned LUNs

An administrator can provision a 4,000-GB volume with five thinly provisioned LUNs with 1,000 GB of space for each LUN as shown in the following table.

Table 1: Thinly provisioned LUNs on a 4,000-GB volume

LUN name	Space actually used by the LUN	Configured space available to the LUN
lun1	100 GB	1,000 GB
lun2	100 GB	1,000 GB
lun3	100 GB	1,000 GB
lun4	100 GB	1,000 GB
lun5	100 GB	1,000 GB
Totals	500 GB	5,000 GB

All 5 LUNs use 100 GB of storage, but each LUN has the possibility of using 1,000 GB of storage. In this configuration, the volume is overcommitted by 1,000 GB, but because the actual space used by the LUNs is 500 GB, the volume still has 3,500 GB available space. Thin provisioning allows LUNs to grow at different rates. From the pool of available space, a LUN can grow as blocks of data are written to that LUN.

If all the LUNs used all their configured space, then the volume would run out of free space. The storage administrator needs to monitor the SVM and increase the size of the volume as needed.

You can have thinly provisioned and space-reserved LUNs on the same volume and the same aggregate. For example, you can use space-reserved LUNs for critical production applications, and thin provisioned LUNs for other types of applications.

When to use thinly provisioned LUNs

Thinly provisioned LUNs provide the most flexibility for storage utilization because they do not reserve space; instead, space is only allocated when data is written to the LUN.

You must closely monitor the available space in the aggregate containing the volume because a thinly provisioned LUN configuration can oversubscribe the available space. You can use the volume configuration and volume autosize settings to enable your LUNs grow automatically.

The typical use case for thinly provisioned LUNs without Snapshot reserve involves shared storage infrastructures, test, or development environments. Because utilization rates can be unpredictable, these environments benefit from flexible space allocation for LUNs and Snapshot copies.

For more information, see the technical reports on thin provisioning and storage efficiency TR-3827 and TR-3563.

Related concepts

[What LUN thin provisioning is](#) on page 15

[Volume option best practices for thinly provisioned LUNs](#) on page 17

[How Data ONTAP can automatically provide more space for full FlexVol volumes](#) on page 41

Related tasks

[Configuring volumes for thinly provisioned LUNs](#) on page 43

Related information

[If You're Doing This, Then Your Storage Could Be Underutilized: media.netapp.com/documents/tr-3827.pdf](http://media.netapp.com/documents/tr-3827.pdf)

[NetApp Thin Provisioning Increases Storage Utilization With On Demand Allocation: media.netapp.com/documents/tr-3563.pdf](http://media.netapp.com/documents/tr-3563.pdf)

Why thinly provisioned LUNs go offline

If a thinly provisioned LUN has no available space to accept a write, Data ONTAP takes this LUN offline to maintain data integrity. Free space must be available on the volume before you can bring this LUN back online.

You can add more space on your volume in the following ways:

- Manually add free space to the volume

- Enable volume autogrow
- Enable Snapshot autodelete
- Enable FlexClone LUN automatic deletion

For more information about the snap autodelete command and the volume autosize command, see the *Clustered Data ONTAP Logical Storage Management Guide*.

Related concepts

[Using FlexClone LUNs to protect your data](#) on page 90

[Using Snapshot copies to protect your data](#) on page 84

How to keep your thinly provisioned LUNs online

When your LUNs are thinly provisioned and over-committed, you can use several strategies to prevent your LUNs from going offline.

You can use the following strategies to prevent your LUNs from going offline:

- **Follow the volume options best practices for thin provisioning.** When you follow the best practices, you automate the volume to grow or delete Snapshot copies as needed.
- **Monitor the available space on your volumes and aggregates.** You can monitor your SAN environment manually with System Manager or you can monitor your entire SAN environment automatically.
- **Use space reclamation when possible.** If a LUN contains deleted blocks, you can use space reclamation to put those blocks back into the general pool of storage.
- **Understand the rate of change of your data.** How much your data changes over time helps you determine what type of LUN would be most beneficial for your environment.

Related concepts

[Volume option best practices for thinly provisioned LUNs](#) on page 17

[How to monitor your LUNs effectively](#) on page 18

[Space Reclamation](#) on page 19

Related tasks

[Calculating Rate of Change](#) on page 12

Volume option best practices for thinly provisioned LUNs

Data ONTAP provides volume configuration options that make managing your thinly provisioned LUN easier.

In respect to the best practices presented here, thinly provisioned refers specifically to non-space reserved LUNs. When you configure a non-space reserved LUN, you can run out of space on the volume containing the LUN. To help minimize the risk of running out of space on a volume, Data ONTAP has the following configurable options that you can use:

- volume autosize - allows your volume to grow automatically.
- autodelete - deletes Snapshot copies automatically
- snap reserve - allocates space for Snapshot copies when needed

Volume option	best practice setting	command to use
volume autosize	on	volume autosize
Snapshot autodelete	true	volume snapshot autodelete modify
snap reserve	0	volume modify

You might need to adjust these volume options for your environment. For example, you might need to set the "Snapshot autodelete" option to `false` based on your business requirements.

For more information about thinly provisioned LUNs, see TR 3827 and TR 3563.

Related concepts

Considerations for setting fractional reserve for FlexVol volumes on page 39

How Data ONTAP can automatically provide more space for full FlexVol volumes on page 41

What Snapshot reserve is on page 41

What Snapshot autodelete is on page 40

How the rate of change in your data determines LUN type on page 11

Related tasks

Configuring volumes for space-reserved LUNs with Snapshot reserve on page 50

Configuring volumes for space-reserved LUNs without Snapshot reserve on page 46

Configuring volumes for thinly provisioned LUNs on page 43

Related information

If You're Doing This, Then Your Storage Could Be Underutilized - media.netapp.com/documents/tr-3827.pdf

NetApp Thin Provisioning Increases Storage Utilization With On Demand Allocation - media.netapp.com/documents/tr-3563.pdf

How to monitor your LUNs effectively

Using System Manager OnCommand Insight to manage your LUN resources can help prevent your LUNs from going offline due to lack of storage space on your LUNs, volumes, and aggregates.

- **System Manager 2.0 or higher:** enables you to manually manage storage needs at the cluster level. You can view if your storage is running out of space in the following ways:

- **Storage view**- to see if alarms are present indicating that the space on the aggregate is running low.
- **Vserver level** - to determine whether your volumes are running out of space.
- **OnCommand Insight**: presents a single view of your storage infrastructure. OnCommand Insight also provides tools to manage your environment automatically with policy-driven monitoring, alerts, and reporting when your LUNs, volumes, and aggregates are running out of storage space. See the *OnCommand Insight Getting Started Guide* for more information.

Space Reclamation

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

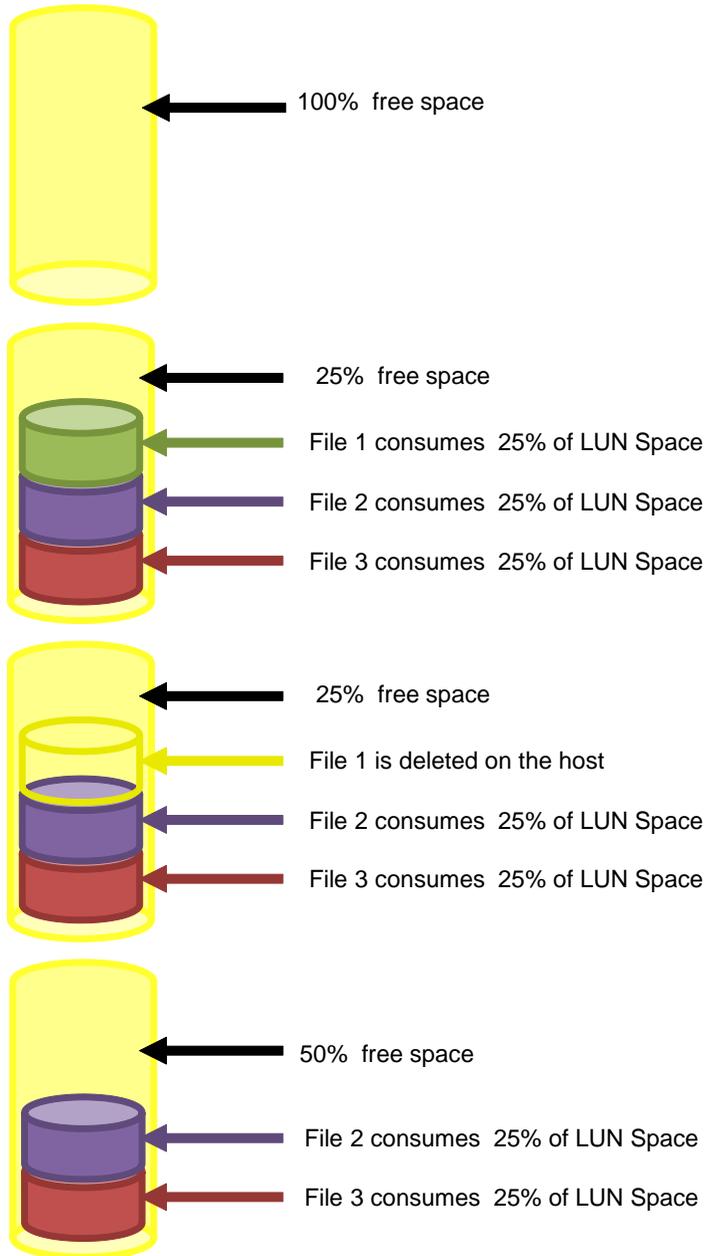
A host file system contains metadata to keep track of which blocks are available to store new data and which blocks contain valid data and must not be overwritten. This metadata is stored within the LUN. When a file is deleted in the host file system, the file system metadata is updated to mark that 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 than 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 and 100 GB of unused space. The blocks containing the 50 GB of data deleted by the host are not freed on the storage system until they are reclaimed through space reclamation.

One supported method of reclaiming space is built into SnapDrive for Windows. The SnapDrive implementation of space reclamation, called SpaceReclaimer, is a process that runs on the host. Each block in the file system is examined and compared against the corresponding block in the storage system. If the space reclamation process finds a block that is marked as free in the host file system, but not free in the storage system, the space reclamation process issues a special SCSI command to the storage system identifying which block can be freed. After the process has completed, the amount of free space reported by the host and the amount of free space inside the LUN as reported by the storage system will be identical.

Space Reclamation Process



Determining the volume size when using Snapshot autodelete

Before you create a volume for use with Snapshot autodelete, you should estimate how large it needs to be.

About this task

Snapshot autodelete is a volume-level option that allows you to automatically delete Snapshot copies when a pre-defined threshold called a "trigger" is met. You can set the trigger for autodelete when the volume is nearly full, when the snap reserve space is nearly full, or when the overwrite reserved space is consumed. Using Snapshot autodelete is recommended in most SAN configurations, but is particularly useful when:

- You do not want your volumes to automatically grow, because automatic growth consumes space in the aggregate.
- Ensuring availability of your LUNs is more important to you than maintaining old Snapshot copies.

Steps

1. Calculate the Rate of Change (ROC) of your data per day.

This value depends on how often you overwrite data. It is expressed as GB per day.

2. Calculate the amount of space you need for Snapshot copies by multiplying your ROC by the number of days of Snapshot copies you intend to keep.

Space required for Snapshot copies = ROC x number of days of Snapshot copies.

Example

You need a 200 GB LUN, and you estimate that your data changes at a rate of about 10 percent each day, which in this case is 20 GB per day. You want to take one Snapshot copy each day and want to keep three weeks' worth of Snapshot copies, for a total of 21 days of Snapshot copies. The amount of space you need for Snapshot copies is 21×20 GB, or 420 GB.

3. Calculate the required volume size by adding together the total LUN size and the space required for Snapshot copies.

The total LUN size = the size of all the LUNs in the volume.

Example

The following example shows how to calculate the size of a volume based on the following information:

- You need to create two 200 GB LUNs. The total LUN size is 400 GB.
- You take one Snapshot copy each day and you want to keep 10 days of Snapshot copies. This means you need 400 GB of space for Snapshot copies (40 GB ROC \times 10 Snapshot copies).

- You rate of change varies due to periodic increases. You do not want to add additional space to your volume to accommodate the variable rate. In this case, you can configure Snapshot autodelete with a volume space trigger to delete snapshots, so that space remains available for additional overwrites even when your rate of change increases more than usual.

You would calculate the size of your volume as follows:

Volume size = Total data size + Space required for Snapshot copies.

The size of the volume in this example is 800 GB (400 GB + 400 GB).

For more information about the autodelete function, see the *Clustered Data ONTAP Logical Storage Management Guide*, and for more information about working with FlexVol volumes, see the *Clustered Data ONTAP Physical Storage Management Guide* and *Clustered Data ONTAP Logical Storage Management Guide*.

Related information

Data ONTAP documentation on the NetApp Support Site-support.netapp.com

How host operating systems can automatically reclaim space and keep LUNs online

Starting with Data ONTAP 8.2, you can use the `space-allocation` option to reclaim space and notify the host when a thinly provisioned LUN cannot accept writes. The space allocation option enables the Logical Block Provisioning feature as defined in the SCSI SBC-3 standard.

When you enable the `space-allocation` on a thinly provisioned LUN, the following two SCSI features are enabled:

- **Reclaims space automatically when your host deletes data.** When a host that does not support the space allocation functionality deletes data on a LUN, the storage system is unaware of the deletion, which results in poor space efficiency. If the host supports this functionality, Data ONTAP reclaims this space automatically.

Hosts that support automatic space reclamation. The following hosts currently support automatic space reclamation when you enable space allocation:

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

See the host utilities documentation for more information about which hosts support automatic space reclamation.

Note: The space reclamation process issues one or more SCSI `UNMAP` commands to free blocks on the storage system after identifying the blocks that can be freed on the host file system.

- **Notifies the host when the LUN cannot accept writes due to lack of space on the volume.** On hosts that do not support the space allocation functionality, when the volume containing LUN runs out of space and cannot automatically grow, Data ONTAP takes the LUN offline. Free space must be available on the volume before you can bring the LUN back online.

However, when you enable the `space-allocation` option, Data ONTAP notifies the host when the volume containing the LUN is running out of space and cannot grow. If the LUN cannot accept writes due to the lack of space on the volume, the LUN stays online. The host cannot write to the LUN, but the host can still read the LUN.

You must add more space to your volume before the LUN can accept writes. You can add more space on your volume in the following ways:

- Manually add free space to the volume
- Enable volume autogrow
- Enable Snapshot autodelete
- Enable FlexClone LUN automatic deletion

Hosts that support out-of-space notifications. The following hosts currently support out-of-space notifications that a LUN cannot accept writes when you enable space allocation:

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

See the host utilities documentation for more information about which hosts support out-of-space notifications.

For more information about the `snap autodelete` command and the `volume autosize` command, see the *Clustered Data ONTAP Logical Storage Management Guide*.

Related concepts

[How to monitor your LUNs effectively](#) on page 18

Related tasks

[Enabling space allocation for thinly provisioned LUNs](#) on page 23

Related references

[Logical Block Provisioning feature of the SCSI SBC-3 standard](#) on page 25

Related information

[Host Utilities documentation: www.support.netapp.com](http://www.support.netapp.com)

Enabling space allocation for 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. The `space-allocation` option is set to disabled by default.

About this task

The space allocation functionality uses the Logical Block Provisioning feature as defined in the SCSI SBC-3 standard.

Steps

1. Use the `lun modify` command to take the LUN offline.

Example

```
lun modify -vserver vs1 -volume vol2 -lun lun1 -state offline
```

2. Use the `lun modify` command to set the `-space-allocation` parameter to enabled.

Example

```
lun modify -vserver vs1 -volume vol2 -lun lun1 -space-allocation enabled
```

3. Use the `lun show` command to verify that you enabled space allocation.

Example

```
lun show -vserver vs1 -volume vol2 -lun lun1 -fields space-allocation
```

```

vserver path                space-allocation
-----
vs1      /vol/vol2/lun1  enabled

```

4. Use the `lun modify` command to bring the LUN back online.

Example

```
lun modify -vserver vs1 -volume vol2 -lun lun1 -state online
```

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

After you finish

You must actively monitor your volume and aggregate to ensure enough space is available for your LUNs.

Related concepts

How host operating systems can automatically reclaim space and keep LUNs online on page 22
How to monitor your LUNs effectively on page 18

[Using FlexClone LUNs to protect your data](#) on page 90

[Using Snapshot copies to protect your data](#) on page 84

Related references

[Logical Block Provisioning feature of the SCSI SBC-3 standard](#) on page 25

Logical Block Provisioning feature of the SCSI SBC-3 standard

The space allocation functionality, also known as 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 the space allocation functionality in Data ONTAP.

When you enable the space allocation functionality, you turn on the following thin provisioning features for standard SCSI features:

- Unmapping and reporting space usage for space reclamation
- Reporting resource exhaustion errors
- Reporting low space warnings for thin provisioning thresholds

Related concepts

[How host operating systems can automatically reclaim space and keep LUNs online](#) on page 22

Related tasks

[Enabling space allocation for thinly provisioned LUNs](#) on page 23

What to do if your LUNs go offline

When no space is available for writes, thinly provisioned LUNs go offline to preserve data integrity. LUNs can run out of space for several reasons. The following table contains a list of what to check for.

	What you can do.
If the aggregate is full, you can:	<ul style="list-style-type: none"> • Add more disks. See the <i>Clustered Data ONTAP Physical Storage Management Guide</i> for more information. • Use the <code>volume modify</code> command to shrink a volume that has available space. See the man page for this command or the <i>Clustered Data ONTAP Logical Storage Management Guide</i> for more information. • If you have space-reserved volumes that have available space change the volume space guarantee to none or file with the <code>volume modify</code> command. For more information, see the <i>Clustered Data ONTAP Logical Storage Management Guide</i>.
If your volume is full but there is space available in the containing aggregate, you can:	<ul style="list-style-type: none"> • For space reserved volumes: Use the <code>volume modify</code> command to increase the size of your volume. See the man page for this command or the <i>Clustered Data ONTAP Logical Storage Management Guide</i> for more information. • 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. See the man page for this command or the <i>Clustered Data ONTAP Data Protection Guide</i> for more information.

Additional considerations for space-reserved LUNs

You can have space-reserved LUNs with or without Snapshot reserve. Additionally, you can have space-reserved LUNs on a thinly provisioned volume.

What space-reserved LUNs are

When you use space-reserved LUNs, the LUN space is reserved on the volume but not pre-allocated. Space is allocated only when data is written to the space-reserved LUN.

When to use space-reserved LUNs without Snapshot reserve

Space-reserved LUNs without Snapshot reserve remove the variable of LUN growth rate from space calculations because all the space for any given LUN is reserved for that LUN. Removing the LUN growth rate reduces the need to carefully monitor this environment.

LUNs have pre-allocated space, but Snapshot copies do not. Overwrites for the Snapshot copies are limited by available free space. Although space for Snapshot copies might be oversubscribed, space for active LUN data is already allocated and available to those LUNs.

In this scenario, large database environments would benefit from using space-reserved LUNs without Snapshot copies. These environments tend to have a low overall rate of change in LUN data, and a high or predictable utilization rate.

For more information about the storage efficiency of space-reserved LUNs without Snapshot reserve, see the technical report TR-3827 on storage efficiency.

Related concepts

[How file and LUN reservations work for FlexVol volumes](#) on page 57

Related tasks

[Configuring volumes for space-reserved LUNs without Snapshot reserve](#) on page 46

Related information

[If You're Doing This, Then Your Storage Could Be Underutilized - media.netapp.com/documents/tr-3827.pdf](http://media.netapp.com/documents/tr-3827.pdf)

When to use space-reserved LUNs with Snapshot reserve

Space-reserved LUNs and Snapshot copies have space reserved. This reserved space is not available to any other LUNs or Snapshot copies within the volume.

Pre-allocating space for LUNs and Snapshot copies is least efficient in terms of storage utilization because the configured size of the LUN or Snapshot copy reserve could be much larger than what is actually required. You do not need to monitor this configuration as closely as you do thinly provisioned LUNs or LUNs without Snapshot reserve because the space for the LUNs and Snapshot copies is guaranteed for those LUNs and Snapshot copies.

Small installations may benefit from space-reserved LUNs with Snapshot copies because it is often more important to guarantee the space for LUNs and Snapshot copies than to configure for maximum efficiency. For these environments, it is more efficient to guarantee space for a small number of LUNs and Snapshot copies beforehand, which also eases storage system monitoring requirements.

For more information about storage efficiency of space-reserved LUNs with Snapshot reserve, see the technical report TR-3827 on storage efficiency.

Related concepts

[How file and LUN reservations work for FlexVol volumes](#) on page 57

Related tasks

[Configuring volumes for space-reserved LUNs with Snapshot reserve](#) on page 50

Related information

[If You're Doing This, Then Your Storage Could Be Underutilized - media.netapp.com/documents/tr-3827.pdf](http://media.netapp.com/documents/tr-3827.pdf)

Space-reserved LUNs in a thinly provisioned volume

In thinly provisioned volumes, you can use space-reserved LUNs. However, if the thinly provisioned volume is over-committed on the aggregate, the amount of free space for Snapshot copies needs to be monitored carefully.

You can configure your volume to reserve space for your Snapshot copies from either the active file system or from the Snapshot reserve. However, because the space for the Snapshot reserve is available for all the LUNs in that volume, the Snapshot reserve space is not guaranteed.

For example, space-reserved LUN 1 and LUN 2 reside on volume A and both LUNs created Snapshot copies from the active file system. If LUN 1 has a higher rate of change or schedules more Snapshot copies than LUN 2, most of the active file system will be consumed by Snapshot copies from LUN 1. It is possible that LUN 2 might not have enough room for its Snapshot copies, especially if aggregate is over-committed.

Also if you have no overwrite protection, you can fill up your available free space with Snapshot copies. To prevent your available space from being consumed by Snapshot copies, you should do the following:

- Monitor your aggregate.
- Understand your rate of change in your data.
- Use Snapshot autodelete aggressively.
- Use FlexClone autodelete if appropriate.
- Consider overwrite protection for your LUNs.

Related concepts

[How the rate of change in your data determines LUN type](#) on page 11

[How Data ONTAP can automatically provide more space for full FlexVol volumes](#) on page 41

[How to monitor your LUNs effectively](#) on page 18

Related tasks

[Calculating Rate of Change](#) on page 12

Understanding space usage of LUNs and Snapshot copies

If you have trouble understanding how disk space usage differs between thinly provisioned LUNs versus space-reserved LUNs, you can use the following tasks to see how various commands affect disk space on a volume. These tasks also include how LUNs with Snapshot reserve and without Snapshot reserve affect disk space usage.

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

Related information

[Documentation on the NetApp Support Site: support.netapp.com](http://support.netapp.com)

Showing disk space usage for LUNs without Snapshot copies

You can provision storage and experiment with different commands to see the effect on disk space usage for LUNs that are space reserved or thinly provisioned.

About this task

For this example, none of the LUNs use thin provisioning at first. The volume contains enough free space for all LUNs.

Steps

1. Use the `volume create` to create a new volume named `volSPACE` that has approximately 20 GB.

Example

```
volume create -vserver vs5 -volume volSPACE -aggregate aggr3 -size 20GB
```

2. Use the `df` command to observe the effect of the previous step on disk space.

Example

```
df -vserver vs3 -volume volSPACE -r
```

Filesystem	kbytes	used	avail	reserved	Mounted on
/vol/volSPACE/	19922944	144	19922800	0	---
/vol/volSPACE/.snapshot	1048576	0	1048576	0	---

2 entries were displayed.

3. Use the `volume modify` command to set the percentage of snap reserve space to 0.

Example

```
volume modify -vserver vs5 -volume volSPACE -percent-snapshot-space 0
```

4. Use the `df` command to observe the effect of the previous step on disk space.

Example

```
df -vserver vs3 -volume volSPACE -r
```

Filesystem	kbytes	used	avail	reserved	Mounted on
/vol/volSPACE/	20971520	144	20971376	0	---
/vol/volSPACE/.snapshot	0	0	0	0	---

2 entries were displayed.

The command moved the space from the snapshot reserve back to the available space.

5. Use the `lun create` command to create a LUN named `lun5`.

Example

```
lun create -vserver vs5 -volume volSPACE -lun lun5 -size 5GB -ostype windows
```

6. Use the `df` command to observe the effect of creating `lun5` on disk space.

By default when you create a LUN, the LUN is online and space reservation is enabled.

Example

```
df -vserver vs5 -volume volSPACE -r
```

Filesystem	kbytes	used	avail	reserved	Mounted on
/vol/volospace/	20971520	5266056	15705464	0	---
/vol/volospace/.snapshot	0	0	0	0	---

2 entries were displayed.

The available space decreases and the used space increases by the size of the LUN.

- Use the `lun igroup create` command to create an igroup named `igr5`.

This command has no effect on disk space.

Example

```
lun igroup create -vserver vs5 -igroup igr5 -protocol fcp -ostype windows -initiator 10:00:00:00:c9:73:5b:91
```

- Use the `lun map` command to map the LUN to the igroup.

This command has no effect on disk space.

Example

```
lun map -vserver vs5 -volume volospace -lun lun5 -igroup igr5
```

- Use the `vserver fcp create` to enable the host to listen for FC traffic.

Example

```
vserver fcp create -vserver vs5-status-admin up
```

- Use the `network interface create` command to create a LIF so you can support data access between the initiator and the Storage Virtual Machine (SVM).

Example

```
network interface create -vserver vs5 -lif lif5 -role data -data-protocol fcp -home-node innovate-01 -home-port 0a
```

- From the host, discover the LUN, format it, make the file system available to the host, and write data to the file system.

For information about these procedures, see your Host Utilities documentation. These commands have no effect on disk space.

- Use the `df` command to ensure that creating the file system on the LUN and writing data to it have no effect on space on the SVM.

If you enabled deduplication, compression, or Snapshot copies on the volume, then writing data in the LUN could affect disk space.

Example

```
df -vserver vs5 -volume volospace -r
```

```

Filesystem                kbytes    used    avail    reserved  Mounted on
/vol/volspace/            20971520  5266056  15705464    0    ---
/vol/volspace/.snapshot    0         0         0         0    ---
2 entries were displayed.

```

From the SVM, the amount of space used by the LUN remains 5 GB.

13. Use the `lun modify` command to turn off space reservations.

Example

```
lun modify -vserver vs5 -volume volspace -lun lun5 -space-reserve
disabled
```

14. Use the `df` command to observe the effect of the previous step on disk space.

Example

```
df -vserver vs5 -volume volspace -r
```

```

Filesystem                kbytes    used    avail    reserved  Mounted on
/vol/volspace/            20971520  168    20971352    0    ---
/vol/volspace/.snapshot    0         0         0         0    ---
2 entries were displayed.

```

Note: The 168 kbytes in the code example shown here is the amount of data written to the LUN in step 11.

When you disable space reservation, the LUN is thinly provisioned, and the 5 GB of space for the LUN is no longer reserved. The space is not counted as used space; it is now available space. Any other requests to consume data from the volume, such as a creating a space-reserved LUN or writing to a thinly provisioned LUN, will consume all of the available space, including the 5 GB that the LUN expects to have. If you attempt to write data to the LUN and there is no more space available on the volume, write operations to that LUN will fail and the LUN is taken offline. To restore the reserved space for the LUN, you need to enable space reservation again.

Related concepts

[How file and LUN reservations work for FlexVol volumes](#) on page 57

[What LUN thin provisioning is](#) on page 15

Showing disk space usage for LUNs with Snapshot copies

You can provision storage and experiment with different commands to see the effect of that provisioning on disk space usage for LUNs.

About this task

This task starts by creating a new volume. The example LUN requires 5 GB of disk space, and the fractional overwrite reserve is set to 100 percent.

If you use deduplication on a volume, you would see similar disk space behavior.

Steps

1. Use the `volume create` command to create a new volume named `vol5` that has approximately 20 GB.

```
volume create -vserver vs5 -volume vol5 -aggregate aggr3 -size 20GB
```

2. Use the `df` command volume to observe the effect on disk space.

Example

```
df -vserver vs5 -volume vol5 -r
```

```
Filesystem          kbytes  used   avail  reserved  Mounted on
/vol/vol5/          19922944  144   19922800    0    ---
/vol/vol5/.snapshot 1048576    0   1048576    0    ---
2 entries were displayed.
```

There is a snap reserve of 5 percent on the volume, even though the volume will be used for LUNs, because snap reserve is set to 5 percent by default.

3. Use the `volume modify` command to set the percentage of snap reserve space to 0.

Example

```
volume modify -vserver vs5 -volume vol5 -percent-snapshot-space 0
```

4. Use the `df` command volume to observe the effect on disk space.

Example

```
df -vserver vs5 -volume vol5 -r
```

```
Filesystem          kbytes  used   avail  reserved  Mounted on
/vol/vol5/          20971520  144   20971376    0    ---
/vol/vol5/.snapshot    0      0      0      0    ---
2 entries were displayed.
```

Data ONTAP now counts all the space for SnapReserve as available space for the volume.

5. Use the `lun create` command to create a LUN named `lun5`.

Example

```
lun create -vserver vs5 -volume vol5 -lun lun5 -size 5GB -ostype windows
```

6. Use the `df` command to observe the effect on disk space after creating `lun5`.

Example

```
df -vserver vs5 -volume vol5 -r
```

Filesystem	kbytes	used	avail	reserved	Mounted on
/vol/vol5/	20971520	5266056	15705464	0	---
/vol/vol5/.snapshot	0	0	0	0	---

2 entries were displayed.

About five GB of space is used because this is the amount of space you specified for lun5. By default when you create a LUN, the LUN is online and space reservation is enabled.

7. Use the `lun igroup create` command to create an igroup named `igr5`.

This command has no effect on disk space.

Example

```
lun igroup create -vserver vs5 -igroup igr5 -protocol mixed -ostype windows -initiator 10:00:00:00:c9:73:5b:91
```

8. Use the `lun map` command to map the LUN to the igroup.

Depending on your host, you might need to create WWNN persistent bindings. This command has no effect on disk space.

Example

```
lun map - vs5 -volume vol5 -lun lun5 -igroup igr5
```

9. Use the `vserver fcp create` to enable the host to listen for FC traffic.

Example

```
vserver fcp create -vserver vs5-status-admin up
```

10. Use the `network interface create` command to create a LIF so you can access data between the initiator and the Storage Virtual Machine (SVM).

Example

```
network interface create -vserver vs5 -lif lif5 -role data -data-protocol fcp -home-node innovate-01 -home-port 0a
```

11. From the host, discover the LUN, format it, make the file system available to the host, and write data to the file system.

For information about these procedures, refer to your Host Utilities documentation.

12. Ensure that the active file system is in a quiesced or synchronized state.

13. Use the `lun show` command to display the amount of space consumed by the LUN.

Example

```
lun show -vserver vs5 -volume vol5 -lun lun5 -fields size-used
```

```

vserver path          size-used
-----
vs5      /vol/vol5/lun5 1.51GB

```

14. Use the `volume snapshot create` command to take a Snapshot copy of the active file system.

Example

```
volume snapshot create -vserver vs5 -volume vol5 -snapshot snap_vol5
```

15. Use the `df` command to observe the effect on disk space after taking a Snapshot copy.

Example

```
df -vserver vs5 -volume vol5 -r
```

```

Filesystem          kbytes      used      avail  reserved
Mounted on
/vol/vol5/          20971520   6851372  14120148  1584752 ---
/vol/vol5/.snapshot      0          68          0          0 ---
2 entries were displayed.

```

The reserved amount is the space required to support the next overwrite because the LUN is space reserved.

16. Write approximately 1.5 GB of data to the LUN.
17. Use the `lun show` command to display the amount of space consumed by the lun.

Example

```
lun show -vserver vs5 -volume vol5 -lun lun5 -fields size-used
```

```

vserver path          size-used
-----
vs5      /vol/vol5/lun5 2.97GB

```

18. Use the `df` command to observe the effect on disk space after writing one GB of data to the volume.

Example

```
df -vserver vs5 -volume vol5 -r
```

```

Filesystem          kbytes      used      avail  reserved
Mounted on

```

```

/vol/vol5/                20971520    8387028    12584472    3119520    ---
/vol/vol5/.snapshot      0           540        0           0          ---
2 entries were displayed.

```

After the 1.5 GB write operation to the LUN, the used space increases by about 1.5 GB and the available space decrease by 1.5 GB.

19. Ensure that the active file system is in a quiesced or synchronized state.
20. Use the `volume snapshot create` command to take another Snapshot copy of the active file system.

Example

```
volume snapshot create -vserver vs5 -volume vol5 -snapshot snap2_vol5
```

21. Use the `df` command to observe the effect on disk space after creating lun5.

Example

```
df -vserver vs5 -volume vol5 -r
```

```

Filesystem                kbytes      used      avail    reserved
Mounted on
/vol/vol5/                20971520    8387152    12584368    3119520    ---
/vol/vol5/.snapshot      0           616        0           0          ---
2 entries were displayed.

```

22. Ensure that the active file system is in a quiesced or synchronized state.
23. From the host, overwrite approximately 0.7 GB of data to the LUN.
24. Use the `df` command to observe the effect on disk space after overwriting 0.7 GB of data to lun5.

Example

```
df -vserver vs5 -volume vol5 -r
```

```

Filesystem                kbytes      used      avail    reserved
Mounted on
/vol/vol5/                20971520    10638188    10333332    4653448    ---
/vol/vol5/.snapshot      0           717100     0           0          ---
2 entries were displayed.

```

After the overwrite, the Snapshot copy used space increases by 0.7 GB.

Related concepts

How file and LUN reservations work for Flex Vol volumes on page 57

What LUN thin provisioning is on page 15

SAN configurations always use FlexVol volumes

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.

Volume configuration options

Before you create a FlexVol volume, you must estimate the amount of space you need for LUNs and Snapshot copies. You must also determine the amount of space you want to reserve so that applications can continue to write data to the LUNs in that volume.

The type of LUN you use in your volume, space-reserved or thinly provisioned, determines how you should configure your volume. While thinly provisioned LUNs use less space on the volume, you need to monitor your disk space more closely than you would for space-reserved LUNs.

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

Related concepts

[How Data ONTAP can automatically provide more space for full FlexVol volumes](#) on page 41

[Required changes to Snapshot copy default settings](#) on page 85

[Deciding which LUN type to use in SAN environments](#) on page 11

Considerations for setting fractional reserve for FlexVol volumes

Fractional reserve, also called *LUN overwrite reserve*, enables you to control the size of the overwrite reserve for reserved LUNs and files in a FlexVol volume. By using this volume attribute correctly you can maximize your storage utilization, but you should understand how it interacts with other technologies.

The fractional reserve setting is expressed as a percentage; the only valid values are 0 and 100 percent. You use the `volume modify` command to set fractional reserve.

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`, when any of the following technologies and Data ONTAP features are in use:

- Deduplication
- Compression
- FlexClone files
- FlexClone LUNs
- Virtual environments

If you are using one or more of these technologies with no fractional reserve and you need to prevent errors due to running out of space, you must use all of the following configuration settings for the volume:

- Volume guarantee of `volume`
- File or LUN reservations enabled
- Volume Snapshot copy automatic deletion enabled with a commitment level of `destroy` and a destroy list of `lun_clone, vol_clone, cifs_share, file_clone, sfsr`

Note: 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 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 do not want to monitor aggregate free space or have volume Snapshot copies automatically deleted, you can 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
File	100	100

What Snapshot autodelete is

Snapshot autodelete is a volume-level option that allows you to define a policy for automatically deleting Snapshot copies based on a definable threshold.

You must choose one of these triggers to automatically delete Snapshot copies:

- The volume is nearly full
- The snap reserve space is nearly full
- The overwrite reserved space is full

Using Snapshot autodelete is recommended in most SAN configurations.

For more information about using Snapshot autodelete to automatically delete Snapshot copies, see the *Data ONTAP Data Protection Online Backup and Recovery Guide*.

Related tasks

[Determining the volume size when using Snapshot autodelete](#) on page 21

What Snapshot reserve is

Snapshot reserve is the amount of space reserved for Snapshot copies.

When this space is full, Snapshot copies begin taking space from the available space on a volume. When you use thin provisioning, you should set the Snapshot reserve to zero. You set this variable at the volume level.

If you use Snapshot reserve, you are removing free space from the volume before it is needed, and you cannot use this reserved space for anything other than Snapshot copies. Thin provisioning allocates space only when it is needed.

How Data ONTAP can automatically provide more space for full FlexVol volumes

Data ONTAP uses two methods for automatically providing more space for a FlexVol volume when that volume is nearly full: allowing the volume size to increase, and deleting Snapshot copies (with any associated storage objects). If you enable both of these methods, you can specify which method Data ONTAP should try first.

Data ONTAP can automatically provide more free space for the volume by using one of the following methods:

- Increase the size of the volume when it is nearly full (known as the *autogrow* feature).
This method is useful if the volume's containing aggregate has enough space to support a larger volume. You can configure Data ONTAP to increase the size in increments and set a maximum size for the volume. The increase is automatically triggered based on the amount of data being written to the volume in relation to the current amount of used space and any thresholds set.
- Delete Snapshot copies when the volume is nearly full.
For example, you can configure Data ONTAP to automatically delete Snapshot copies that are not linked to Snapshot copies in cloned volumes or LUNs, or you can define which Snapshot copies you want Data ONTAP to delete first—your oldest or newest Snapshot copies. You can also determine when Data ONTAP should begin deleting Snapshot copies—for example, when the volume is nearly full or when the volume's Snapshot reserve is nearly full.

If you enable both of these methods, you can specify which method Data ONTAP tries first when a volume is nearly full. If the first method does not provide sufficient additional space to the volume, Data ONTAP tries the other method next. By default, Data ONTAP tries to increase the size of the volume first.

Selecting the first method to increase space for full FlexVol volumes

If you enable both the volume autosize capability and automatic Snapshot deletion for the same volume, you can also specify which method Data ONTAP tries first when that volume needs additional free space. How you configure the volume depends on whether you would prefer that the volume continue to grow or that Snapshot copies are deleted.

About this task

If the first method that Data ONTAP tries to use for providing additional free space does not result in addressing the free space needs of the volume, then Data ONTAP tries the other method.

In most cases, the default configuration (growing the volume first) is preferable, because when a Snapshot copy is deleted, it cannot be restored. However, in certain situations, you might want to avoid growing the size of a volume when possible. In this case, you can configure Data ONTAP to delete Snapshot copies before increasing the size of the volume.

Step

1. Select the first method that Data ONTAP should use to provide free space to a volume by using the `volume modify` command with the `-space-mgmt-try-first` option.

To specify increasing the size of the volume first (the default), use `volume_grow`. To specify deleting Snapshot copies first, use `snap_delete`.

Example

The following command configures Data ONTAP to delete Snapshot copies before increasing the volume size for the volume `vol0001`:

```
volume modify -vserver vs2 -volume vol0001 -space-mgmt-try-first  
snap_delete
```

Configuring volumes in a SAN environment

After you decide how you want to reserve space for the LUNs and Snapshot copies in your volumes, you can begin configuring your volumes for your SAN environment. You should configure your volumes before you set up your LUNs.

Depending on the requirements of your system, you might need to modify some of the configurations in these tasks. If you have any questions about the impact of these volume configurations on your environment, contact technical support.

For more information about volumes, see the *Clustered Data ONTAP Logical Storage Management Guide*.

Related concepts

[Guidelines for provisioning storage in a SAN environment](#) on page 56

Configuring volumes for thinly provisioned LUNs

When you configure your volume for thinly provisioned LUNs, you get excellent storage utilization because you can add space to your volume, LUN, and Snapshot copies as needed. Volume configuration of this kind enables you to manage your volumes and LUNs more effectively by growing your volume or deleting Snapshot copies if there is not enough free space.

Before you begin

You have created a volume.

About this task

You should use host-based Snapshot copy creation software such as SnapDrive to create your Snapshot copies.

Steps

1. Use the `volume autosize` command to turn on volume autosize.

Example

```
volume autosize -vserver vs3 -volume voll -is-enabled on
```

```
vol autosize: Flexible volume "vs3:voll" autosize settings UPDATED.
```

2. Use the `volume autosize` command to specify the maximum volume size and the increment size.

Example

```
volume autosize -vserver vs3 -volume voll -maximum-size 40GB -increment-size 5GB
```

Note: The actual values for `-maximum-size` and `-increment-size` should be set based on your environment.

```
vol autosize: Flexible volume "vs3:voll" autosize settings UPDATED.
```

3. Use the `volume modify` command to set the space management first try option to `volume_grow` (`autosize`).

Example

```
volume modify -vserver vs3 -volume voll -space-mgmt-try-first volume_grow
```

```
Volume modify successful on volume: voll
```

4. Use the `volume modify` command to change the snap reserve setting to 0.

Example

```
volume modify -vserver vs3 -volume voll -percent-snapshot-space 0
```

```
Volume modify successful on volume: voll
```

5. Use the `volume modify` command to disable the scheduled creation of Snapshot copies.

Example

```
volume modify -snapshot-policy none -vserver vs3 -volume voll
```

```
Warning: You are changing the Snapshot policy on volume voll to none. Any Snapshot copies on this volume from the previous policy will not be deleted by this new Snapshot policy. Do you want to continue? {y|n}: y
Volume modify successful on volume: voll
```

6. Use the `volume show` command to verify the changes you made for volume `autosize` and `snap reserve`.

Example

```
volume show -vserver vs3 -volume voll -instance
```

```

Vserver Name: vs3
Volume Name: voll
Aggregate Name: aggr1
Volume Size: 100MB
Available Size: 98.88MB
Total User-Visible Size: 100MB
Used Size: 1.12MB
Used Percentage: 1%
Autosize Enabled (for flexvols only): true           Maximum
Autosize (for flexvols only): 40MB                 Autosize Increment (for
flexvols only): 5MB
Total Files (for user-visible data): 3033
Files Used (for user-visible data): 111
Space Guarantee Style: volume
Space Guarantee in Effect: true
Space Reserved for Snapshots: 0%
Snapshot Reserve Used: 0%
Snapshot Policy: none
Clone Volume: false
NVFAIL Option: on
Is File System Size Fixed: false
Extent Option: off
Reserved Space for Overwrites: 0B
Fractional Reserve: 100%
Automatic Snapshots: true
Snapshot Cloning Dependency: off
Primary Space Management Strategy: volume_grow
Read Reallocation Option: off
Naming Scheme for Automatic Snapshots: create_time
Inconsistency in the File System: false

```

Note: Some rows in the example output have been deleted to improve readability.

7. Use the `volume snapshot autodelete modify` command to enable Snapshot autodelete.

Example

```
volume snapshot autodelete modify -vserver vs3 -volume voll -enabled
true
```

```
Volume modify successful on volume: voll
```

Note: You should review the Snapshot autodelete policy settings and configure these settings based on your business needs. For more information, see the man page for `volume snapshot autodelete modify` command or the *Clustered Data ONTAP Logical Storage Management Guide*.

8. Use the `volume snapshot autodelete show` command to verify the change you made for Snapshot autodelete.

Example

```
volume snapshot autodelete show -vserver vs3 -volume vol1
```

```
Vserver Name: vs3
Volume Name: vol1
  Enabled: true
  Commitment: try
  Defer Delete: user_created
  Delete Order: oldest_first
Defer Delete Prefix: (not specified)
  Target Free Space: 20%
    Trigger: volume
  Destroy List: none
```

Result

Volume vs3 is configured for thinly provisioned LUNs without Snapshot reserve. You can now create your thinly provisioned LUNs for your volume.

Related concepts

[Guidelines for provisioning storage in a SAN environment](#) on page 56

[What LUN thin provisioning is](#) on page 15

[Volume option best practices for thinly provisioned LUNs](#) on page 17

[When to use thinly provisioned LUNs](#) on page 16

Related tasks

[Determining the volume size when using Snapshot autodelete](#) on page 21

Configuring volumes for space-reserved LUNs without Snapshot reserve

When you configure a space-reserved LUN, this space is reserved and not available to other LUNs or Snapshot copies within the volume. However, when Snapshot reserve is enabled, Snapshot copies are limited by the amount of available free space on the volume.

Before you begin

You have created a volume.

About this task

You should use host-based Snapshot copy creation software such as SnapDrive to create your Snapshot copies. The following configuration steps apply at the volume level:

Steps

1. Use the `volume modify` command to set space guarantee to volume.

Example

```
volume modify -vserver vs3 -volume vol1 -space-guarantee volume
```

```
Volume modify successful on volume: vol1
```

2. Use the `volume modify` command to set fractional reserve to 0.

Example

```
volume modify -vserver vs3 -volume vol1 -fractional-reserve 0
```

```
Volume modify successful on volume: vol1
```

3. Use the `volume autosize` command to enable volume autosize.

Example

```
volume autosize -vserver vs3 -volume vol1 -is-enabled on
```

```
vol autosize: Flexible volume 'vs3:vol1' autosize settings UPDATED.
```

4. Use the `volume autosize` command to specify the maximum volume size and the increment size.

Example

```
volume autosize -vserver vs3 -volume vol1 -maximum-size 40GB -increment-size 5GB
```

Note: The actual values for `-maximum-size` and `-increment-size` should be set based on your environment.

```
vol autosize: Flexible volume 'vs3:vol1' autosize settings UPDATED.
```

5. Use the `volume modify` command to set the `-space-mgmt-try-first` option to volume grow (autosize).

Example

```
volume modify -vserver vs3 -volume vol1 -space-mgmt-try-first volume_grow
```

```
Volume modify successful on volume: vol1
```

- Use the `volume modify` command to set Snapshot reserve to 0.

Example

```
volume modify -vserver vs3 -volume voll -percent-snapshot-space 0
```

```
Volume modify successful on volume: voll
```

- Use the `volume modify` command to disable the scheduled creation of Snapshot copies.

Example

```
volume modify -snapshot-policy none -vserver vs3 -volume voll
```

```
Volume modify successful on volume: voll
```

- Use the `volume show` command to verify your settings.

Example

```
volume show -vserver vs3 -volume voll -instance
```

```

Vserver Name: vs3
Volume Name: voll
Aggregate Name: aggr3
Volume Size: 20MB
Available Size: 19.90MB
Total User-Visible Size: 20MB
Used Size: 104KB
Used Percentage: 0%
Autosize Enabled (flexvols only): true

```

```

Maximum Autosize (flexvols only): 40GB Autosize Increment (flexvols
only): 5GB

```

```

Total Files (for user-visible data): 566
Files Used (for user-visible data): 96
Space Guarantee Style: volume
Space Guarantee in Effect: true
Space Reserved for Snapshots: 0%

```

```

none Snapshot Reserve Used: 0% Snapshot Policy:
Creation Time: Wed May 18 20:23:11 2011
Fractional Reserve: 0%
Automatic Snapshots: false
Snapshot Cloning Dependency: off
Primary Space Management Strategy: volume_grow

```

Note: Some rows in the example output have been deleted for improved readability.

- Use the `volume snapshot autodelete modify` command to enable Snapshot autodelete.

Example

```
volume snapshot autodelete modify -vserver vs3 -volume voll -enabled true
```

```
Volume modify successful on volume: voll
```

Note: You should review the Snapshot autodelete policy settings and configure these settings based on your business needs. For more information, see the man page for `volume snapshot autodelete modify` command or the *Clustered Data ONTAP Logical Storage Management Guide*.

- Use the `volume snapshot autodelete modify` command to set the autodelete trigger.

Example

```
volume snapshot autodelete modify -vserver vs3 -volume voll -trigger volume
```

```
Volume modify successful on volume: voll
```

- Use the `volume snapshot autodelete modify` command to set the delete order to delete oldest Snapshot copy first.

Example

```
volume snapshot autodelete modify -vserver vs3 -volume voll -delete-order oldest_first
```

```
Volume modify successful on volume: voll
```

- Use the `volume snapshot autodelete show` to verify your settings.

Example

```
volume snapshot autodelete show -vserver vs3 -volume voll -instance
```

```
Vserver Name: vs3
Volume Name: voll
  Enabled: true
  Commitment: try
  Defer Delete: user_created
  Delete Order: oldest_first
Defer Delete Prefix: (not specified)
Target Free Space: 20%
```

```

Trigger: volume
Destroy List: none

```

Result

The volume is now configured for space-reserved LUNs without Snapshot reserve. You can now create LUNs for your volume.

Related concepts

[Guidelines for provisioning storage in a SAN environment](#) on page 56

[How file and LUN reservations work for FlexVol volumes](#) on page 57

[When to use space-reserved LUNs without Snapshot reserve](#) on page 27

Related tasks

[Determining the volume size when using Snapshot autodelete](#) on page 21

Configuring volumes for space-reserved LUNs with Snapshot reserve

When you reserve space for LUNs and Snapshot copies, you guarantee that the space is used only for those LUNs and Snapshot copies. This reserved space for the LUNs and the Snapshot copies is not available to any other LUN or Snapshot copy within that same volume.

Before you begin

You have created a volume.

About this task

You should use host-based Snapshot copy creation software such as SnapDrive to create your Snapshot copies. The following configurations apply at the volume level:

Steps

1. Use the `volume modify` command to set space guarantee to `volume`.

Example

```
volume modify -vserver vs3 -volume vol1 -space-guarantee volume
```

```
Volume modify successful on volume: vol1
```

2. Use the `volume modify` command to set `fractional reserve` to 100.

Example

```
volume modify -vserver vs3 -volume voll -fractional-reserve 100
```

```
Volume modify successful on volume: voll
```

3. Use the `volume autosize` command to disable volume autosize.

Example

```
volume autosize -vserver vs3 -volume voll -is-enabled off
```

```
vol autosize: Flexible volume 'vs3:voll' autosize settings UPDATED.
```

Note: If you need to protect Snapshot copy creation, set volume autosize to on. This setting ensures that your volume has space to grow.

4. Use the `volume modify` command to change the snap reserve setting to 0.

Example

```
volume modify -vserver vs3 -volume voll -percent-snapshot-space 0
```

```
Volume modify successful on volume: voll
```

5. Use the `volume modify` command to disable the scheduled creation of Snapshot copies.

Example

```
volume modify -snapshot-policy none -vserver vs3 -volume voll
```

```
Volume modify successful on volume: voll
```

6. Use the `volume show` command to verify your changes.

Example

```
volume show -vserver vs3 -volume voll -instance
```

```

Vserver Name: vs3
Volume Name: voll
Aggregate Name: aggr1
Volume Size: 100MB
Available Size: 98.88MB
Total User-Visible Size: 100MB
Used Size: 1.12MB
Used Percentage: 1%
Autosize Enabled (for flexvols only): false
Maximum Autosize (for flexvols only): 120MB
Autosize Increment (for flexvols only): 5MB

```

```

Total Files (for user-visible data): 3033
Files Used (for user-visible data): 111
Space Guarantee Style: volume
Space Guarantee in Effect: true
Space Reserved for Snapshots: 0%
Snapshot Reserve Used: 0%
Snapshot Policy: none
Clone Volume: false
NVFAIL Option: on
Is File System Size Fixed: false
Extent Option: off
Reserved Space for Overwrites: 0B
Fractional Reserve: 100%
Automatic Snapshots: true
Snapshot Cloning Dependency: off
Primary Space Management Strategy: volume_grow
Read Reallocation Option: off

```

Note: Some rows in the example output have been deleted to improve readability.

7. Use the `volume snapshot autodelete modify` command to disable autodelete.

Example

```

volume snapshot autodelete modify -vserver vs3 -volume voll -enabled
false

```

```

Volume modify successful on volume: voll

```

Note: You should review the Snapshot autodelete policy settings and configure these settings based on your business needs. For more information, see the man page for `volume snapshot autodelete modify` command or the *Clustered Data ONTAP Logical Storage Management Guide*.

8. Use the `volume snapshot autodelete show` command to verify your change to autodelete.

Example

```

volume snapshot autodelete show -vserver vs3 -volume voll -instance

```

```

Vserver Name: vs3
Volume Name: voll
Enabled: false
Commitment: try
Defer Delete: user_created
Delete Order: oldest_first
Defer Delete Prefix: (not specified)
Target Free Space: 20%
Trigger: volume
Destroy List: none

```

Result

Your volume is configured for space-reserved LUNs with Snapshot reserve. You can now create LUNs for your volume.

Related concepts

Guidelines for provisioning storage in a SAN environment on page 56

How file and LUN reservations work for FlexVol volumes on page 57

When to use space-reserved LUNs with Snapshot reserve on page 28

Related tasks

Determining the volume size when using Snapshot autodelete on page 21

LUN setup

After you have configured your volume for a specific LUN type, LUN setup provides LUN information and workflow tasks based on iSCSI, FC, or FCoE protocols. You also can use the Storage Virtual Machine (SVM) setup wizard to create your LUNs. After you complete the LUN workflow tasks, your LUNs are ready to serve data.

Data ONTAP 8.1 and later releases support the use of nodes in cluster SAN environments by enabling LUN access to nodes through LIFs and SVMs, provided you have configured LUNs for a cluster environment.

If you are setting up LUNs for the first time, you need to answer the following questions before setting up your LUNs:

- How many SVMs do you need?
- How many LIFs per SVM for each node?
- What are your LUN sizes?
- Are you using space-reserved or thinly provisioned LUNs?
- Do you want to limit LUN access by initiators through port sets?
- Which kind of storage provisioning do you want to use?
- Are you using FCoE?

Note: SAN cluster functionality is only available when all nodes in a cluster use Data ONTAP 8.1 or later.

Related concepts

What LUN thin provisioning is on page 15

Volume option best practices for thinly provisioned LUNs on page 17

Related tasks

Setting up LUNs using the SVM setup wizard on page 79

Setting up LUNs manually with iSCSI block access on page 64

Setting up LUNs manually with FC block access on page 71

LUN guidelines for SAN environments

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

Before you create a LUN, you must make sure that the host is supported. For information about supported hosts, see the Interoperability Matrix.

Related information

[Interoperability Matrix: support.netapp.com/matrix](http://support.netapp.com/matrix)

Recommended and supported SAN configurations

Storage systems and hosts can be connected through an FC fabric or a TCP/IP network.

You should create multiple data LIFs to LUNs to ensure access to your LUNs in case of failure. For information about recommended and supported SAN configurations, see the *Clustered Data ONTAP SAN Configuration Guide*.

FC fabrics and FCoE networks must be zoned by the WWPN of the LIF, rather than the WWPN of the port. For more information about zoning, see the *Clustered Data ONTAP SAN Configuration Guide*.

Related information

[Interoperability Matrix: support.netapp.com/matrix](http://support.netapp.com/matrix)

[Data ONTAP documentation on the NetApp Support Site-support.netapp.com](http://support.netapp.com)

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/qtreen1/file1` and the LUN path must be `/vol/vol1/qtreen1/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/qtreen1/subdir1/file1` and the LUN path must be `/vol/vol1/qtreen1/lun1`.

Considerations about LUN identification numbers

You can assign a number for the LUN ID, or you can accept the default LUN ID. However, your Host Utilities could have additional guidelines on how to assign LUN identification numbers.

Typically, the default LUN ID begins with 0 and increments by 1 for each additional LUN as it is mapped. 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.

Note: For detailed information, see the documentation provided with your Host Utilities.

Related concepts

What Host Utilities are on page 177

Related information

Host Utilities documentation: www.support.netapp.com

Guidelines for provisioning storage in a SAN environment

When provisioning storage in a SAN environment, there are several best practices you should follow to ensure that your systems run smoothly.

You should follow these guidelines when creating FlexVol volumes that contain LUNs, regardless of which provisioning method you choose:

- You cannot create any LUNs in the system's root volume. You should not create LUNs in the Storage Virtual Machine (SVM) root volume.
Data ONTAP uses this volume to administer the SVM.
- You must ensure that no other files or directories exist in a volume that contains LUNs.
If this is not possible and you are storing LUNs and files in the same volume, you can use a separate qtree to contain the LUNs.
- If multiple hosts share the same volume, you can create a qtree on the volume to store all LUNs for the same host.
This is a recommended best practice that simplifies LUN administration and tracking.
- You can make the required changes to the snap reserve default settings.
You can change the `-percent-snapshot-space` setting for the volume to 0, run the `volume snapshot policy modify` command to disable the Snapshot schedule so that no controller-based Snapshot copies are taken, and delete all existing Snapshot copies.
- To simplify management, you should use naming conventions for LUNs and volumes that reflect their ownership or the way that they are used.

For more information about creating volumes, see the *Clustered Data ONTAP Logical Storage Management Guide* and *Clustered Data ONTAP Physical Storage Management Guide*.

Related information

Data ONTAP documentation on the NetApp Support Site-support.netapp.com

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 map a LUN to multiple igroups that contain the same initiator.
- 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.

Guidelines for LUN layout and space allocation

When you create LUNs, you should follow certain guidelines for LUN layout and space allocation.

- Group LUNs according to their rates of change.
If you plan to take Snapshot copies, do not create LUNs with a high rate of change in the same volumes as LUNs with a low rate of change. When you calculate the size of your volume, the data for rate of change enables you to determine the amount of space you need for Snapshot copies. If you calculate your volume size based on a low rate of change, and then create LUNs with a high rate of change in that volume, you might not have enough space for Snapshot copies.
- Keep backup LUNs in separate volumes.
Keep backup LUNs in separate volumes if the primary data has a high rate of change. For example, you might copy all the data in a LUN to a backup LUN and then move the backup LUN to tape each day. All of the data in the backup LUN changes daily. If you want to keep backup LUNs in the same volume, you must calculate the size of the volume based on a high rate of change in your data.

Understanding LUN functions

To set up your LUNs successfully, you need to understand how LUN configurations affect your environment.

You need to understand the following:

- How the host operating system affects usable space for a LUN
- How you make a LUN available on specific target ports
- How to avoid igroup mapping conflicts
- How you map read-only LUNs to hosts at SnapMirror destinations

Before you create a LUN, you must make sure that the host is supported. For information about supported hosts, see the Interoperability Matrix.

You also can review the man page for additional information about a specific LUN command.

Related information

[Interoperability Matrix: support.netapp.com/matrix](http://support.netapp.com/matrix)

How file and LUN reservations work for FlexVol volumes

When reservations are enabled for one or more LUNs, Data ONTAP reserves enough space in the volume so that writes to those LUNs do not fail because of a lack of disk space.

Reservations are an attribute of the LUN; they are persistent across storage system reboots, takeovers, and givebacks. Reservations are enabled for new LUNs by default, but you can create a

LUN with reservations disabled or enabled. After you create a LUN, you change the reservation attribute by using the `lun modify` command.

When a volume contains one or more LUNs with reservations enabled, operations that require free space, such as the creation of Snapshot copies, are prevented from using the reserved space. If these operations do not have sufficient unreserved free space, they fail. However, writes to the LUNs with reservations enabled continue to succeed.

You can enable reservations for LUNs contained by volumes with volume guarantees of any value. However, if the volume has a guarantee of `none`, reservations do not provide protection against out-of-space errors.

Example

If you create a 100-GB space-reserved LUN in a 500-GB volume, that 100 GB of space is immediately allocated, leaving 400 GB remaining in the volume. In contrast, if space reservation is disabled on the LUN, all 500 GB in the volume remain available until writes are made to the LUN.

How the host operating system affects usable space for a LUN

The disk geometry used by the operating system determines the minimum and maximum size values of LUNs.

The usable space in the LUN depends on host or application requirements for overhead. Partition tables and metadata on the host file system reduce the usable space for applications. In general, when you format and partition LUNs as a disk on a host, the actual usable space on the disk depends on the overhead required by the host.

For information about the maximum sizes for LUNs and disk geometry, see the vendor documentation for your host OS. If you are using third-party volume management software on your host, consult the vendor's documentation for more information about how disk geometry affects LUN size.

LUN availability on specific target LIFs

When you map a LUN to an igroup, the LUN is available on all of the storage system's FC or iSCSI target LIFs if the igroup is not bound to a port set.

By binding an igroup to a port set, you can make the LUN available on a subset of the system's target LIFs. Any host in the igroup can access the LUNs only by connecting to the target LIFs in the port set.

You can define port sets for FC and iSCSI target LIFs or have mixed port sets.

Related concepts

[What igroups are](#) on page 123

[How to limit LUN access in a virtualized environment](#) on page 63

How Data ONTAP avoids igroup mapping conflicts

When you add an initiator to an igroup, Data ONTAP verifies that there are no igroup conflicts within the cluster.

If you try to add an initiator to an igroup where the mapped LUN ID is already in use, Data ONTAP returns an error message.

For example, initiator1 is in igroup1 and is connected to "lun1" with a mapped LUN ID of zero. Initiator2 is in igroup2 and is connected to "lun2" with a mapped LUN ID of zero.

If you try to add initiator1 to igroup2, Data ONTAP returns an error message. To avoid this error, you can use the `lun map` command and change the `lun-id` option to manually specify the mapped LUN ID.

SnapMirror destinations and read-only LUNs

When you use SnapMirror to replicate a source volume to a destination volume, you replicate all LUNs in that source volume to the destination volume. You cannot replicate an individual LUN in a source volume to the destination volume.

You can use SnapMirror to replicate a destination volume on a different Storage Virtual Machine (SVM) from the source volume. However you must first set up SVM peering relationship between the SVMs before you can establish a SnapMirror relationship. For example, if your source volume is on SVM1, the destination volume can be on SVM2 or the destination volume can be on the same SVM as the source volume. For more information about SVM peering relationships, see the *Clustered Data ONTAP System Administration Guide for Cluster Administrators*.

After the first SnapMirror is initialized or updated, the LUNs in the destination volume are unmapped, online, and are assigned a new serial number. The LUNs are read-only on the destination volume. You need to use the `lun map` command to map the LUNs in the destination volume to the appropriate hosts to enable read-only access to the LUNs.

You can create new igroups on the destination, map the destination LUN to those igroups, or use any existing igroups.

As long as the SnapMirror relationship exists between the source and destination volumes, you cannot create or delete LUNs on the destination volume. You can do the following to the destination LUNs:

- Map a LUN
- Unmap a LUN
- Bring a LUN online
- Take a LUN offline
- Make changes to SCSI-2 reservations and SCSI-3 persistent reservations

After the mirror relationship is broken, the LUN transparently migrates to a read/write state. Hosts might need to remount the device to notice the change.

Attention: Attempts to write to read-only LUNs fail, and might cause applications and hosts to fail as well. Before mapping read-only LUNs to hosts, you must ensure that the operating system and application support read-only LUNs.

For more information about read-only LUNs and SnapMirror, see the *Clustered Data ONTAP Data Protection Guide*.

Introduction to LUNs in a virtualized SAN environment

Storage Virtual Machines (SVMs) and LIFs enable you to scale your SAN environment in a cluster. For a single SVM, you can create many LIFs on a port for a node. How you set up igroups and port sets determines which initiators have access to which LUNs.

The advantages of SVMs and LIFs in a cluster are the following:

- **Distributed management:** You can login to any node in the SVM to administer all 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 that SVM.
- **Controlled LUN access:** If you use port sets, you can limit which LIFs an initiator can use to access LUNs.

Related concepts

[What ALUA is](#) on page 125

[MPIO and ALUA](#) on page 126

[How to limit LUN access in a virtualized environment](#) on page 63

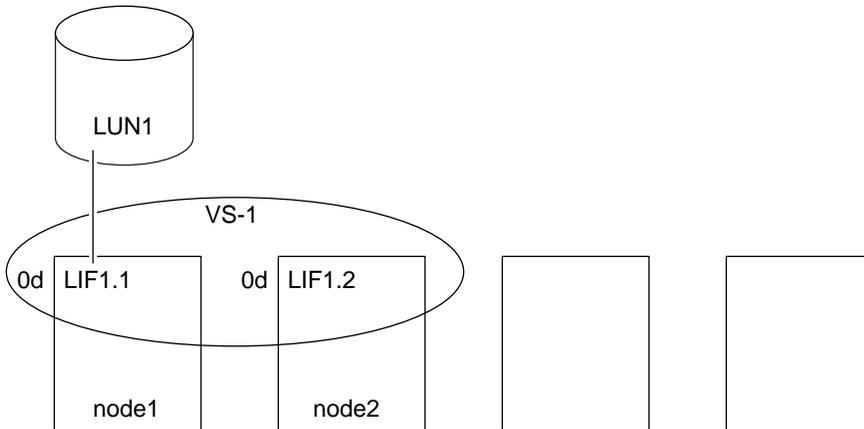
How LUN access works in a virtualized environment

LIFs enable Storage Virtual Machines (SVMs) to access LUNs through optimized and unoptimized paths.

A LIF is a logical interface that connects an 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.

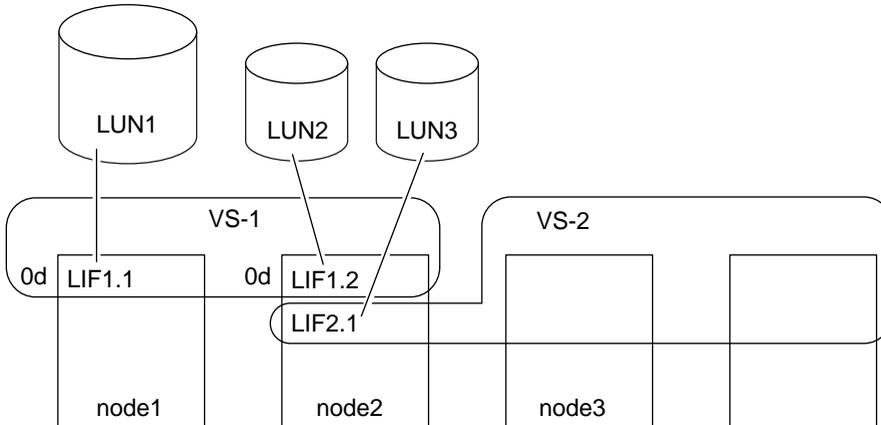
Single SVM with a single LIF

In the following example, LIF1.1 connects SVM VS-1 to node1:0d. LIF1.1 belongs only to VS-1.



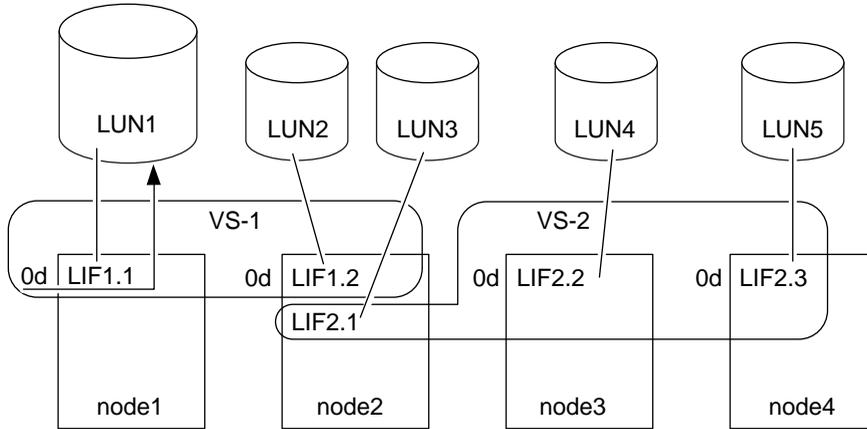
Two SVMs with multiple LIFs

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 VS-1 has LIF1.2 and SVM VS-2 has LIF2.1, both of which are on node2:0d. SVM VS-1 connects LIF1.2 to LUN2 and SVM VS-2 connects LIF2.1 to LUN3.



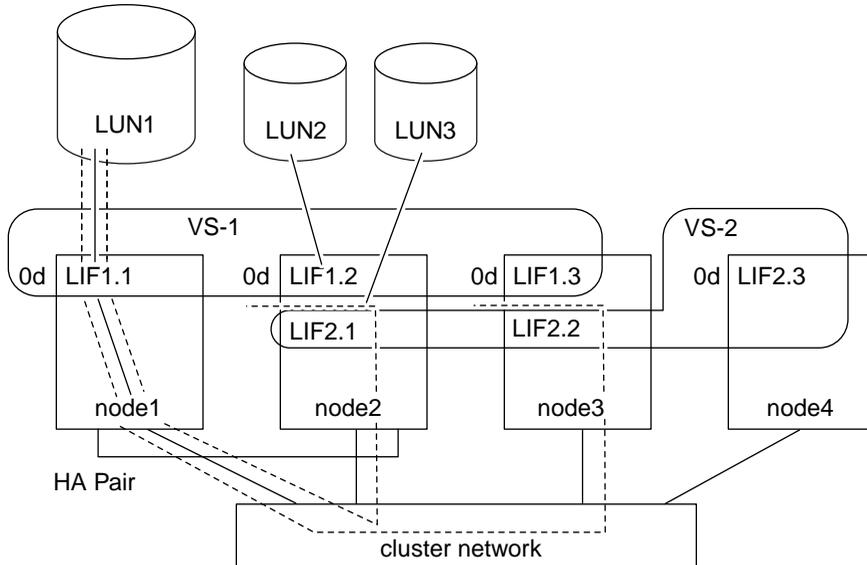
An optimized path to a LUN

In an optimized path, the data traffic does not travel over the cluster network and 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.



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 in order 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.

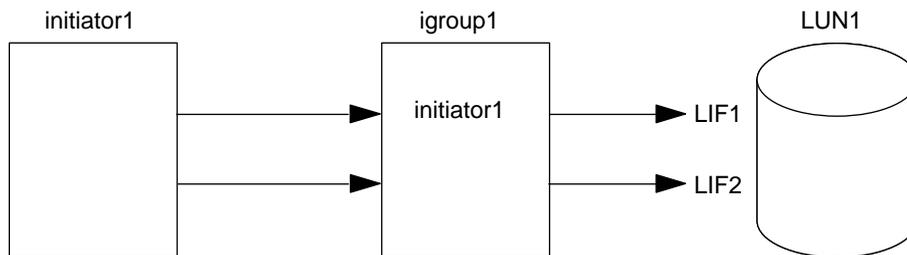


How to limit LUN access in a virtualized environment

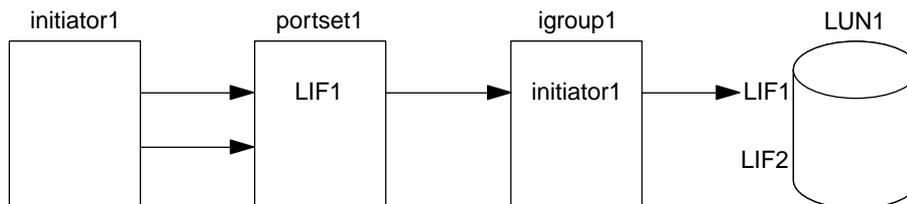
You can limit access to your LUNs through igroups and port sets. Until you map LUNs to an igroup, you cannot access your LUNs. You use port sets to limit which LIFs an initiator can use to access your LUNs.

A port set is a collection of LIFs. If you do not bind those igroups to a port set, any initiator within an igroup can access any LUN mapped to that igroup through any LIF.

The way an initiator connects to a LUN is through an igroup. In the following example, initiator1 can access LUN1 through LIF1 and LIF2. Without a port set, initiator1 can access LUN1 over any LIF.



You can limit access to a LUN by binding igroups to port sets. In the following example, initiator1 can access LUN1 through LIF1. However, initiator1 cannot access LUN1 through LIF2 because LIF2 is not in portset1.



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.
- Some options are not applicable for iSCSI or FC. For example, you cannot use IP addresses with FC.

Setting up LUNs

You have the option of setting up LUNs by using the Storage Virtual Machine (SVM) wizard or manually setting up your LUNs.

If you want to set up your LUNs manually, you must do the following:

- Set up block access with either iSCSI or FC protocol.
- Create LIFs that connect an SVM to a port.
- Create your LUNs and igroups.
- Map the LUN to an igroup.
- Optional: Create port sets and bind the port sets to igroups.

Related concepts

[Introduction to LUNs in a virtualized SAN environment](#) on page 60

Setting up LUNs manually with iSCSI block access

When you set up your LUNs manually, you must complete several required tasks before you can access your LUNs. After you complete all the required tasks, you can access your LUNs.

Before you begin

Before you set up your LUNs manually, you should first consider the following:

- Your SAN configuration must be supported.
For more information about supported configurations, see the Interoperability Matrix Tool.
- Your SAN environment must meet the SAN host and controller configuration limits.
For more information, see 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.

Steps

1. [Verifying the license for FC or iSCSI](#) on page 65
Before you can enable block access for a Storage Virtual Machine (SVM) with FC or iSCSI, you must have a license.
2. [Enabling block access for an SVM with iSCSI](#) on page 65
You must assign an iSCSI protocol to your Storage Virtual Machine (SVM) and create your LIFs for that SVM.
3. [Creating LUNs and mapping igroups for iSCSI](#) on page 68
As part of configuring your SAN environment, you have to create LUNs and map them to initiator groups.
4. [Creating port sets and binding igroups to port sets](#) on page 69
You have the option of using port sets to limit which LIFs can be used by an initiator to access a LUN.
5. [Enabling block access for a specific host with iSCSI](#) on page 70
Although the actual steps for enabling block access for your host could vary depending on your host, use this procedure as a guide for what you must do to connect your host to the Storage Virtual Machine (SVM).

Related information

Interoperability Matrix: support.netapp.com/matrix

Host Utilities documentation: www.support.netapp.com

Documentation on the NetApp Support Site: support.netapp.com

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

1. 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	-

2. 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 an SVM with iSCSI

You must assign an iSCSI protocol to your Storage Virtual Machine (SVM) and create your LIFs for that SVM.

Before you begin

You need an iSCSI license.

Note: If you do not enable the iSCSI service, the LIFs and SVMs appear to be online but the operational status is down. You must enable the iSCSI service in order for your LIFs and SVMs to be operational.

About this task

You need a minimum of two LIFs per node. A host needs at least two paths to the LUN through each node.

At a minimum, you need paths to both nodes of the HA pair on which you create the LUN. You can have paths to a maximum of four nodes. If you have more than four nodes in your cluster, you can use port sets to limit the number of paths.

Note: Starting with Data ONTAP 8.2.0, LIFs support IPv4 and IPv6 addressing.

Steps

1. Use the `vserver iscsi create` command to enable the SVMs to listen for iSCSI traffic.

Example

```
vserver iscsi create -vserver vs2 -target-alias vs2
```

The operational status of iSCSI service on the specified SVM is up and ready to serve data.

2. Use the `network interface create` command to create a LIF for a specific SVM.

Example

```
network interface create -vserver vs2 -lif lif1 -role data -data-protocol iscsi -home-node node1 -home-port e0b -address 192.168.2.72 -netmask 255.255.255.0
```

Note: To create an iSCSI LIF, you must specify the `-role` parameter as `data` and the `-protocol` parameter as `iscsi`.

Note: You should create at least one LIF for each node in all SVMs serving data with the iSCSI protocol.

3. If you route your iSCSI traffic, use the `network routing-groups route create` command to create the default route for each SVM.

Example

```
network routing-groups route create -vserver vs2 -routing-group d192.168.1.0/24 -destination 0.0.0.0/0 -gateway 192.168.1.1
```

4. Use the `network routing-groups route show` command to verify the default route for each SVM.

Example

```
network routing-groups route show -vserver vs2
```

Vserver	Routing Group	Destination	Gateway	Metric
---------	---------------	-------------	---------	--------

```
-----
vs2          d192.168.1.0/24  0.0.0.0/0      192.168.1.1    20
-----
```

Note: Some of the columns in the example have been truncated to improve readability.

- Use the `network interface show` command to verify that you set up your LIFs correctly.

Example

```
network interface show -vserver vs2
```

```
-----
Logical      Status      Network      Current      Current      Is
Vserver      Interface   Admin/Oper   Address/Mask  Node         Port         Home
-----
vs2
    lif0      up/up       192.168.1.72/24  node1        e0b         true
    lif1      up/up       192.168.2.72/24  node1        e4a         true
    lif2      up/up       192.168.1.71/24  node2        e0b         true
    lif3      up/up       192.168.2.71/24  node2        e4a         true
4 entries were displayed.
```

- From your host (initiator), create iSCSI sessions to your LIFs.

After you finish

- You must create and map your LUNs.
- Optional: You can set up iSCSI security with CHAP.

Related concepts

[Considerations for LIFs in cluster SAN environments](#) on page 63

[What CHAP authentication is](#) on page 136

Related tasks

[Defining a security policy method for an initiator](#) on page 129

Related information

[Documentation on the NetApp Support Site: support.netapp.com](#)

[Interoperability Matrix: support.netapp.com/matrix](#)

Creating LUNs and mapping igroups for iSCSI

As part of configuring your SAN environment, you have to create LUNs and map them to initiator groups.

Before you begin

You must have completed the following tasks:

1. Created your aggregates, volumes, and Storage Virtual Machines (SVMs)
2. Enabled block access with iSCSI
3. Created your LIFs

About this task

When you create a LUN, you must specify the LUN OS type. After the LUN is created, you cannot modify the LUN host operating system type.

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.

Note: You can grow a LUN to 10 times its original size. For example, if you create a 100 GB LUN, you can grow that LUN to 1,000 GB. However, you cannot exceed 16 TB, which is the maximum LUN size limit.

Note: The actual size of the LUN may be slightly different based on the OS type of the LUN.

Note: ALUA is always enabled. You cannot change the ALUA setting.

Steps

1. Use the `lun create` command to create your LUNs.

If you want...	Then use...
a space-reserved LUN	space-reserve enabled
a thinly provisioned LUN	space-reserve disabled

Example

```
lun create -vserver vs3 -volume vol3 -lun lun0 -size 50GB -ostype windows -space-reserve disabled
```

2. Use the `igroup create` command to create your igroups.

Example

```
igroup create -vserver vs3 -igroup igroup1 -protocol iscsi -ostype
windows -initiator iqn.1991-05.com.microsoft:example
```

3. Use the `lun map` command to map your LUNs to igroups.

Example

```
lun map -vserver vs3 -volume vol3 -lun lun0 -igroup igroup1
```

4. Use the `lun show` command to verify that you configured your LUNs correctly.

Example

```
lun show -vserver vs3
```

Vserver	Path	State	Mapped	Type	Size
vs3	/vol/vol3/lun0	online	mapped	windows	50.01GB
vs3	/vol/vol3/lun1	online	mapped	windows	50.01GB
vs3	/vol/vol3/lun2	online	mapped	windows	50.01GB

3 entries were displayed.

After you finish

If you want to limit access to your LUNs to specific LIFs, you must create port sets and bind the port sets to the igroup.

Related concepts

[Considerations when creating a LUN from an existing file](#) on page 55

[How LUN access works in a virtualized environment](#) on page 60

[How to limit LUN access in a virtualized environment](#) on page 63

[Guidelines for mapping LUNs to igroups](#) on page 56

[How the host operating system affects usable space for a LUN](#) on page 58

[How Data ONTAP avoids igroup mapping conflicts](#) on page 59

[What ALUA is](#) on page 125

Creating port sets and binding igroups to port sets

You have the option of using port sets to limit which LIFs can be used by an initiator to access a LUN.

Before you begin

You must have created a LIF and an igroup.

About this task

If you do not bind a port set to an igroup, then all the initiators in the igroup can access mapped LUNs through all the LIFs.

Steps

1. Use the `portset create` command to create a port set containing the appropriate LIFs.

Example

```
portset create -vserver vs3 -portset portset0 -protocol iscsi -port-name
lif0 lif1
```

2. Use the `lun igroup bind` command to bind an existing igroup to a port set.

```
lun igroup bind -vserver vs3 -igroup igroup1 -portset portset0
```

3. Use the `portset show` command to verify your port sets and LIFs.

Example

```
portset show -vserver vs3
```

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

Enabling block access for a specific host with iSCSI

Although the actual steps for enabling block access for your host could vary depending on your host, use this procedure as a guide for what you must do to connect your host to the Storage Virtual Machine (SVM).

Before you begin

Before you can connect your initiator to the storage system, the following conditions must be met:

- Ensure network connectivity between the host and the LIFs on the SVM.
- Verify that the iSCSI service is on and operational.
- Create igroups.
- Add your iSCSI initiator to the igroup.
- Map LUNs to igroups.

Step

1. Use the Host Utilities to complete the iSCSI mapping and to discover your LUNs on the host.

For more information, see the Host Utilities documentation.

Result

You have connectivity to your LUNs.

Related information

Interoperability Matrix: support.netapp.com/matrix

Host Utilities documentation: www.support.netapp.com

Setting up LUNs manually with FC block access

When you set up your LUNs manually, you must complete several required tasks before you can access your LUNs. After you complete all the required tasks, you can access your LUNs.

Before you begin

Before you set up your LUNs manually, you should first verify the following:

- Your SAN configuration is supported.
For more information about supported configurations, see the Interoperability Matrix Tool.
- Your SAN environment meets the SAN host and controller configuration limits.
For more information about limits, see the *Clustered Data ONTAP SAN Configuration Guide* for your version of Data ONTAP software.
- A supported version of Host Utilities is installed.
For more information, see the host utilities documentation.
- You are familiar with the FC best practices reviewed in TR-4017.

Steps

1. [Configuring FCoE](#) on page 72
If you are using FCoE on your Ethernet infrastructure, FCoE must be configured at the switch level before your FC service can run over the existing Ethernet infrastructure.
2. [Verifying the license for FC or iSCSI](#) on page 73
Before you can enable block access for a Storage Virtual Machine (SVM) with FC or iSCSI, you must have a license.
3. [Enabling block access for an SVM with FC](#) on page 74
You must create LIFs for a Storage Virtual Machine (SVM) and assign the FC protocol to those LIFs.
4. [Creating LUNs and mapping igroups for FC](#) on page 76
As part of configuring your SAN environment, you have to create LUNs and map them to initiator groups.
5. [Optional: Creating port sets and binding igroups to port sets](#) on page 78
You have the option of using port sets to limit which LIFs can be used by an initiator to access a LUN.
6. [Enabling block access for a specific host with FC](#) on page 79

While the actual steps for enabling block access for your host could vary depending on your host, use this procedure as a guide for what needs to happen to connect your host to the Storage Virtual Machine (SVM).

Related information

Interoperability Matrix: support.netapp.com/matrix

Host Utilities documentation: www.support.netapp.com

Documentation on the NetApp Support Site: support.netapp.com

SAN Configuration Guide for Cluster-Mode: support.netapp.com

Fibre Channel SAN Best Practices - <http://media.netapp.com/documents/tr-4017.pdf>

Configuring FCoE

If you are using FCoE on your Ethernet infrastructure, FCoE must be configured at the switch level before your FC service can run over the existing Ethernet infrastructure.

About this task

You must install a Unified Target Adapter (UTA) on your storage system and a converged network adapter (CNA) on your host. These adapters are required for running FCoE traffic over your Ethernet network. For more information about FCoE deployments, see TR-3800.

Your SAN configuration must be supported. For more information about supported configurations, see the Interoperability Matrix Tool.

Steps

1. Install the UTA or CNA.
2. Configure your switches for FCoE. See the switch documentation for more information.
3. On the storage controller, use the `dcb show` command to verify that the DCB settings for each node in the cluster has been correctly configured.

Example

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

Interface	PGID	Priority	Application	Bandwidth
e3a	0	0 1 2 4 5 6 7	unassigned	50%
	1	3	FCoE	50%
e3b	0	0 1 2 3 4 5 6 7	unassigned	0%
e4a	0	0 1 2 4 5 6 7	unassigned	50%
	1	3	FCoE	50%

```
e4b
      0      0 1 2 3 4 5 6 7      unassigned      0%
```

Note: Data ONTAP displays the switch settings for FCoE based on how you set up your switches. If these switch settings are wrong, you need to consult your switch documentation.

4. On the storage controller, use the `fc adapter show` command to verify that FCoE login is working when the FC target port status is online true state.

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

```
node      adapter  physical      status      fabric      speed state
-----  -
node-1    4a           ethernet      online true   auto  up
node-2    4b           ethernet      online true   auto  up
```

Note: If the FC target port status is online false, you need to consult your switch documentation.

Result

After you have verified that the DCB settings have been correctly configured, you can continue manually setting up your FC block access.

Related information

Brocade switch FCoE set up procedures - <http://www.brocade.com/products/all/index.page>

Cisco switch FCoE set up procedures - http://www.cisco.com/en/US/products/ps9670/products_installation_and_configuration_guides_list.html

Interoperability Matrix: support.netapp.com/matrix

Host Utilities documentation: www.support.netapp.com

Technical Report: Fibre Channel over Ethernet (FCoE) End-to-End Deployment Guide: media.netapp.com/documents/TR-3800.pdf

Data ONTAP documentation on the NetApp Support Site-support.netapp.com

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

1. 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 an SVM with FC

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.

About this task

You need a minimum of two LIFs per node. A host needs at least two paths to the LUN through each node.

At a minimum, you need paths to both nodes of the HA pair on which you create the LUN. You can have paths to a maximum of four nodes. If you have more than four nodes in your cluster, you can use port sets to limit the number of paths.

Steps

- Use the `vserver fcp create` command to enable FC service on the SVM.

Example

```
vserver fcp create -vserver vs3 -status-admin up
```

- Use the `network interface create` command to create the LIFs for a specific SVM.

Example

```
network interface create -vserver vs3 -lif lif2 -role data -data-protocol fcp -home-node node1 -home-port 0a
```

Note: You should create at least one LIF for each node in all SVMs serving data with FC protocol.

Note: To create an FC LIF, you must specify the `-role` parameter as `data` and the `-protocol` parameter as `fc`.

- Use the `network interface show` command to verify your LIFs.

Example

```
network interface show -vserver vs3
```

Vserver	Logical Interface	Status Admin/Oper	Network Address/Mask	Current Node	Current Port	Is Home
vs3	lif2	up/up	20:0b:00:a0:98:27:db:a2	node1	0a	true
	lif0	up/up	20:08:00:a0:98:27:db:a2	node1	0b	true
	lif1	up/up	20:09:00:a0:98:27:db:a2	node2	0b	true
3 entries were displayed.						

- Use the `fc topology show` command to verify that the N_Port ID Virtualization (NPIV) is enabled for each node in the cluster.

Example

```
run -node node1 -command fc topology show
```

```
Switches connected on adapter 0c:
```

```
Switch Name: sjlab-sw3
Switch Vendor: Cisco Systems, Inc.
Switch Release: 4.2(1)N1(1)
Switch Domain: 75
Switch WWN: 23:e9:00:0d:ec:b4:94:01
```

Port	Port WWPN	State	Type	Attached WWPN

vfc1	20:00:00:0d:ec:b4:94:3f	Offline	none	
vfc2	20:01:00:0d:ec:b4:94:3f	Offline	none	
vfc3	20:02:00:0d:ec:b4:94:3f	Offline	none	
vfc4	20:03:00:0d:ec:b4:94:3f	Offline	none	
vfc5	20:04:00:0d:ec:b4:94:3f	Offline	none	
vfc6	20:05:00:0d:ec:b4:94:3f	Offline	none	
vfc7	20:06:00:0d:ec:b4:94:3f	Offline	none	
vfc8	20:07:00:0d:ec:b4:94:3f	Offline	none	
2/1	20:41:00:0d:ec:b4:94:00	Offline	none	
2/2	20:42:00:0d:ec:b4:94:00	Online	F-Port	50:0a:09:80:00:05:a2:40
2/3	20:43:00:0d:ec:b4:94:00	Online	F-Port	50:0a:09:85:8d:9d:bf:ce

```
20:01:00:a0:98:13:d7:fc
2/4    20:44:00:0d:ec:b4:94:00    Online    F-Port 50:0a:09:85:8d:ad:c0:02
```

All WWPNs that start with a 50 are physical WWPNs, and all WWPNs that start with a 20 are LIF WWPNs. If you do not see LIF WWPNs logged into the FC switch, but you *do* see the physical WWPNs, then you do not have NPIV enabled. Consult your switch documentation to enable NPIV.

Note: Some columns have been removed for readability.

After you finish

You can create and map your LUNs.

Related concepts

Considerations for LIFs in cluster SAN environments on page 63

Related information

Documentation on the NetApp Support Site: support.netapp.com

Interoperability Matrix: support.netapp.com/matrix

Creating LUNs and mapping igroups for FC

As part of configuring your SAN environment, you have to create LUNs and map them to initiator groups.

Before you begin

You must have done the following tasks:

1. You have created your aggregates, volumes, and Storage Virtual Machines (SVMs).
2. You have enabled block access with FC.
3. You have created your LIFs.

About this task

When you create a LUN, you must specify the LUN OS type. After the LUN is created, you cannot modify the LUN host operating system type.

The metadata for each LUN requires approximately 64K of space in the containing aggregate. When you create a LUN, you must ensure that the containing aggregate has enough space for that 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.

Note: You can grow a LUN to approximately 10 times its original size. For example, if you create a 100 GB LUN, you can grow that LUN to approximately 1,000 GB. However, you cannot exceed 16 TB, which is the approximate maximum LUN size limit.

Note: The actual size of the LUN may be slightly different based on the OS type of the LUN.

Note: ALUA is always enabled. You cannot change the ALUA setting.

Steps

1. Use the `lun create` command to create your LUNs.

If you want	Then use ...
a space-reserved LUN	space-reserve enabled
a thinly provisioned LUN	space-reserve disabled

Example

```
lun create -vserver vs3 -volume vol3 -lun lun0 -size 50GB -ostype windows -space-reserve disabled
```

2. Use the `igroup create` command to create your igroups.

Example

```
igroup create -vserver vs3 -igroup igroup1 -protocol fcp -ostype windows -initiator 20:00:00:0d:ec:b4:94:3f
```

3. Use the `lun map` command to map your LUNs to igroups.

Example

```
lun map -vserver vs3 -volume vol3 -lun lun0 -igroup igroup1
```

4. Use the `lun show` command to verify that you configured your LUNs correctly.

Example

```
lun show -vserver vs3
```

```

Vserver    Path                State  Mapped  Type    Size
-----
vs3        /vol/vol3/lun0     online mapped  windows 50.01GB
vs3        /vol/vol3/lun1     online mapped  windows 50.01GB
vs3        /vol/vol3/lun2     online mapped  windows 50.01GB
3 entries were displayed.
```

After you finish

If you want to limit access to your LUNs, you must create port sets.

Related concepts

- [Considerations when creating a LUN from an existing file](#) on page 55
- [How LUN access works in a virtualized environment](#) on page 60
- [How to limit LUN access in a virtualized environment](#) on page 63
- [Guidelines for mapping LUNs to igroups](#) on page 56
- [How the host operating system affects usable space for a LUN](#) on page 58
- [How Data ONTAP avoids igroup mapping conflicts](#) on page 59
- [What ALUA is](#) on page 125

Optional: Creating port sets and binding igroups to port sets

You have the option of using port sets to limit which LIFs can be used by an initiator to access a LUN.

Before you begin

You must have created LIFs and mapped igroups to your LUNs.

About this task

If you do not bind a port set to an igroup, then all the initiators in the igroup can access mapped LUNs through all the LIFs.

Steps

1. Use the `portset create` command to create a port set containing the appropriate LIFs.

Example

```
portset create -vserver vs3 -portset portset0 -protocol fcp -port-name lif0 lif1
```

2. Use the `lun igroup bind` command to bind an existing igroup to a port set.
`lun igroup bind -vserver vs3 -igroup igroup1 -portset portset0`
3. Use the `portset show` command to verify your port sets and LIFs.

Example

```
portset show -vserver vs3
```

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

Related concepts

- [How to limit LUN access in a virtualized environment](#) on page 63

Enabling block access for a specific host with FC

While the actual steps for enabling block access for your host could vary depending on your host, use this procedure as a guide for what needs to happen to connect your host to the Storage Virtual Machine (SVM).

Before you begin

Before you can connect your host to the storage system, you need to do the following:

- Ensure network connectivity between the host and the LIFs on the SVM.
- Verify that the FC service is on and operational.
- Create initiator groups (igroups).
- Add your FC initiator to the igroup.
- Create and map your LUNs.

About this task

Note: You can create a WWPN alias for each initiator for ease of use.

Step

1. Use the Host Utilities and *Clustered Data ONTAP System Administration Guide for Cluster Administrators* to complete the FC mapping and to discover your LUNs on the host.

For more information, see the Host Utilities documentation.

Result

You have connectivity to your LUNs.

Related information

Interoperability Matrix: support.netapp.com/matrix

Host Utilities documentation: www.support.netapp.com

Setting up LUNs using the SVM setup wizard

The Storage Virtual Machine (SVM) setup wizard sets up and configures block access, SVMs, volumes, LUNs, and igroups.

Before you begin

- Your SAN configuration must be supported.
For more information about supported configurations, see the Interoperability Matrix Tool.
- Your SAN environment must meet the SAN host and controller configuration limits.

For more information, see 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.

About this task

You can re-enter the SVM setup wizard to add services. For example, if SVM vs5 already exists, you can use the SVM setup wizard to add additional LUNs to vs5.

When you use the SVM setup wizard, LUNs are automatically thinly provisioned. If you want to change the thinly provisioned LUN to a space-reserved LUN, you can use the `lun modify` command and set the `-space-reserve` parameter to enabled.

You cannot configure FCoE using the `vserver setup` command. See the topic *Verifying that DCB is configured correctly* to verify your FCoE configuration.

Steps

1. Use the `vserver setup` command to start the wizard.

Example

```
vserver setup
```

2. Follow the instructions on the screen.

After you finish

If you want to limit which LIFS an initiator can use to access LUNs, you need to create port sets and then bind the port sets to the igroups.

Related concepts

[Volume option best practices for thinly provisioned LUNs](#) on page 17

[What LUN thin provisioning is](#) on page 15

[Unified Ethernet network management](#) on page 149

Related tasks

[Verifying that DCB is configured correctly](#) on page 150

Related information

[Interoperability Matrix: support.netapp.com/matrix](http://support.netapp.com/matrix)

[Documentation on the NetApp Support Site: support.netapp.com](http://support.netapp.com)

[Host Utilities documentation: www.support.netapp.com](http://www.support.netapp.com)

Managing data protection in your SAN environment

Data ONTAP provides a variety of methods for protecting your data. These methods are based on Snapshot technology in Data ONTAP, which enables you to maintain online multiple read-only versions of LUNs per volume.

You can protect the data in your LUNs in the following ways:

- Clone LUNs from an active volume or a Snapshot copy
- Copy all LUNs from one node to another
- Restore LUNs from a Snapshot copy
- Back up LUNs through a host backup system

See the *Clustered Data ONTAP Data Protection Guide* for more information about data protection.

Related references

[Data protection methods](#) on page 81

Related information

Data ONTAP documentation on the NetApp Support Site-support.netapp.com

Data protection methods

Data ONTAP provides multiple methods for protecting your data.

Snapshot copy

Make point-in-time copies of a volume.

volume copy command

Perform fast block-copy of data from one volume to another.

In systems running clustered Data ONTAP, this method is valid only for FlexVol volumes.

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.

SnapVault backups (SnapVault license required)

- Back up data by using Snapshot copies on the storage system and transferring them on a scheduled basis to a secondary storage system.
- Store these Snapshot copies on the secondary storage system for weeks or months, allowing recovery operations to occur nearly instantaneously from the secondary storage system to the original storage system.

Data protection mirror copies (SnapMirror license required)

- Replicate data or asynchronously mirror data from one storage system to another over local or wide area networks (LANs or WANs).
- Transfer Snapshot copies taken at specific points in time to other storage systems. These replication targets can be in the same data center through a LAN, or distributed across the globe connected through metropolitan area networks (MANs) or WANs. Because SnapMirror operates at the changed block level instead of transferring entire files or file systems, it usually reduces bandwidth and transfer time requirements for replication.

SnapRestore (license required)

- Restore a LUN or file system to an earlier preserved state in less than a minute without rebooting the storage system, regardless of the size of the LUN or volume being restored.
- Recover from a corrupted database or a damaged application, file system, LUN, or volume by using an existing Snapshot copy.

SnapDrive for Windows or UNIX (SnapDrive license required)

- Manage storage system Snapshot copies directly from a Windows or UNIX host.
- Manage storage (LUNs) directly from a host.
- Configure access to storage directly from a host.

Note: For more information about SnapDrive as well as the supported Windows and UNIX environments, see the *SnapDrive for Windows Installation and Administration Guide* or *SnapDrive for UNIX Installation and Administration Guide*.

Native tape backup and recovery

Data ONTAP supports native tape backup and recovery. Support for most existing tape drives is included, as well as a method for tape vendors to dynamically add support for new devices. In addition, Data ONTAP supports the Remote Magnetic Tape (RMT) protocol, enabling backup and recovery to any capable system. For more information about tape backup and recovery, see the *Clustered Data ONTAP Data Protection Tape Backup and Recovery Guide*.

NDMP

Control native backup and recovery facilities in storage systems and other file servers. Backup application vendors provide a common interface between backup applications and file servers.

NDMP is an open standard for centralized control of enterprise-wide data management. For more information about how NDMP-based topologies can be used by storage systems to protect data, see the *Clustered Data ONTAP Data Protection Tape Backup and Recovery Guide*.

Related information

Data ONTAP documentation on the NetApp Support Site: support.netapp.com

Copying all LUNs on a volume with the volume copy command

You can use the `volume copy start` command to copy all the LUNs in a volume. When you use this command, everything on the volume is copied including files, if any exist, on the volume. You can copy a volume to a different aggregate on a different node within a cluster.

Before you begin

Applications accessing the LUNs should be quiescent and offline prior to the copy operation. The contents of host file system buffers must be saved to disk.

About this task

The `volume copy start` command enables you to copy all the LUNs from one volume to another, either within the same storage system or to a different storage system. The size of the volume could impact the copy operation of the volume.

If the copying takes place between two storage systems, you can enter the `volume copy start` command on either the source or destination storage system. You cannot, however, enter the command on a third storage system that does not contain the source or destination volume.

Steps

1. Use the command `volume copy start` to copy all the LUNs in the specified volume to the same or different storage system.

Example

```
volume copy start -vserver vs3 -volume sm_vol3 -destination-volume  
sm_vol3_copy -destination-aggregate aggr2
```

```
[Job 775] Job is queued: Copy sm_vol3 in Vserver vs3 to sm_vol3_copy.
```

2. Optional: Use the `job show` command to view the status of the copy operation.

Example

```
job show -instance -id 775
```

```

        Job ID: 775
    Owing Vserver: innovate
        Name: Volume Copy
    Description: Copy sm_vol3 in Vserver vs3 to sm_vol3_copy
    Priority: High
        Node: node01
    Affinity: Cluster
    Schedule: @now
    Queue Time: 04/14 15:27:07
    Start Time: 04/14 15:27:07
    End Time: 04/14 15:27:11
    Drop-dead Time: -
    Restarted?: false
        State: Success
    Status Code: 0
    Completion String:
        Job Type: COPY
    Job Category: VOPL
    Execution Progress: Complete: [0]
    User Name:

```

3. Use the `volume show` command to verify that the copy operation is complete.

Example

```
volume show
```

Vserver	Volume	Aggregate	State	Type	Size	Available	Used%
node01	vol3	aggr0	online	RW	352.6GB	327.7GB	7%
vs1	lun_volumel	aggr1	online	RW	1.46TB	1.34TB	8%
vs1	root_vol1	aggr1	online	RW	20MB	18.88MB	5%
vs2	lun_vol_2	aggr1	online	RW	100GB	94.82GB	5%
vs2	root_vol2	aggr1	online	RW	20MB	18.88MB	5%
vs2	sm_vol1	aggr1	online	RW	100MB	94.79MB	5%
vs3	root_vol	aggr1	online	RW	20MB	18.88MB	5%
vs3	sm_vol3	aggr1	online	RW	100MB	94.85MB	5%
vs3	sm_vol3_copy	aggr2	online	RW	100MB	94.85MB	5%
vs3	sm_vol1	aggr1	online	RW	100MB	94.85MB	5%
vs3	sm_vol1_copy	aggr2	online	RW	100MB	94.85MB	5%
vs3	vol3	aggr1	online	RW	100MB	94.81MB	5%

12 entries were displayed.

Using Snapshot copies to protect your data

A Snapshot copy is a copy of a FlexVol volume representing the volume's contents at a particular point in time. You can view the contents of the Snapshot copy and use the Snapshot copy to restore data that you lost.

See the *Clustered Data ONTAP Data Protection Guide* for more information.

Related information

[Data ONTAP documentation on the NetApp Support Site-support.netapp.com](http://support.netapp.com)

Required changes to Snapshot copy default settings

When you create a volume, Data ONTAP automatically schedules Snapshot copies and reserves space for them. You must modify these default settings to ensure that overwrites to LUNs in the volume do not fail.

Data ONTAP Snapshot copies are the basis of many optional features, such as the SnapMirror feature, SyncMirror feature, and tape backup features.

Data ONTAP automatically performs the following operations:

- Reserves 5 percent of the space for Snapshot copies
- Schedules Snapshot copies

Because the internal scheduling mechanism for taking Snapshot copies within Data ONTAP has no means of ensuring that the data within a LUN is in a consistent state, you should change these Snapshot copy settings by performing the following tasks:

- Turn off the automatic Snapshot copy schedule.
- Delete all existing Snapshot copies.
- Set the percentage of space reserved for Snapshot copies to zero (0).

Restoring a single LUN from a Snapshot copy

You can use `volume snapshot restore-file` command to restore a single LUN without restoring the entire volume that contains the single LUN.

About this task

You should unmount the LUN from the host before you run the `volume snapshot restore-file` command on the volume that contains the LUN that you want to restore.

When restoring a single LUN with SnapRestore, you need as much available space as the size of the LUN you are restoring. For example, if you are restoring a 10 GB LUN, then you need 10 GB of available space on the volume to perform the SnapRestore.

However, in some circumstances, a LUN could require more or less space:

- If you are restoring a space-reserved LUN where the fractional reserve is 0%, you need 1 times the size of the restored LUN.
- If you are restoring a space-reserved LUN where the fractional reserve is 100%, you need 2 times the size of the restored LUN.
- If you are restoring a thinly provisioned LUN, you only need the actual space used for the restored LUN.

Note: If your destination LUN is a different OS type than your source LUN, your host can lose data access to the destination LUN after the restore operation.

Attention: If the restore operation fails, your destination LUN might be truncated, resulting in data loss. You must create a Snapshot copy of the destination LUN before restoring the destination LUN. If the destination LUN is truncated, you can use the latest Snapshot copy of the destination LUN.

In rare cases, the LUN restore can fail, leaving the original LUN unusable. Before restoring a LUN, make an additional Snapshot copy of the LUN you want to restore. Then if the restore fails, you can at least return the LUN to the state just before the restore attempt.

Steps

1. Check available space. SnapRestore does not revert the Snapshot copy if sufficient space is not available.
2. From the host, stop all host access to the LUN.
3. Unmount the LUN on its host to ensure that the host does not access the LUN.
4. Optional: Use the `lun unmap` command to unmap the LUN.

Example

```
lun unmap -vserver vs3 -volume vol3 -lun lunX_00 -igroup vs3_igrp
```

5. Use the `volume snapshot show` command to determine the snapshot copy you want to restore your LUN to.

Example

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

Vserver	Volume	Snapshot	Size	Total%	Used%
vs3	vol3	weekly.2013-04-10_0015	120KB	0%	37%
		snap_vol3	84KB	0%	29%

2 entries were displayed.

6. Optional: Use the `volume snapshot create` command to create a Snapshot copy of the LUN prior to restoring the LUN.

```
volume snapshot create -vserver vs3 -volume vol3 -snapshot snap_vol3
```

7. Use the `volume snapshot restore-file` command to restore the specified LUN in a volume.

Example

```
volume snapshot restore-file -vserver vs3 -volume vol3 -snapshot
snap_vol3 -path /vol/vol3/lunx_00
```

8. Follow the steps on the screen.
9. Use the `lun modify` command if the LUN is not online.

Example

```
lun modify -vserver vs3 -path /vol/vol3/lunx_00 -state online
```

10. Optional: Use the `lun map` command if your LUN is unmapped.


```
lun map -vserver vs3 -volume vol3 -lun lunX_00 -igroup vs3_igrp
```
11. From the host, remount the LUN.
12. From the host, restart access to the LUN.

Restoring all the 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.

About this task

When restoring a volume using SnapRestore, you only need as much available space as the size of the volume you are restoring. For example, if you are restoring a 10 Gb volume, then you only need 10 Gb of available space to perform the SnapRestore.

Attention: When you restore a volume, you must stop all data access to all LUNs in that volume. On the host, you should unmount the LUNs to stop data access before you run the `volume snapshot restore` command. On the storage system, you should unmap the LUNs. Using SnapRestore on a volume that contains LUNs without stopping all host access to those LUNs can cause data corruption and system errors.

Steps

1. Check available space. SnapRestore does not revert the Snapshot copy if sufficient space is not available.
2. From the host, stop all host access to the LUNs.
3. Unmount the LUNs on that host to ensure that host does access the LUNs.
4. Optional: Use the `lun unmap` command to unmap your LUNs.

Example

```
lun unmap -vserver vs3 -volume vol3 -lun lunX_00 -igroup vs3_igrp
```

- Use the `volume snapshot show` command to determine the Snapshot copy to which you want to restore your volume.

Example

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

Vserver	Volume	Snapshot	Size	Total%	Used%
vs3	vol3	weekly.2013-02-10_0015	120KB	0%	37%
		snap_vol3	84KB	0%	29%

2 entries were displayed.

- Use the `set` command to change your privilege setting to advanced.

Example

```
set -privilege advanced
```

- Use the `volume snapshot restore` command to restore your data.

Example

```
volume snapshot restore -vserver vs3 -volume vol3 -snapshot snap_vol3
```

- Follow the instructions on the screen.
- If you unmapped your LUNs, use the `lun map` command to remap your LUNs.

Example

```
lun map -vserver vs3 -volume vol3 -lun lunX_00 -igroup vs3_igrp
```

- Use the `lun show` command to verify that your LUNs are online.

Example

```
lun show -vserver vs3 -path /vol/vol3/lun lunX_00 -fields state
```

vserver	path	state
vs3	/vol/vol3/lunX_00	online

- Use the `lun modify` command if any of the LUNs are not online.

Example

```
lun modify -vserver vs3 -path /vol/vol3/lun lunx_00 -state online
```

- Use the `set` command to change your privilege setting to admin.

Example

```
set -privilege admin
```

13. From the host, remount the LUNs.
14. 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, writable 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 thinly provisioned, the FlexClone LUN is also thinly provisioned. However, you can create a thinly provisioned 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.

When you create a FlexClone LUN, the FlexClone LUN automatic deletion function is enabled by default.

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 will be deleted.

If you do not want a particular a FlexClone LUN to be automatically deleted, you must set the auto delete parameter to disable for that FlexClone LUN.

Note: When you create a FlexClone LUN from a Snapshot copy, a background split between the FlexClone LUN and the Snapshot copy is automatically triggered. If this background split has not completed and this Snapshot copy is autodeleted, that FlexClone LUN will be deleted even if you have disabled the FlexClone auto delete function for that FlexClone LUN. After the background split completes, the FlexClone LUN is not deleted even if that Snapshot copy is deleted.

The FlexClone LUN or parent LUN does not consume additional disk space until changes are made to the FlexClone LUN or the parent LUN. For more information about FlexClone LUNs and volumes, see the *Clustered Data ONTAP Logical Storage Management Guide*.

Related information

[Data ONTAP documentation on the NetApp Support Site-support.netapp.com](http://Site-support.netapp.com)

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.

What FlexClone LUN autodelete is

FlexClone LUN autodelete is a volume-level option that allows you to define a policy for deleting FlexClone LUNs based on a definable threshold. By default all FlexClone LUNs are enabled for auto delete.

When you configure FlexClone LUN autodelete, FlexClone LUNs are automatically deleted and space reclaimed when the available space in a volume goes below a pre-defined threshold. If a particular FlexClone LUN should not be deleted automatically, you can change this default with the `volume file clone autodelete` command.

How a FlexVol volume can reclaim free space from FlexClone LUNs

Starting with Data ONTAP 8.2, you can configure the autodelete settings of a FlexVol volume to automatically delete FlexClone LUNs when the free space in a volume decreases below a particular threshold value.

A FlexVol volume that has the autodelete capability enabled resorts to automatic deletion of FlexClone LUNs only in the following situations:

- If the volume does not have Snapshot copies for automatic deletion.
- If the volume has Snapshot copies but the automatic deletion of Snapshot copies does not create sufficient free space in the volume.

Because the autodelete settings enable a FlexVol volume to directly delete FlexClone LUNs when the volume requires free space, you can preserve certain FlexClone LUNs by preventing them from getting automatically deleted.

Note: If the FlexVol volume contains FlexClone LUNs created using Data ONTAP versions earlier than 8.2 and if you want to delete them to increase the amount of free space in the volume, you can specify those FlexClone LUNs for automatic deletion.

Configuring a FlexVol volume to automatically delete FlexClone LUNs

Starting with Data ONTAP 8.2, you can enable a FlexVol volume to automatically delete FlexClone LUNs when the free space in the volume decreases below a particular threshold. You can also enable automatic deletion for FlexClone LUNs created using Data ONTAP versions earlier than 8.2.

Before you begin

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

Steps

1. Enable automatic deletion of FlexClone LUNs in the FlexVol volume by using the `volume snapshot autodelete modify` command.

Example

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

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

Volume modify successful on volume:voll
```

The FlexClone LUNs created using Data ONTAP 8.2 or later in the FlexVol volume are automatically enabled for deletion when you enable FlexClone LUN autodelete in the volume.

Note: While enabling FlexClone LUNs for automatic deletion, if you set the value of the commitment parameter to `destroy`, all the FlexClone LUNs created using Data ONTAP 8.2 or later, including those marked for preservation can get deleted when the free space in the volume decreases below the specified threshold value.

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

Example

The following example shows how you can determine if volume `voll` is enabled for automatic deletion of FlexClone LUNs:

```
cluster1::> volume snapshot autodelete show -vserver vs1 -volume voll
Vserver Name: vs1
Volume Name: voll
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, sfsr
    Is Constituent Volume: false

```

3. If the FlexVol volume has FlexClone LUNs created using Data ONTAP versions earlier than 8.2, and if you want to enable these FlexClone LUNs for automatic deletion, then perform the following steps:

- a) Use the volume file clone autodelete command to enable automatic deletion of a particular FlexClone LUN.

Example

The following example shows how you can enable the automatic deletion of a FlexClone LUN lun1_clone contained in volume vol1:

```

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

```

- b) Use the lun show command to verify if the FlexClone LUN is enabled for automatic deletion.

Example

The following example shows how you can verify if the FlexClone LUN lun1_clone is enabled for automatic deletion:

```

cluster1::> lun show -instance -lun lun1_clone

Vserver Name: vs1
    LUN Path: /vol/vol1/lun1_clone
    Volume Name: vol1
    Qtree Name: ""
    LUN Name: lun1
    LUN Size: 50.01GB
    OS Type: windows
    Space Reservation: disabled
    Serial Number: BZHV0]C9Pj3U
    Comment:
    Space Reservations Honored: true
    Space Allocation: disabled
    State: online
    LUN UUID: b02b01da-438d-11e2-821a-123478563412
    Mapped: unmapped
    Block Size: 512
    Device Legacy ID: -
    Device Binary ID: -
    Device Text ID: -
    Read Only: false

```

```

Inaccessible Due to Restore: false
      Used Size: 0
      Maximum Resize Size: 502.0GB
      Creation Time: 12/11/2012 12:24:19
      Class: -
      Clone: -
Clone Autodelete Enabled: true
      QoS Policy Group: -

```

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 writable 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, 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
```

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

Result

The FlexClone LUN is created from the LUN in the active volume. To prevent a specific FlexClone LUN from being automatically deleted when the volume needs more room, you must disable Flex Clone LUN autodelete for that LUN.

Related concepts

[Reasons for using FlexClone LUNs](#) on page 90

[Using FlexClone LUNs to protect your data](#) on page 90

Cloning LUNs from a Snapshot copy in a volume

You can create copies of LUNs by cloning LUNs from the Snapshot copies in a volume. These FlexClone LUNs are readable and writable copies of the LUNs as the LUN existed when the Snapshot copy was created.

Before you begin

A FlexClone license must be installed.

About this task

The FlexClone LUN inherits the `-space-reserve` parameter of the parent LUN. For example, if the parent LUN is thinly provisioned (`-space-reserve` is disabled) then the FlexClone LUN is thinly provisioned too.

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

Attention: If you enabled Snapshot Autodelete on a volume with a destroy commitment level and the destroy list includes `file_clone` and `lun_clone`, volume Snapshot copies may be deleted while blocking sharing is in progress. This could result in an inconsistent FlexClone LUN.

Steps

1. Use the `volume snapshot create` command to create a Snapshot of the volume that contains the LUNs. You must create a Snapshot copy (the backing Snapshot copy) of the LUN you want to clone.

Example

```
volume snapshot create -vserver vs1 -volume voll -snapshot voll_snapshot
```

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

Example

```
file clone create -vserver vs1 -volume voll -source-path lun1 -snapshot-name voll_snapshot -destination-path lun1_snap_clone
```

3. Use the `lun show` command to verify that you created the FlexClone LUN `lun1_snap_clone`.

Example

```
lun show -vserver vs1
```

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

Result

A FlexClone LUN is created from the LUN in the Snapshot volume. If you do not want this specific FlexClone LUN to be automatically deleted when the volume needs more room, you must disable Flex Clone LUN autodelete for that LUN.

Related concepts

[Reasons for using FlexClone LUNs](#) on page 90

[Using FlexClone LUNs to protect your data](#) on page 90

Preventing a specific FlexClone LUN from being autodeleted

If you configure a FlexVol volume to automatically delete FlexClone LUNs, any FlexClone LUN that fits the criteria you configure might be deleted. If you have specific 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 8.2, when you create a FlexClone LUN, by default the FlexClone LUN can be deleted automatically to reclaim space on the volume if you have configured the volume for

FlexClone autodelete. If you have not configured the volume for FlexClone autodelete, FlexClone LUNs are not deleted automatically.

Note: When you create a FlexClone LUN from a Snapshot copy, a background split between the FlexClone LUN and the Snapshot copy is automatically triggered. If this background split has not completed and this Snapshot copy is autodeleted, that FlexClone LUN will be deleted even if you have disabled the FlexClone autodelete feature for that FlexClone LUN. After the background split is complete and if that Snapshot copy is deleted, the FlexClone LUN will not be deleted.

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

Steps

1. Use the `volume file clone autodelete` command to prevent a specific FlexClone LUN from being automatically deleted.

Example

```
volume file clone autodelete -vserver vs1 -volume voll -clone-path
lun1_clone -enable false
```

The FlexClone LUN `lun1_clone` cannot be deleted automatically to reclaim space on the volume.

2. To verify that the FlexClone LUN cannot be deleted automatically, use the `lun show` command.

Example

```
lun show -instance -lun lun1_clone
```

```
Vserver Name: vs1
  LUN Path: /vol/voll/lun1_clone
  Volume Name: voll
  Qtree Name: ""
    LUN Name: lun1_clone
    LUN Size: 50.01GB
    OS Type: windows
  Space Reservation: disabled
  Serial Number: 2FiMg]C2FFHF
  Comment:
Space Reservations Honored: true
  Space Allocation: disabled
    State: online
    LUN UUID: b02b01da-438d-11e2-821a-123478563412
    Mapped: unmapped
  Block Size: 512
  Device Legacy ID: -
  Device Binary ID: -
  Device Text ID: -
```

```

Read Only: false
Inaccessible Due to Restore: false
Used Size: 0
Maximum Resize Size: 502.0GB
Creation Time: 12/11/2013 12:24:19
Class: -
Clone: -
Clone Autodelete Enabled: false
QoS Policy Group: -

```

Result

FlexClone LUN lun1_clone cannot be autodeleted.

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 has a different identity from the source LUN and must first be mapped.

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.

About this task

You can access a LUN directly from the SnapVault secondary volume without having to first restore or clone the volume containing the LUN.

The LUN path and serial number are different from the source LUN. Persistent reservations, LUN mappings, and igroups are not replicated to the SnapVault secondary volume.

If you plan to access LUN copies directly from the SnapVault secondary volume, you should create your igroups on the SnapVault SVM in advance.

Note: 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 example 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 map` command to map the desired LUN copy to the igroup.

Example

```

cluster::> lun map -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.

About this task

The process starts with 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.

If you plan to restore single LUNs, you should create your igroups on the SnapVault SVM in advance.

Steps

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

Example

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

Source Path	Dest Type	Mirror Path	Relation 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.
```

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

Example

This example 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
```

6. Run the `lun map` command to map the desired LUN copy to the `igroup`.

Example

```
cluster::> lun map -vserver vserverB
               -path /vol/dstvolB_clone/lun_C -igroup temp_igroup
```

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

The LUN is read-write and can be used in place of the original LUN. Because the LUN serial number is different, the host interprets it as a different LUN from the original.

8. Use a copy program on the host to copy the LUN contents back to the original LUN.

Restoring all LUNs in a volume from a SnapVault backup

If one or more LUNs in a volume need to be restored from a SnapVault backup, you can restore the entire volume. Restoring the volume affects all LUNs in the volume.

Before you begin

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

About this task

Restoring an entire volume mostly returns the volume to the state it was in when the Snapshot copy was made. If a LUN was added to the volume after the Snapshot copy, that LUN is removed during the restore process.

After restoring the volume, the LUNs remain mapped to the igroups they were mapped to just before the restore. The LUN mapping might be different from the mapping at the time of the Snapshot copy. Persistent reservations on the LUNs from host clusters are retained.

Steps

1. Stop I/O to all LUNs in the volume.
2. Run the `snapmirror show` command to verify the secondary volume that contains the SnapVault secondary volume.

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	-

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 can add or remove LUNs

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
```

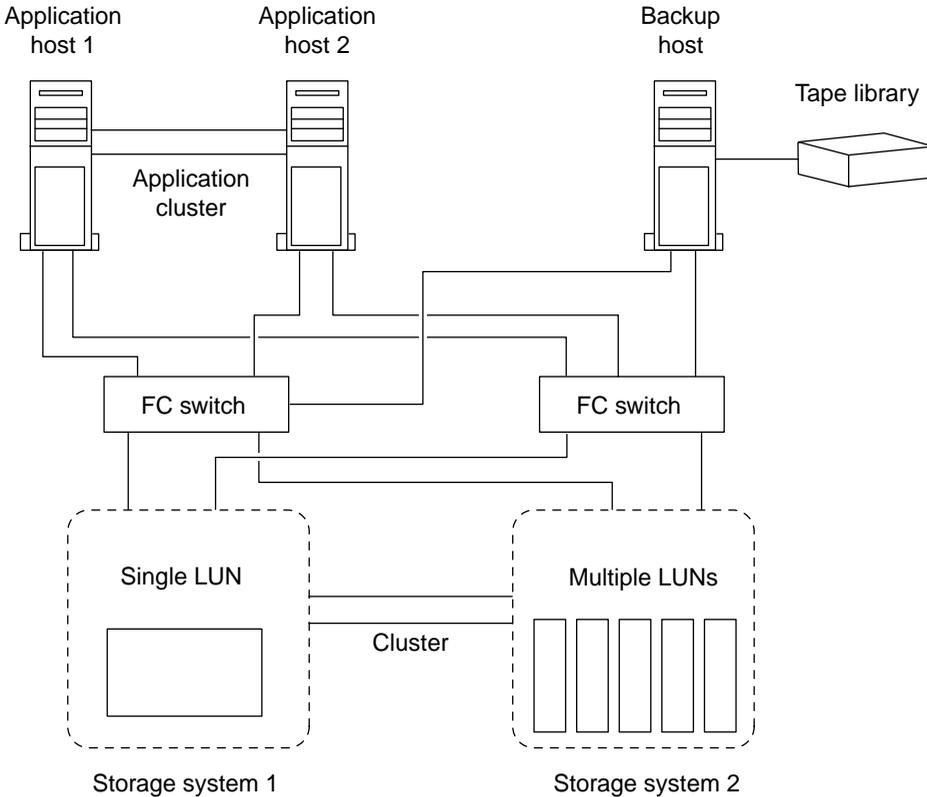
```
6 entries were displayed.
```

After the volume is restored from the SnapVault secondary volume, the source volume no longer contains lun_D. Note that the LUNs in the source volume are still mapped.

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

In most cases, you 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. You must configure volumes as SAN-only or NAS-only, and you must also configure qtrees within a single volume as SAN-only or NAS-only. From the point of view of the SAN host, 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.

For more information, see the *Clustered Data ONTAP Data Protection Tape Backup and Recovery Guide*.

Related tasks

[Backing up a LUN through a host backup system](#) on page 106

Related information

[Documentation on the NetApp Support Site: support.netapp.com](http://support.netapp.com)

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

- You must create a production LUN.

- You must create an igroup that the LUN belongs to.
The igroup must include the WWPN or initiator node name of the application server.
- You must map the LUN to the igroup.
- You must format the LUN and make it accessible to the host.

About this task

Step 1 can be part of your SAN backup application's pre-processing script. Steps 5 and 6 can be part of your SAN backup application's post-processing script.

Steps

1. When you are ready to start the backup (usually after your application has been running for some time in your production environment), save the contents of host file system buffers to disk using one of the following: the command provided by your host operating system, or using SnapDrive for Windows or SnapDrive for UNIX.

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 map` command to map the LUN clone you created in Step 3 to the backup host.

Example

```
lun map -vserver vs3 -volume vol3 -lun lun1_backup -igroup igroup3
```

6. From the host, discover the new LUN and make the file system available to the host.
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 105

Managing LUNs

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

You can use one of the following methods to create LUNs and igroups:

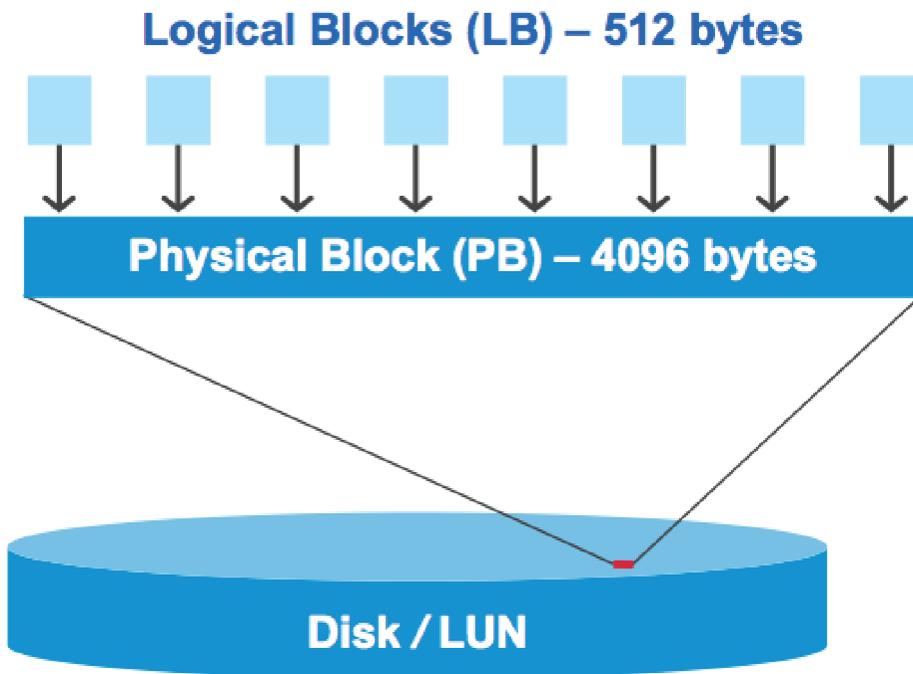
- Entering the `vserver setup` command
The Storage Virtual Machine (SVM) setup wizard prompts you through the process of creating a LUN, creating an igroup, and mapping the LUN to the igroup. You can reenter the SVM setup wizard to add services. For example, if SVM vs5 already exists, you can use the SVM setup wizard to add additional LUNs to vs5.
- Entering a series of individual commands (such as `lun create`, `igroup create`, and `lun map`)
You can use this method to create one or more LUNs and igroups in any order.

You can use the man page for a command to review all the parameter options.

Misaligned I/O can occur on properly aligned LUNs

Data ONTAP can 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 to the LUN. However, the presence of misaligned I/O does not necessarily mean 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.

Related concepts

[LUN alignment in virtual environments](#) on page 171

The effect of ostype on I/O alignment

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

The partition scheme employed by the host operating system is a major contributing factor to I/O misalignments. Some Data ONTAP LUN ostyles 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 ostyles with a prefix value greater than 0, a custom partition might actually create misaligned I/O.

LUN ostyle	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
wmware	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

There are various factors you need to consider when determining whether you should use the `solaris` ostyle or the `solaris_efi` ostyle..

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

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, files, and FlexVol volumes 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.

Note: Storage QoS is supported on clusters that have up to eight nodes.

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.

Increasing the size of a LUN

You can increase the size of your thinly provisioned or space-reserved LUNs with the `lun resize` command.

About this task

If the configured size of a LUN is filled and you cannot write additional blocks to that LUN, you can increase the size of the LUN with the `lun resize` command, provided your volume contains enough space. For example, if you configure the size of a LUN to be 50 GB and all 50 GB of that LUN is filled with data, you can use the `lun resize` command to increase the size of your LUN provided the volume has available space.

You do not have to take your LUN offline to increase the size of that LUN. When you increase the size of your LUN, Data ONTAP automatically notifies the initiator that the LUN size has increased.

You can grow a LUN to approximately 10 times its original size. For example, if you create a 100 GB LUN, you can grow that LUN to approximately 1,000 GB. However, you cannot exceed 16 TB, which is the approximate maximum LUN size limit.

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

Note: The actual size of the LUN may be slightly different based on the OS type of the LUN.

Steps

1. Use the `lun resize` command to increase the size of your LUN.

Example

```
lun resize -vserver vs3 -volume vol3 -lun lun1 -size 75GB
```

2. Use the `lun show` command to verify the LUN size.

Example

```
lun show -vserver vs3
```

Vserver	Path	State	Mapped	Type	Size
vs3	/vol/vol3/lun1	online	mapped	windows	75.00GB

Some columns have been removed for improved readability.

Decreasing the size of a LUN

You can decrease the size of your space-reserved or thinly provisioned LUN with the `lun resize` command.

About this task

You can use `lun resize` command to decrease the size of your LUN. However, you need to make sure when you decrease the size of your LUN that you do not truncate data on that LUN. Before you decrease the size of a LUN, the host needs to migrate the blocks containing the data into a new LUN boundary to prevent data loss. After you decrease the size of your LUN, Data ONTAP automatically notifies the initiator that the LUN size has decreased.

Note: The actual size of the LUN may be slightly different based on the OS type of the LUN.

Steps

1. Use the `lun resize` command to decrease the size of your LUN.

Example

```
lun resize -vserver vs5 -volume vol5 -lun lun5 -size 4GB -force
```

2. Use the `lun show` command to verify the LUN size.

Example

```
lun show -vserver vs5 -volume vol5 -lun lun5
```

Vserver	Path	State	Mapped	Type	Size
vs5	/vol/vol5/lun5	online	mapped	windows	4.01GB

Note: Some of the columns in this example were removed for readability.

Deleting LUNs

You can delete a LUN from a Storage Virtual Machine (SVM) if you no longer need the LUN.

About this task

If you try to delete a LUN that is mapped to an igroup, an error message is returned and the LUN is not deleted. If you want to delete a LUN that is mapped to igroup, you must use the force option or unmap the LUN.

Note:

- If you create a LUN from a file, you cannot remove the file while the LUN is linked to it. If you want to remove the file, you must first delete the LUN.
- 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 unmap` to unmap the LUN from the igroup.

Example

```
lun unmap -vserver vs5 -volume vo5 -lun lun5 -igroup igr5
```

3. Use the `lun modify` command to take the LUN offline.

```
lun modify -vserver vs5 -volume vol5 -lun lun5 -state offline
```

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

```
lun delete -vserver vs5 -volume vol5 -lun lun5
```

5. 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/vol116/lun8	online	mapped	windows	10.00GB

Examining configured and used space of a LUN

You can determine the configured space and the actual space used for thinly provisioned and space reserved LUNs.

About this task

By examining how space has been allocated and written for a set of LUNs, you can determine the following:

- The amount of space that can be reclaimed when doing space reclamation.
- The amount of reserved space that contains data.
- 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

```
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       0B
vs3      /vol/vospace/lun0        5.00GB   enabled       4.50GB
4 entries were displayed.

```

Commands for managing LUNs

Data ONTAP provides commands to 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 map</code>
Determine whether a LUN is mapped	<code>lun show</code>
Move a LUN within the same volume	<code>lun move</code>
Rename a LUN	<code>lun move</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>

If you want to...	Use this command...
Display detailed LUN information	<code>lun show -instance</code>
Display LUN paths and LUN IDs	<code>lun mapped show</code>
Unmap a LUN from an igroup	<code>lun unmap</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>
View a man page for a command	<code>man <i>command name</i></code>

See the man page for each command for more information.

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

[LUN availability on specific target LIFs](#) on page 58

[Considerations for LIFs in cluster SAN environments](#) on page 63

LIF migration in a SAN environment

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 Fibre Channel fabric or create new iSCSI sessions between the attached hosts of your cluster and its new target interface.

Note: In Data ONTAP 8.1 operating in Cluster-Mode, non-disruptive LIF migration requires advanced privileges. In clustered Data ONTAP 8.2, normal administrative privileges are sufficient.

Using LIF migration, you can do a non-disruptive upgrade if you are:

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

A LIF that has been assigned a SAN data protocol cannot be migrated using the `network interface migrate` command. LIF migrations for a SAN data protocol 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 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.

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 migrate a SAN LIF to preserve its configuration information, such as its WWPN, to avoid re-zoning the switch fabric. Because a SAN LIF must be taken offline before it is migrated, host traffic must rely upon host multipathing software to ensure non-disruptive access to the LUN. You can migrate SAN LIFs to any node in a cluster, but you cannot migrate the SAN LIFs between Storage Virtual Machines (SVMs).

Before you begin

If the LIF is a member of a port set, you must remove the LIF from the port set before you can move the LIF 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 Fibre Channel 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.

Note: When you take a LIF offline, an I/O interruption to the hosts occurs. Ensure that there are other paths that the host MPIO can use to access LUNs on the SVM.

Steps

1. Use the `network interface show` command to view the `Status Admin` of the LIF and the current port:

```
network interface show -vserver vs1
```

Logical Vserver	Status Interface	Network Admin/Oper	Current Node	Current Port	Is Home
vs1	lif1	up/up	node01	1a	true

The administration status of lif1 is up (online). The current node is node01 and the current port for lif1 is 1a.

2. Use the `network interface modify` command to change the `Status Admin` of the LIF to down (offline):

```
network interface modify -vserver vs1 -lif lif1 -status-admin down
```

3. Use the `network interface modify` command to assign lif1 a new node and port:

```
network interface modify -vserver vs1 -lif lif1 -home-node node02 -home-port 1b
```

4. Use the `network interface modify` command to change the `status admin` of the LIF to up (online):

```
network interface modify -vserver vs1 -lif lif1 -status-admin up
```

5. Use the `network interface show` command to verify your changes:

```
network interface show -vserver vs1
```

Logical Vserver	Status Interface	Network Admin/Oper	Current Node	Current Port	Is Home
vs1	lif1	up/up	node02	1b	true

LIF lif1 has been reassigned to `Current Node node02` and `Current Is Port 1b`.

Related tasks

[Removing a SAN LIF from a port set](#) on page 119

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 119

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.

Understanding igroup functions

To manage your igroups successfully, you have to understand some concepts about igroups, iSCSI initiator node names, WWPNs, MPIO, and ALUA. You also can review the man pages for additional information about igroup commands.

The information that you need to understand includes the following:

- What igroups are
- How you work with iSCSI initiator node names
- How you work with FC protocol initiator WWPNs
- What ALUA is
- How MPIO and ALUA work together

Note: When running clustered Data ONTAP 8.2.0 on a storage system attached to AIX hosts, you must make sure all the nodes in the cluster are running Data ONTAP 8.2.0.

Related concepts

[What ALUA is](#) on page 125

[MPIO and ALUA](#) on page 126

[What iSCSI nodes are](#) on page 179

[How FC nodes are identified](#) on page 182

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 HBAs or software initiators to have access to a LUN. If you are using multipathing software or have clustered hosts, each HBA 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.

Note: An initiator cannot be a member of igroups of differing ostyles.

Related concepts

[Managing igroups](#) on page 127

igroup example

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.

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 (connected to Host4) 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/ qtree1/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/ qtree1/lun4 /vol/vol2/ qtree1/lun5

About iSCSI initiator node names

You can specify the node names of the initiators when you create an igroup. You can also add or remove node names later.

To know which node names are associated with a specific host, see the Host Utilities documentation for your host. These documents describe commands that display the host's iSCSI node name.

Related concepts

What iSCSI nodes are on page 179

What Host Utilities are on page 177

Related information

Host Utilities documentation: www.support.netapp.com

FC protocol initiator WWPN

You can specify the WWPNs of the initiators when you create an igroup. You can also add them or remove them later.

For instructions on obtaining the host identifiers (WWPN or IQN), see the Host Utilities documentation for your host operating system. 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 182

What Host Utilities are on page 177

Related information

Host Utilities documentation: www.support.netapp.com

What ALUA is

Data ONTAP uses Asymmetric Logical Unit Access (ALUA) for iSCSI and FC to identify optimized paths.

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.

Attention: Host support for ALUA is required for cluster failover to work correctly.

Related information

[Interoperability Matrix: support.netapp.com/NOW/products/interoperability](http://support.netapp.com/NOW/products/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 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>
Modify attributes for an igroup	<code>lun igroup modify</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 123

Configuring your iSCSI 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).

Managing your iSCSI service

You must ensure that the iSCSI service is licensed and running on your system. You can manage the target node name, target alias, and iSCSI services.

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 135

[Guidelines for using CHAP authentication](#) on page 136

[What CHAP authentication is](#) on page 136

Enabling iSCSI session error recovery levels 1 and 2

By default, Data ONTAP is configured to use only 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 specify the maximum error recovery level allowed by the storage system.

About this task

There might be a minor performance reduction for sessions running error recovery level 1 or 2.

The `max_error_recovery_level` parameter specifies the maximum error recovery level allowed by the storage system. You can specify 0, 1, or 2, or you can accept the default value.

This parameter specifies the maximum error recovery level supported by the storage system. The initiator and storage system negotiate the actual error recovery level used for a session when the session is created; this is the smaller of the initiator's maximum and the storage system's maximum.

Note: The modified session error recovery level affects only the newly created sessions and does not affect existing sessions.

Steps

1. Use the `set -privilege advanced` command to enter the advanced mode.

Example

```
set -privilege advanced
```

2. Use the `iscsi show` command to verify the current setting.

Example

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

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

3. Use the `iscsi modify` command if you need to change the error recovery level.

Example

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

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
```

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 131
- [Commands for managing storage system interfaces on iSCSI](#) on page 132
- [Commands for managing iSCSI interface access management](#) on page 133
- [Commands for managing iSCSI initiator security](#) on page 133
- [Commands for managing iSCSI sessions](#) on page 133
- [Commands for managing iSCSI connections](#) on page 134
- [Commands for managing initiators](#) on page 134

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	<p><code>network interface delete</code></p> <p>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.</p>
Display a node name	<code>vserver iscsi show</code>
Display the target alias	<code>vserver iscsi show</code>
Change the target node name	<p><code>vserver iscsi modify</code></p> <p>Note: You need to change the privilege to advanced mode to change the target node name.</p> <p>Note: You must stop the iSCSI service before you can change the node name.</p>
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 <i>command</i></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 <i>command</i></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 <i>command</i></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>
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>

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 178

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 136

iSCSI interface access management

By default, all initiators have access to all interfaces.

Access lists are useful in a number of ways:

- **Performance:** In some cases, you might achieve better performance by limiting the number of interfaces an initiator can access.
- **Security:** You can gain better control over access to the interfaces.

Access lists must be explicitly defined. When an initiator begins a discovery session using an iSCSI `SendTargets` command, it receives those IP addresses associated with network interfaces on its access list.

iSNS server registration

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 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.

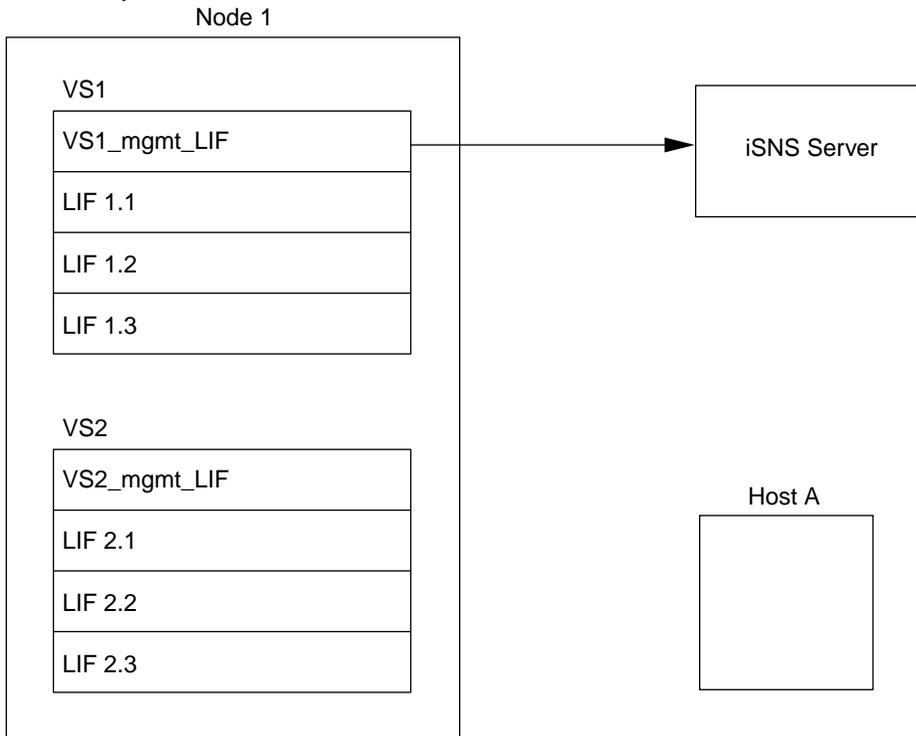
NetApp does not supply or resell iSNS servers. You obtain these servers from a vendor supported by NetApp. Be sure to check the NetApp iSCSI Support Matrix to see which iSNS servers are currently supported.

How SVMs interact with an iSNS server

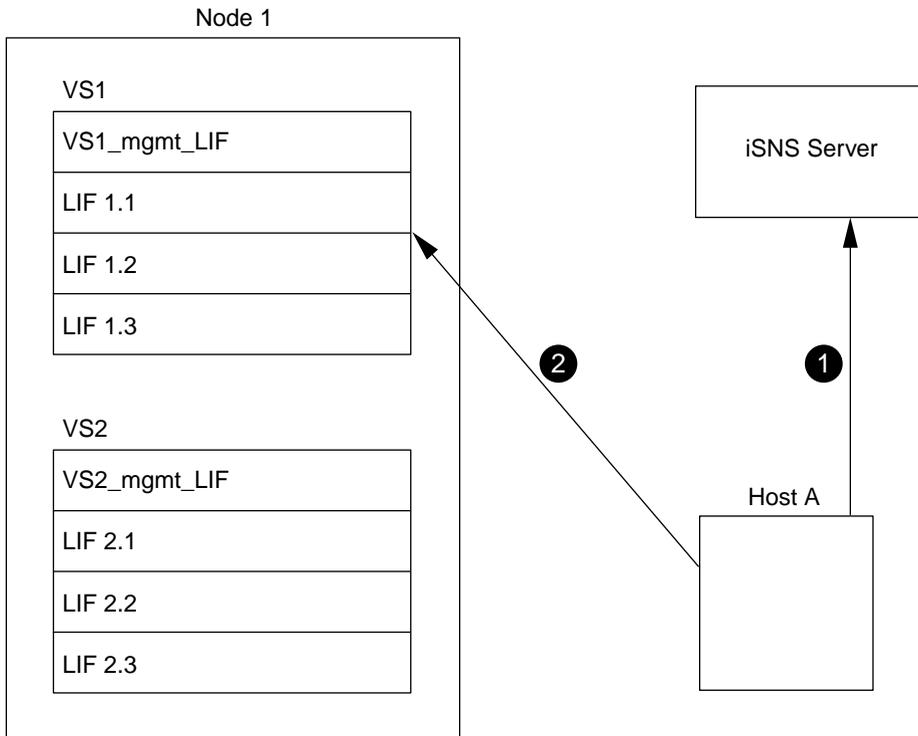
The iSNS server communicates with each Storage Virtual Machine (SVM) through the SVM management LIF. The management LIF registers all iSCSI target node name, alias, and portal information with the iSNS service for a specific SVM.

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 an SVM with an iSNS server](#) on page 141

Related references

[Commands for managing iSNS](#) on page 143

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 an SVM with an iSNS server

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

About this task

The `vserver iscsi isns create` command only 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 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. On the storage system console, use the `network interface create` command to create an SVM management LIF on a data port:

```
network interface create -vserver vs1 -lif vs1_mgmt_lif -role data -
data-protocol none -home-node innovate-01 -home-port e0c -address
192.168.3.72 -netmask 255.255.255.0
```

The SVM management LIF `vs1_mgmt_lif` enables you to register all LIFs for `vs1`. For more information about SVM management LIFs, see *Clustered Data ONTAP System Administration Guide for Cluster Administrators*.

3. Use the `vserver iscsi create` command to create an iSCSI service on your SVM if one does not already exist:

```
vserver iscsi create -vserver vs1
```

4. Use the `vserver iscsi show` command to verify that you created a iSCSI service:

```
iscsi show -vserver vs1
```

```
Vserver: vs1
      Target Name: iqn.
1992-08.com.companyX:sn.a416ff2821a123478563412:vs.3
      Target Alias: vs1
      Administrative Status: up
```

- Verify that a default route exists for the SVM with the `network routing-groups route show` command:

```
cluster1 ::> network routing-groups route show -vserver vs1
```

Vserver	Routing Group	Destination	Gateway	Metric
vs1	dl10.63.0.0/18	0.0.0.0/0	10.63.0.1	20

- If a default route does not exist for the SVM, use the `network routing-groups route create` command to create a default route:

```
network routing-groups route create -vserver vs1 -routing-group dl10.63.0.0/18 -destination 0.0.0.0/0 -gateway 10.63.0.1
```

- Use the `vserver iscsi isns create` command to configure your SVM to register with the iSNS service:

```
vserver iscsi isns create -vserver vs1 -address 10.61.208.215
```

Note: 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.

- Use the `vserver iscsi isns show` command to ensure that the iSNS service is running:

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

```
cluster1::> vserver iscsi isns show -vserver vs1

Vserver Name: vs1
iSNS Server IP Address: 10.61.208.215
Administrative Status: up
iSNS Server Entity Id: isns:00000003
Last Successful Update: 3/1/2012 19:32:11
Last Update Attempt: 3/1/2012 19:32:11
Last Update Result: iSNS_Ok
```

Note: You must configure iSCSI LIFs on your SVM to serve data. Without iSCSI LIFs, iSNS registration does not occur.

- If the iSNS service is not running, use the `vserver iscsi isns start` command:

```
cluster1::> vserver iscsi isns start -vserver vs1
```

Related concepts

[How SVMs interact with an iSNS server](#) on page 138

[What iSNS is](#) on page 138

Related references

[Commands for managing iSNS](#) on page 143

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 138

[What iSNS is](#) on page 138

Related tasks

[Registering an SVM with an iSNS server](#) on page 141

iSCSI troubleshooting tips

You can troubleshoot common problems that occur with iSCSI networks.

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	You should verify that the cables between the host and the storage system are properly connected.
Network connectivity	<p>You should verify that there is TCP/IP connectivity between the host and the storage system.</p> <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	You should 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 in the support.netapp.com/matrix .
Jumbo frames	If you are using jumbo frames in your configuration, you must ensure 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	You should verify that the iSCSI service is licensed and started on the storage system.
Initiator login	You should verify that the initiator is logged in to the storage system. If the command output shows no initiators are logged in, you should check the initiator configuration on the host. You should also verify that the storage system is configured as a target of the initiator.
iSCSI node names	You should 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.

Configuration setting	What to do
LUN mappings	You should 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 mapped show</code> Displays all LUNs and the igroups to which they are mapped. <code>lun mapped show -igroup</code> Displays the LUNs mapped to a specific igroup.
iSCSI LIFs enable	You should check if iSCSI logical interfaces are enabled or disabled.
port set configuration	You should verify whether there is a port set bound to the igroup for mapped LUN. If there is, then verify that the initiator is connecting on a LIF in that port set.

Related concepts

[Managing your iSCSI service](#) on page 129

[Understanding igroup functions](#) on page 123

[Understanding LUN functions](#) on page 57

Related information

[Interoperability Matrix: support.netapp.com/NOW/products/interoperability](http://support.netapp.com/NOW/products/interoperability)

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.

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

Message	Explanation	What to do
<code>ISCSI: network interface identifier disabled for use; incoming connection discarded</code>	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>

Message	Explanation	What to do
ISCSI: Authentication failed for initiator <i>nodename</i>	CHAP is not configured correctly for the specified initiator.	Check CHAP settings. <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. • You cannot use the same user name and password for inbound and outbound settings on the storage system.

Related concepts

Guidelines for using CHAP authentication on page 136

Understanding FC and FCoE SAN functions

To manage your SAN FC and FCoE environments, you must understand some conceptual information about Ethernet network management, Unified Ethernet network management, WWPNs, and WWNNs.

The information that you need to understand includes the following:

- How FC switches are identified
- How worldwide name assignments work
- What Unified Ethernet network management is

Related concepts

How Data ONTAP implements an FC SAN on page 181

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.

Brocade Fibre Channel switch
WWNN: 10:00:00:60:69:51:06:b4

Port numbers:															
0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
<input type="checkbox"/>															

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 182

How worldwide name assignments work

Data ONTAP assigns WWPNs and target names on demand.

While each adapter has a pre-configured WWPN and WWNN, Data ONTAP does not use these pre-configured values. Data ONTAP assigns its own worldwide names, WWPNs or WWNNs, when an FC logical interface is created. You can change this value only after creation of the FC logical interface with the `vserver fcp portname set` command.

The basis of the worldwide names depends on the MAC addresses of the onboard Ethernet ports. If you want to remove a node from a cluster and you have created FC protocol services or data logical interfaces, then you must contact technical support.

Fibre Channel over Ethernet overview

Fibre Channel over Ethernet (FCoE) is a model for connecting hosts to storage systems. As with Fibre Channel (FC), FCoE maintains existing FC management and controls. However, the hardware transport is a lossless 10-Gb Ethernet network.

Setting up an FCoE connection on the host or storage requires one or more supported converged network adapters (CNAs) connected to a supported FCoE switch. The CNA is a consolidation point and effectively serves as both an FC HBA and an Ethernet adapter.

The CNA is presented to the host and target as both an FCoE Initiator HBA and a 10-Gb Ethernet adapter. The FCoE Initiator HBA portion of the CNA handles the FCoE traffic when traffic is sent and received as FC frames mapped into Ethernet packets (FC over Ethernet). The Ethernet adapter portion of the CNA handles the standard Ethernet IP traffic, such as iSCSI, CIFS, NFS, and HTTP, for the host. Both the FCoE and standard Ethernet portions of the CNA communicate over the same Ethernet port, which connects to the FCoE switch.

The FCoE target adapter is also sometimes called a "unified target adapter" or UTA. Like the CNA, the UTA supports both FCoE and regular Ethernet traffic.

You should configure jumbo frames (MTU = 9000) for the Ethernet adapter portion of the CNA. You cannot change the MTU for the FCoE portion of the adapter.

Note: Unified target adapters (UTAs) are 10-Gb converged network adapters (CNAs) that you install in your storage systems.

In general, you configure and use FCoE connections just like traditional FC connections. You can use UTAs for non-FCoE IP traffic such as NFS, CIFS, or iSCSI.

Note: For detailed information about how to set up and configure your host to run FCoE, see your host documentation.

Unified Ethernet network management

A unified Ethernet network runs data and storage traffic, including iSCSI, CIFS, NFS, and Fibre Channel, over your existing Ethernet infrastructure.

Unified Target Adapters (UTAs) and Converged Network Adapters (CNAs) are network adapters that support IP traffic and FCoE traffic on the same 10-Gb Ethernet link. These adapters may also support fibre channel traffic.

Note: UTAs and CNAs are configured and managed just like any other FC or Ethernet port; there are no unique configuration commands. See the *Clustered Data ONTAP File Access and Protocols Management Guide* for information about managing file system protocols.

In addition to the hardware components, Data ONTAP also supports the Data Center Bridging Exchange (DCBX) protocol, which is required for negotiating operating parameters that control transfers of both FC and Ethernet traffic over the Ethernet infrastructure.

Related concepts

[Managing your iSCSI service](#) on page 129

[Managing systems with Fibre Channel adapters](#) on page 152

Related information

[Technical Report: Fibre Channel over Ethernet \(FCoE\) End-to-End Deployment Guide - media.netapp.com/documents/TR-3800.pdf](#)

[Interoperability Matrix: support.netapp.com/matrix](#)

Data center bridging

Data center bridging (DCB) is a collection of extensions to the existing Ethernet standard that provides a lossless transport layer for FCoE, iSCSI, and other NIC traffic sharing the same physical CNA interface. DCB assigns priority and allocates bandwidth to network traffic based on the traffic protocol.

The DCB standard resolves packet loss issues by implementing the following technologies:

Per-priority pause (priority-based flow control)	Enables a device to only inhibit the transmission of frames based on user-defined priorities.
Enhanced transmission selection	Allows administrators to allocate bandwidth on a percentage basis to different priorities.
Congestion notification	Transmits congestion information.
DCB Exchange (DCBX) protocol	Exchanges connection information with directly connected peers and detects misconfigurations.

Although these technologies possess their own independent functions, they operate together to provide an enhanced Ethernet standard that eliminates packet loss due to traffic congestion. For more information about FCoE deployment, see TR-3800.

Note: iSCSI DCB support is only available on the QLogic 8300 series dual port 10GbE CNA adapter.

Related information

Technical Report: Fibre Channel over Ethernet (FCoE) End-to-End Deployment Guide: media.netapp.com/documents/TR-3800.pdf

Data Center Bridging task group: www.ieee802.org/1/pages/dcbridges.html

Documentation on the NetApp Support Site: support.netapp.com

Support for iSCSI DCB

iSCSI data Center bridging (DCB) can assign priority and allocate bandwidth to network traffic based on the traffic protocol. Beginning in Data ONTAP 8.2.1, DCB is supported in iSCSI. DCB can be enabled on networks for FCoE, iSCSI, and other NIC traffic sharing the same physical CNA interface

You can display DCB settings through Data ONTAP, but you must enable and configure your DCB settings through your switch.

Note: iSCSI DCB support is only available on the QLogic 8300 series dual port 10GbE CNA adapter.

Verifying that DCB is configured correctly

If you are using FCoE, you must verify that DCB bridging is configured correctly and that the target adapter has successfully logged into the fabric.

About this task

If FCoE is not configured, you need to consult the switch documentation.

Note: These settings are configured at the switch level, and the storage system discovers and displays these pre-configured settings.

Step

1. Use the `dcb show` command to display the DCB settings for each node in the cluster.

Example

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

Interface	PGID	Priority	Application	Bandwidth
e3a	0	0 1 2 4 5 6 7	unassigned	50%
	1	3	FCoE	50%
e3b	0	0 1 2 3 4 5 6 7	unassigned	0%
e4a	0	0 1 2 4 5 6 7	unassigned	50%
	1	3	FCoE	50%
e4b	0	0 1 2 3 4 5 6 7	unassigned	0%

Related concepts

Unified Ethernet network management on page 149

Fibre Channel over Ethernet overview on page 148

Data center bridging on page 149

Related information

Brocade switch FCoE set up procedures: <http://www.brocade.com/products/all/index.page>

Cisco switch FCoE set up procedures: http://www.cisco.com/en/US/products/ps9670/products_installation_and_configuration_guides_list.html

Managing systems with Fibre Channel adapters

Most systems have onboard FC adapters that you can configure as initiators or targets. You can also use certain FC adapter cards to configure as initiators or targets. Initiators connect to back-end disk shelves, and targets connect to FC switches or other storage controllers.

You should follow the instructions in this section to configure your onboard FC adapters as initiators or targets.

For additional configuration details, see the *Clustered Data ONTAP SAN Configuration Guide*.

Related information

Configuration and hardware guides on the NetApp Support Site: support.netapp.com

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.

Before you begin

You must first move or delete all LIFs that use this adapter as their home port.

About this task

Each onboard FC port can be individually configured as an initiator or a target. Traditionally, ports on FC adapter cards were either initiators or targets, and you could not change the type.

However, each port on the following cards can be individually configured as either a target port or an initiator FC port, just like the onboard FC ports:

- The X2056-R6 4-port 8-Gb FC adapter for the 62xx vertical I/O slots (slots 1, 11, and 12).
- The X1132A-R6 4-port 8-Gb FC adapter for the 62xx and 32xx.
- The X1143A-R6 dual port 16-Gb FC, 10-GbE FCoE adapter for the 62xx and 32xx.

For detailed information about the number of target adapters supported on each hardware platform, see the *Clustered Data ONTAP SAN Configuration Guide*.

Steps

1. Use the `network interface delete` command to remove all the LIFs from the adapter.

Example

```
network interface delete -vserver vs3 -lif lif2,lif0
```

Note: If the LIF is in a port set, you must remove the LIF from port set before you can delete the LIF.

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

2. Use the `network fcp adapter modify` command to take your adapter offline.

Example

```
network fcp adapter modify -node innovate-01 -adapter 0c -state down
```

3. Use the `system hardware unified-connect modify` command to change the adapter from target to initiator.

Example

```
system hardware unified-connect modify -t initiator 0c
```

4. Reboot the node hosting the adapter you changed.
5. Use the `system hardware unified-connect show` command to verify that the FC ports are configured in the correct state for your configuration.

Example

```
ssan-3240-25-a::> system hardware unified-connect show
```

Node	Current Adapter	Current Mode	Pending Mode	Pending Type	Admin Status
ssan-3240-25-a	0c	fc	initiator	-	online
ssan-3240-25-a	0d	fc	initiator	-	offline
ssan-3240-25-a	2a	fc	target	-	online
ssan-3240-25-a	2b	fc	target	-	online
ssan-3240-25-a	2c	fc	target	-	online
ssan-3240-25-a	2d	fc	target	-	online
ssan-3240-25-a	3a	cna	target	-	online
ssan-3240-25-a	3b	cna	target	-	online

6. Use the `storage enable adapter` command to bring the adapter back online.

Example

```
node run -node node1 storage enable adapter -e 0c
```

Related tasks

[Deleting a LIF in a SAN environment](#) on page 121

Related information

Configuration and hardware guides: <https://support.netapp.com/documentation/productsatoz/index.html>

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. Traditionally, ports on FC adapter cards were either initiators or targets, and you could not change the mode.

However, each port on the following cards can be individually configured as either a target port or an initiator FC port, just like the onboard FC ports:

- The X2056-R6 4-port 8-Gb FC adapter for the 62xx vertical I/O slots (slots 1, 11, and 12).
- The X1132A-R6 4-port 8-Gb FC adapter for the 62xx and 32xx.
- The X1143A-R6 dual port 16-Gb FC, 10-GbE FCoE adapter for the 62xx and 32xx.

For detailed information about the number of target adapters supported on each hardware platform, see the *Clustered Data ONTAP SAN Configuration Guide*.

Steps

1. Use the `storage disable adapter` command to take the adapter offline.

Example

```
node run -node node1 storage disable adapter -d 0c
```

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

2. Use the `system hardware unified-connect modify` command to change the adapter from initiator to target.

Example

```
system hardware unified-connect modify -t target 0c
```

3. Reboot the node hosting the adapter you changed.
4. Use the `system hardware unified-connect show` command to verify that the FC ports are configured in the correct state for your configuration.

Example

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

Node	Current Adapter		Current Mode	Pending Mode	Pending Type	Admin Status
ssan-3240-25-a	0c	fc	initiator	-	-	online
ssan-3240-25-a	0d	fc	initiator	-	-	offline
ssan-3240-25-a	2a	fc	target	-	-	online
ssan-3240-25-a	2b	fc	target	-	-	online
ssan-3240-25-a	2c	fc	target	-	-	online
ssan-3240-25-a	2d	fc	target	-	-	online
ssan-3240-25-a	3a	cna	target	-	-	online
ssan-3240-25-a	3b	cna	target	-	-	online

5. Use the `network fcp adapter` command to verify that the target port has the correct connection type (ptp or loop) for your configuration.

Example

```
network fcp adapter show -node node1
```

Connection Node	Host Adapter	Established	Port Address
node1	0a	ptp	10c00
node1	0c	ptp	10d00
node2	0a	ptp	10e00
node2	0c	ptp	10f00

6. Use the `network fcp adapter modify` command to bring your adapter online.

Example

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

Related information

Configuration and hardware guides: <https://support.netapp.com/documentation/productsatoz/index.html>

Displaying the status of onboard FC adapter

Use the `system hardware unified-connect show` command to determine the configured and operational status of the FC onboard adapters.

About this task

This command also displays other information, including the configuration status of the adapter and whether it is configured as a target or initiator.

Note: Onboard FC adapters are set to initiator mode by default.

Step

1. Use the `system hardware unified-connect show` command to determine the status of the FC onboard adapters.

Example

```
node run -node node1 system hardware unified-connect show
```

Node	Current Adapter	Current Mode	Pending Mode	Pending Type	Admin Status
ssan-3240-25-a	0c	fc	initiator	-	online
ssan-3240-25-a	0d	fc	initiator	-	offline
ssan-3240-25-a	2a	fc	target	-	online
ssan-3240-25-a	2b	fc	target	-	online
ssan-3240-25-a	2c	fc	target	-	online
ssan-3240-25-a	2d	fc	target	-	online
ssan-3240-25-a	3a	cna	target	-	online
ssan-3240-25-a	3b	cna	target	-	online

Note: The available Local State values are CONFIGURED, PENDING, and UNCONFIGURED. See the `system hardware unified-connect man` page for a detailed description of each value.

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. Use the `network fcp adapter show` command.

Example

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

```

      Node: nodel
      Adapter: 3a
      Description: Fibre Channel Target Adapter 3a
(QLogic 2532 (2562), rev. 2,
      8G)
      Physical Protocol: fibre-channel
      Maximum Speed: 8
      Status: online
      Substatus: ADAPTER UP
      Host Port Address: 4f2000
      Firmware Revision: 5.8.0
      Data Link Rate (Gbit): 8
      Fabric Established: true
      Fabric Name: 10:00:00:05:33:ed:95:11
      Connection Established: ptp
      Mediatype: ptp
      Configured Speed: auto
      Configured State: up
      Adapter WWNN: 50:0a:09:80:8d:e1:cf:86
      Adapter WWPN: 50:0a:09:83:8d:e1:cf:86
      Switch Port: sdev-broc-gg28a:32
      Form Factor Of Transceiver: SFP
      Vendor Name Of Transceiver: FINISAR CORP.
      Part Number Of Transceiver: FTLF8528P2BCV-QL
      Revision Of Transceiver: A
      Serial Number Of Transceiver: PKL2CTF
      FC Capabilities Of Transceiver: 2,4,8 (Gbit/sec)
      Vendor OUI Of Transceiver: 0:144:101
      Wavelength In Nanometers: 850
      Date Code Of Transceiver: 11:05:11
      Validity Of Transceiver: true
      Connector Used: LC
      Encoding Used: 8B10B
      Is Internally Calibrated: true
      Received Optical Power: 489.0 (uWatts)
      Is Received Power In Range: true
      SPF Transmitted Optical Power: 478.2 (uWatts)
      Is Xmit Power In Range: true

```

System configuration information and adapter information for each slot that is used are displayed on the screen.

Changing the FC adapter speed

You can use the `network fcp adapter modify` command to change the FC adapter speed to match the fabric to which the adapter is connecting. Some fabrics do not support auto negotiation.

Before you begin

Before you can change the speed of an adapter, all LIFs that use this adapter as their home port must be offline.

About this task

The available speeds are dependent on the HBA being used. See the man page for the list of available adapter speeds.

Note: 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, this operation applies to all LIFs in the cluster, which might not be desirable.

Steps

1. Use the `network interface modify` command to take all the LIFs on this adapter offline.

Example

```
network interface modify -vserver * -lif * -home-node node1 -home-port
0c -status-admin down
```

2. Use the `network fcp adapter modify` command to take your adapter offline.

Example

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

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

3. Use the `fcp adapter show` command to determine the maximum speed for the port adapter. You cannot modify the adapter speed beyond the maximum speed.

Example

```
fcp adapter show -instance
```

```
Node: node1
Adapter: 0a
```

```

Description: Fibre Channel Target Adapter 0a (rev. 2)
Physical Protocol: fibre-channel
  Maximum Speed: 4
    Status: online
    Substatus: ADAPTER DOWN
Host Port Address: 10c00
Firmware Revision: 5.5.2
  PCI Bus Width: 64
  PCI Clock Speed: 33
  Cacheline Size: 16
  FC Packet Size: 2048
    SRAM Parity: true
    External GBIC: false
Data Link Rate (Gbit): 4
  Fabric Established: true
Connection Established: ptp
  Mediatype: ptp
  Configured Speed: auto
  Configured State: down

```

4. Use the network `fcp adapter modify` command to change the adapter speed.

Example

```
network fcp adapter modify -node node1 -adapter 0a -speed 8
```

5. Use the network `fcp adapter modify` command to bring your adapter back online.

Example

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

6. Use the network `interface modify` command to bring all the LIFs on this adapter online.

Example

```
network interface modify -vserver * -lif * -home-node node1 -home-port
0c -status-admin up
```

Commands for managing FC adapters

You can use FC commands to manage FC adapters on your storage controller.

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

- [Commands for managing FCP target adapters](#) on page 159
- [Commands for managing FCP initiator adapters](#) on page 160

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 target adapter statistics	<code>fcp stats</code>
Display FC protocol traffic information	<code>run -node <i>node_name</i> sysstat -f</code>
Display how long 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>
Display configuration information that affects core dumps and takeovers	<code>run -node <i>node_name</i> sysconfig -c</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>
Display configuration information that affects core dumps and takeovers	<code>run -node <i>node_name</i> sysconfig -c</code>

Managing systems with the X1143A-R6 adapter

Most systems have onboard FC adapters that you can configure as initiators or targets. You can also use the X1143A-R6 adapter which can be configured for FC mode or CNA mode.

Function of the X1143A-R6 adapter

The X1143A-R6 network adapter is a new generation adapter with ports that can be configured as either 10 Gb Ethernet and FCoE ports or as 16Gb FC initiator or target ports.

Data ONTAP has supported FC and converged network adapter (CNA) modes in previous releases, but used different add-on host bus adapters (HBAs) and onboard devices for each mode. The X1143A-R6 adapter supports both FC and CNA modes on the same device using dual 16-Gb FC and 10-GbE/FCoE ports. The X1143A-R6 is the first device that can be configured into either one of the personality modes within the same hardware.

In FC mode, the X1143A-R6 device behaves just like any existing FC device with speeds up to 16Gbps. In CNA mode, you can use the X1143A-R6 device for concurrent NIC and FCoE traffic sharing the same 10 GbE port.

Data ONTAP discovers different number of PCI functions per chip depending on its mode. Each PCI function is an independent PCI device with its own PCI configuration space. For example, under FC mode, only one PCI function of FC will be available for each port, and the physical port allows the use of either 16G FC SFP+, 8G FC SFP+, or 4G FC SFP+. While under CNA mode, two PCI functions will be available for each port. One is the NIC function while the other is the FCoE function and the physical port only allows the use of 10G FC SFP+. Data ONTAP drivers identify and bind to PCI functions based on PCI Vendor_ID and PCI Device_ID combinations from their PCI configuration space. The X1143A-R6 device will assign unique but different PCI Device_IDs for different types of PCI functions, such as FC, NIC and FCoE.

In addition, from a Data ONTAP device driver perspective, once a PCI function is discovered, a device driver will be attached to use the X1143A-R6. For the NIC function, the NIC driver will be attached. In the case of the FC or FCoE function, a different driver will be used depending on its type. The X1143A-R6 supports the following types in Data ONTAP:

- FCP Initiator
- FCP target

Note: CNA mode only supports FCP target for the FCoE function.

Data ONTAP can attach the proper driver to the X1143A-R6 device during boot-up and initialize the device for FCP target. The Unified Connect Management (UCM) framework provides the infrastructure to manage these personality modes and FC-4 types, enabling you to change

configurations. However UCM itself does not require any one particular configuration default for personality mode and/or FC-4 type. Instead it is capable of working with any default value that is set.

Managing the X1143A-R6 adapter

After the port mode has been configured to CNA mode, you can manage the unified target adapter (X1143A-R6) using the `system hardware unified-connect show` or the `system hardware unified-connect modify` command. You can view your settings, change modules, modify your port settings, and set up the appropriate connections.

Steps

1. Ensure that you have selected the CNA mode and view your current configuration.
2. Modify your port settings.

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.

Changing the optical modules

Depending upon the mode you require, you can switch to Fibre Channel (FC) or Converged Network Adapter (CNA). Refer to configuration information for configuration settings for the X1143A-R6.

Steps

1. Remove the current optical modules from the X1143A-R6 adapter card.
2. Insert the correct modules for your preferred personality mode (FC or CNA) optics.
3. Run the `network fcp adapter show -instance -node -adapter` command to ensure that you have the correct SFP+ installed.

The list of installed optics or Copper cables is displayed.

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

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

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

Default configuration for the X1143A-R6

You must connect to the X1143A-R6 ports using LC fiber optic cables with supported SFP+ optical modules, or by using copper Twinax cables (in 10GbE mode only). You can monitor the activity on the ports through the LEDs on the bezel.

The X1143A-R6 device can be configured as Fibre Channel (FC) or Converged Network Adapter (CNA). By default, the X1143A-R6 ships configured in FC target mode. You must use the `system hardware unified-connect modify` command to change the operational mode (FC or Ethernet) of the card. See the man page for details.

Note: Both ports that are connected to the same chip must operate in the same mode (FC or CNA).

Data ONTAP provides a mechanism to manage these configurations, which enables you to reconfigure the device into different modes. For the FC type, there is an existing command in Data ONTAP, `system hardware unified-connect modify`, which can be used to manage FC types of ports between FCP initiator and FCP target.

Configuration rules and supported settings

The X1143A-R6 device supports configuration of two adjacent ports or all ports on the same chip in the same personality mode. Data ONTAP supports configuration of specific initiator or target types on a port.

The Data ONTAP framework manages both personality modes and FC types for Unified Connect adapters, including the UI. Data ONTAP not only manages the personality (FC/CNA) for all UC devices, but also their FC types. This framework extends management of initiator or target types of existing FC devices to provide the same functionality as `fcadmin config` does currently so that the transition is seamless. UCM provides unified configuration management for both UC devices and existing FC devices.

UCM also helps device drivers to determine device personality mode (initiator or target) during the device discovery phase of Data ONTAP so that correct device drivers will be attached.

Configuration resources

You can find the most up-to-date installation, cabling, and configuration information about this card and your platform online.

For supported platform models and slot assignment, see the *Hardware Universe* at hwu.netapp.com.

Specific information about installing this card in your platform can be found in the PCIe replacement flyer for your platform at support.netapp.com/documentation/productsatoz/index.html.

You can refer to the following documentation available at support.netapp.com when configuring the X1143A-R6 card for your specific release of Data ONTAP:

- *Clustered Data ONTAP Network Management Guide*
- *Clustered Data ONTAP SAN Administration Guide*
- *Clustered Data ONTAP SANConfiguration Guide*

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.

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

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 148

[How WWPNs are used](#) on page 182

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
```

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 166
- [Commands for displaying active FC initiators](#) on page 167
- [Commands for managing FC logical interfaces](#) on page 167
- [Commands for managing FC initiator WWPN aliases](#) on page 167

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>

If you want to....	Use this command...
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 <i>command</i> <i>_name</i></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 <i>command</i> <i>_name</i></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</i> <i>_name</i></code>

Commands for managing FC initiator WWPN aliases

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

Volume management

Data ONTAP allows you to manage your volumes through support for such capabilities as moving, deleting, and monitoring available space on your volumes.

See the *Clustered Data ONTAP Logical Storage Management Guide* for information on how to manage your volume space.

Before performing a volume move in SAN

You should have a minimum of 2 paths per LUN connecting to each node in your cluster. This eliminates single points of failure and enables the system to survive component failures.

You should do the following before performing a volume move:

- Make sure that at least 1 LIF with the iSCSI or FC protocol exists in the specified SVM on the volume move destination node.
- Make sure that at least 1 LIF with the iSCSI or FC protocol exists in the specified SVM on the volume move destination partner node.

Note: If you do not have the minimum number of recommended paths per LUN or at least 1 LIF with the iSCSI or FC protocol in the SVM on the destination node and destination partner node, then you might experience an additional delay before your LUNs are accessible after your volume move operation completes.

For more information about creating LIFs, see the *Clustered Data ONTAP Network Management Guide*.

Virtualization of storage

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.

In virtualized environments, you also need to ensure that your LUNs are aligned. Unaligned LUNs can degrade performance.

LUN alignment in virtual environments

LUN alignment problems, which can lead to lower performance for your storage system, are common in virtualized server environments. To avoid LUN alignment problems, it is essential to follow best practices for proper LUN alignment.

See the technical report TR 3747 for detailed guidelines and background information on provisioning storage in virtualized server environments.

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*

ESX boot LUNs

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 normally a performance problem because the total amount of misaligned I/O is small. Assuming the LUN was correctly provisioned with the `ostype` option value of `vmware`, no action is needed.

Related concepts

[Misaligned I/O can occur on properly aligned LUNs](#) on page 109

Related information

[Best Practices for File System Alignment in Virtual Environments: media.netapp.com/documents/tr-3747.pdf](http://media.netapp.com/documents/tr-3747.pdf)

[Documentation on the NetApp Support Site: support.netapp.com](http://support.netapp.com)

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, there are significant performance benefits 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

ODX can only be utilized if your system meets certain requirements. ODX can also only be used with certain volume types for intra-cluster copies.

Note: 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:

- 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 on 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
- FlexCache volumes
- Source volumes that are not deduplicated
- Compressed volumes
- Sparse volumes
- Read-only 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. The copy is performed by using FlexClone file technology, which provides additional remote copy performance benefits.
- **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.

How hosts connect to storage systems

Hosts can connect to block storage 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 Fibre Channel HBAs or CNAs.

Related information

Host Utilities documentation: www.support.netapp.com

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

Interoperability Matrix: support.netapp.com/NOW/products/interoperability

Host Utilities documentation: www.support.netapp.com

About SnapDrive for Windows and UNIX

SnapDrive software is an optional management package for Microsoft Windows and UNIX hosts. SnapDrive can simplify some of the management and data protection tasks associated with iSCSI and FC storage.

SnapDrive for Windows is a server-based software solution that provides advanced storage virtualization and management capabilities for Microsoft Windows environments. It is tightly integrated with Microsoft NTFS and provides a layer of abstraction between application data and physical storage associated with that data. SnapDrive runs on Windows Server hosts and complements native NTFS volume management with virtualization capabilities. It enables administrators to easily create virtual disks from pools of storage that can be distributed among several storage systems.

SnapDrive for UNIX provides simplified storage management, reduces operational costs, and improves storage management efficiency. It automates storage provisioning tasks and simplifies the process of creating Snapshot copies and clones from Snapshot copies consistent with host data.

Related information

Host Utilities documentation: www.support.netapp.com

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 135

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 iSCSI works in a Storage Virtual Machine (SVM)

An iSCSI SVM provides high availability through MPIO and ALUA. When a node fails, traffic is rerouted to the other nodes within SVM.

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.

Related concepts

[Understanding FC and FCoE SAN functions](#) on page 147

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.

Related concepts

[Understanding FC and FCoE SAN functions](#) on page 147

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

[How to limit LUN access in a virtualized environment](#) on page 63

Copyright information

Copyright © 1994–2014 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, Network Appliance, the Network Appliance logo, Akorri, ApplianceWatch, ASUP, AutoSupport, BalancePoint, BalancePoint Predictor, Bycast, Campaign Express, ComplianceClock, Customer Fitness, Cryptainer, CryptoShred, CyberSnap, Data Center Fitness, Data ONTAP, DataFabric, DataFort, Decru, Decru DataFort, DenseStak, Engenio, Engenio logo, E-Stack, ExpressPod, FAServer, FastStak, FilerView, Fitness, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexSuite, FlexVol, FPolicy, GetSuccessful, gFiler, Go further, faster, Imagine Virtually Anything, Lifetime Key Management, LockVault, Manage ONTAP, Mars, MetroCluster, MultiStore, NearStore, NetCache, NOW (NetApp on the Web), Onaro, OnCommand, ONTAPI, OpenKey, PerformanceStak, RAID-DP, ReplicatorX, SANscreen, SANshare, SANtricity, SecureAdmin, SecureShare, Select, Service Builder, Shadow Tape, Simplicity, Simulate ONTAP, SnapCopy, Snap Creator, SnapDirector, SnapDrive, SnapFilter, SnapIntegrator, SnapLock, SnapManager, SnapMigrator, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapSuite, SnapValidator, SnapVault, StorageGRID, StoreVault, the StoreVault logo, SyncMirror, Tech OnTap, The evolution of storage, Topio, VelocityStak, vFiler, VFM, Virtual File Manager, VPolicy, WAFL, Web Filer, and XBB are trademarks or registered trademarks of NetApp, Inc. in the United States, other countries, or both.

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. A complete and current list of other IBM trademarks is available on the web at www.ibm.com/legal/copytrade.shtml.

Apple is a registered trademark and QuickTime is a trademark of Apple, Inc. in the United States and/or other countries. Microsoft is a registered trademark and Windows Media is a trademark of Microsoft Corporation in the United States and/or other countries. RealAudio, RealNetworks, RealPlayer, RealSystem, RealText, and RealVideo are registered trademarks and RealMedia, RealProxy, and SureStream are trademarks of RealNetworks, 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.

NetApp, Inc. is a licensee of the CompactFlash and CF Logo trademarks.

NetApp, Inc. NetCache is certified RealSystem compatible.

How to send your comments

You can help us to improve the quality of our documentation by sending us your feedback.

Your feedback is important in helping us to provide the most accurate and high-quality information. If you have suggestions for improving this document, send us your comments by email to doccomments@netapp.com. To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

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

-space-mgmt-try-first option
for selecting first method of automatic FlexVol
volume space increases [42](#)

A

access lists

about [137](#)
commands [133](#)
iSNS [138](#)

active initiators

available commands [134](#)
displaying [134](#)
showing [134](#)

active volume

cloning [94](#)

adapters

active initiators [167](#)
changing [165](#)
changing speed [158](#)
configuring
initiator mode [152](#)
target mode [154](#)
displaying
configurations [159](#)
FC logical interfaces [167](#)
FC target [156](#)
targets [159](#)

FC initiator

displaying [160](#)

initiator status [160](#)

managing UTAs [162](#)

Unified Target [161](#)

WWPN [165](#)

ALUA

defined [125](#)

attaching

FC host [79](#)

iSCSI host [70](#)

authentication

displaying security information [133](#)

iSCSI [135](#)

auto negotiation [158](#)

autodelete

FlexClone LUN [94](#)

volume size [21](#)

autogrow

how Data ONTAP can add space for FlexVol
volumes automatically [41](#)

autosizing

how Data ONTAP can add space for FlexVol
volumes automatically [41](#)

availability

LUN [58](#)

available space

command effects [29](#)

B

backups

about configuring and using SnapVault, in SAN
environments [98](#)
accessing read-only LUN copies from SnapVault [99](#)
LUNs through host system [106](#)
restoring all LUNs in a volume from SnapVault [102](#)
restoring single LUNs from SnapVault [100](#)

best practices

storage provisioning [56](#)

block access

enabling [65](#), [70](#), [74](#), [79](#)

FC hosts [79](#)

iSCSI [65](#)

iSCSI hosts [70](#)

C

capacity [15](#)

CHAP

defined [136](#)

guidelines [136](#)

iSCSI authentication [135](#)

Cisco switches

FC and FCoE zoning requirement [122](#)

commands

FC

active initiators [167](#)

initiator adapter [160](#)

logical interfaces [167](#)

service [166](#)

target adapter [159](#)

WWPN alias [167](#)

- igroup [127](#)
- initiator [134](#)
- iSCSI
 - connections [134](#)
 - initiator security [133](#)
 - interface access [133](#)
 - security [129](#)
 - service [131](#)
 - sessions [133](#)
- iSNS service [143](#)
- LUN [116](#)
- port sets [117](#)
- storage system interfaces [132](#)
- communication
 - iSNS [138](#)
- COMPARE WRITE feature
 - VAAI feature [172](#)
- configuration limits
 - volume and SVM requirements [38](#)
- configuration options
 - volumes [11](#)
- configuring
 - thin provisioning [17](#)
 - unified target adapter ports [163](#)
 - unified target adapters [163](#)
- considerations
 - LUNs and files [55](#)
- copy offload
 - use cases for [175](#)
 - See also* ODX
- copy offloads
 - system requirements [174](#)
- creating
 - iSCSI initiator security method [133](#)
 - LUNs [116](#)
 - port sets [117](#)

D

- data center bridging
 - iSCSI support [150](#)
- Data ONTAP options
 - max_error_recovery_level [130](#)
- data protection
 - methods of [81](#)
 - Snapshot copies [84](#)
 - using FlexClone LUNs [90](#)
- DCB
 - FCoE switching [148](#)
 - iSCSI support [150](#)

- verifying [150](#)
- decreasing
 - LUN size [114](#)
- default security authentication
 - defining [133](#)
- defining
 - default security authentication [133](#)
 - security password [133](#)
- deleting
 - FlexClone LUNs automatically [96](#)
 - LUNs [114](#), [116](#)
 - port sets [117](#)
 - Snapshots automatically [40](#)
- destroying
 - LUNs [116](#)
 - port sets [117](#)
- disk space
 - managing [169](#)
- disk space usage
 - no Snapshot copies [29](#)
 - showing [29](#), [32](#)
 - Snapshot copies [32](#)
- displaying
 - authentication security information [133](#)
 - igroup [117](#)
 - LUN mapping information [116](#)
 - onboard FC adapters [156](#)
 - port sets [117](#)

E

- error recovery levels
 - enabling [130](#)
- ethernet [147](#)
- Ethernet
 - 10-Gb adapters [149](#)
- eui type designator [180](#)
- examples
 - igroup mapping [63](#)
 - optimized path [60](#)
 - port sets [63](#)
 - SVMs [60](#)
 - thin provisioning [15](#)
- extended copy feature
 - environment [172](#)
 - invoked automatically [172](#)
 - VAAI feature [172](#)
 - when the standard copy operation is used [172](#)

- F**
- FC
 - active initiators
 - commands [167](#)
 - displaying [167](#)
 - adapter initiator
 - configuring [152](#)
 - adapter speed
 - changing [158](#)
 - adapters
 - displaying [156](#)
 - adding [166](#), [167](#)
 - adding target alias [166](#)
 - assigning [167](#)
 - available commands [160](#)
 - block access [74](#)
 - commands [159](#)
 - configuring [154](#)
 - core dumps [159](#), [160](#)
 - deleting [167](#)
 - displaying
 - target adapters [159](#)
 - displaying target alias [166](#)
 - functions [147](#)
 - initiator adapter [160](#)
 - initiator adapters [160](#)
 - license [65](#), [73](#)
 - LIF [166](#)
 - logical interfaces [167](#)
 - managing systems with onboard adapters [152](#)
 - onboard adapters
 - displaying [156](#)
 - takeovers [159](#), [160](#)
 - target adapter [159](#)
 - target adapters
 - displaying [156](#)
 - modifying [159](#)
 - statics [159](#)
 - traffic information [159](#)
 - WWPN alias [167](#)
- FC commands [167](#)
- FC license
 - adding [166](#)
 - disabling [166](#)
 - enabling [166](#)
 - verifying [166](#)
- FC LIF
 - zoning restrictions for Cisco switches [122](#)
- FC logical interfaces
 - available commands [167](#)
- FC service
 - commands [166](#)
 - deleting [166](#)
 - modifying [166](#)
 - starting [166](#)
 - stopping [166](#)
- FC target
 - creating [166](#)
- FC target alias
 - adding [166](#)
- FC WWPN alias
 - available commands [167](#)
 - creating [167](#)
 - deleting [167](#)
 - destroying [167](#)
 - displaying [167](#)
 - modifying [167](#)
- FCoE
 - configure [72](#)
 - data center bridging [149](#)
 - DCB support [150](#)
 - functions [147](#)
 - switches
 - configuring FCoE [72](#)
 - target adapters [149](#)
 - Unified Ethernet
 - management [147](#)
- FCoE LIF
 - zoning restrictions for Cisco switches [122](#)
- FCP
 - defined [181](#)
 - node connection [182](#)
 - node identification [182](#)
 - nodes defined [182](#)
 - switch nodes [147](#)
- FCP commands
 - fcp adapter show [159](#)
 - fcp nodename [159](#)
 - fcp show [159](#)
 - fcp stats [159](#)
 - storage show adapter [160](#)
 - sysconfig [160](#)
- FCP service
 - destroying [166](#)
- FlexClone LUN
 - autodelete
 - FlexClone LUN [91](#)
- FlexClone LUNs
 - autodelete [96](#)

- enabling automatic deletion of [92](#)
 - for data protection [90](#)
 - reasons for using [90](#)
- FlexVol volumes
 - considerations for setting fractional reserve [39](#)
 - how Data ONTAP can automatically add space for [41](#)
 - how they reclaim free space from FlexClone LUNs [91](#)
 - required for SAN [38](#)
 - selecting first method for automatic space increases [42](#)
- fractional reserve
 - considerations for setting [39](#)
- free space
 - how Data ONTAP can increase automatically for FlexVol volumes [41](#)
 - selecting first method to automatically increase FlexVol volume [42](#)

G

- geometry
 - LUN [58](#)
- guidelines
 - CHAP authentication [136](#)
 - LUN layout [57](#)
 - LUN mapping [56](#)
 - provisioning [56](#)
 - space allocation [57](#)
 - thin provisioning [17](#)

H

- HBA [123](#), [135](#), [148](#), [177](#), [178](#)
- head swap
 - changing WWPNs [165](#)
- high availability [126](#)
- host
 - configuring [70](#), [79](#)
 - FC [79](#)
 - initiator
 - node name [125](#)
 - iSCSI [70](#)
 - iSCSI implementation [135](#)
 - iSNS [138](#)
 - LUN usable space [58](#)
 - storage system connection [177](#)
 - supported configurations [55](#)
 - tape backup [105](#)

- Host Utilities
 - defined [177](#)
- hosts
 - backups [106](#)
 - moving LIF [119](#)

I

- I/O
 - alignment considerations for Linux [111](#)
 - alignment considerations for Solaris [111](#)
 - LUN ostypes [110](#)
 - misaligned on properly aligned LUNs [109](#)
- igroup
 - WWPN [182](#)
- igroups
 - available commands [127](#)
 - binding [69](#), [78](#), [117](#)
 - creating [127](#)
 - deleting [127](#)
 - destroying [127](#)
 - displaying [127](#)
 - manually mapping [68](#)
 - modifying [127](#)
 - modifying OS type [127](#)
 - removing [127](#)
 - renaming [127](#)
 - unbinding [117](#)
- Infinite Volumes
 - not supported for SAN [38](#)
- initiator
 - node
 - name [125](#)
 - node name
 - login [180](#)
 - removing [127](#)
 - initiator configuration
 - displaying [160](#)
 - initiator groups
 - defined [123](#)
 - initiator security list
 - removing [133](#)
 - initiator status
 - displaying [160](#)
- initiators
 - adding [127](#)
 - binding [127](#)
- initiators adapters
 - configuring [152](#)
- interaction

- SVM, LIF [60](#)
- iqn type designator [179](#)
- iSCSI
 - access lists [137](#)
 - adding
 - LIFs [131](#)
 - target alias [131](#)
 - available commands [131](#), [133](#)
 - block access [65](#)
 - commands [132](#), [133](#)
 - configuring network for best performance [128](#)
 - connections
 - available commands [134](#)
 - deleting [134](#)
 - shutdown [134](#)
 - creating
 - targets [131](#)
 - data center bridging [149](#)
 - DCB support [150](#)
 - default TCP port [180](#)
 - disabling [132](#), [133](#)
 - displaying
 - connections [134](#)
 - node names [131](#)
 - target alias [131](#)
 - enabling [130](#), [132](#), [133](#)
 - error messages [145](#)
 - error recovery levels [130](#)
 - explained [178](#)
 - host implementation [135](#)
 - how communication sessions work [181](#)
 - how nodes are identified [179](#)
 - implementation on the storage system [179](#)
 - initiator security
 - method [133](#)
 - initiator security method [133](#)
 - interface access [133](#)
 - iSNS [138](#)
 - license
 - adding [131](#)
 - verifying [131](#)
 - modifying
 - service [131](#)
 - target alias [131](#)
 - target node name [131](#)
 - modifying initiator security method [133](#)
 - new service [131](#)
 - node
 - name [125](#)
 - nodes defined [179](#)
 - security [135](#)
 - security policy [129](#)
 - service
 - commands [131](#)
 - deleting [131](#)
 - destroying [131](#)
 - starting [131](#)
 - stopping [131](#)
 - verifying [131](#)
 - showing [133](#)
 - stopping [132](#), [133](#)
 - storage system interfaces [132](#)
 - supported configurations [55](#)
 - SVM [181](#)
 - troubleshooting [144](#)
- ISCSI
 - service
 - creating [131](#)
- iSCSI displaying
 - interface access [133](#)
- iSCSI initiator security commands [133](#)
- iSCSI service
 - deleting [131](#)
- iSCSI sessions
 - available commands [133](#)
 - deleting [133](#)
 - displaying [133](#)
 - showing [133](#)
 - shutdown [133](#)
- iSNS
 - access lists [138](#)
 - communication [138](#)
 - defined [138](#)
 - discovery [138](#)
 - host [138](#)
 - iSCSI service [138](#)
 - LIF [138](#)
 - registration [138](#), [141](#)
 - server versions [141](#)
 - service
 - adding [143](#)
 - configuration [143](#)
 - configuring [143](#)
 - creating [143](#)
 - deleting [143](#)
 - displaying [143](#)
 - modifying [143](#)
 - removing [143](#)
 - showing [143](#)
 - stopping [143](#)

- updating [143](#)
- iSNS service
 - available commands [143](#)
- L**
- license
 - FC [74, 181](#)
 - iSCSI [131](#)
 - verifying
 - FC [65, 73](#)
 - iSCSI [65, 73](#)
- LIF
 - port set
 - removal [119](#)
 - removing IIF
 - from port sets [119](#)
 - setup process [63](#)
 - troubleshooting [14, 26, 144, 145, 170](#)
- LIFs
 - adding [131](#)
 - adding FC [166](#)
 - adding iSCSI [131](#)
 - advantages [60](#)
 - commands
 - FC [166](#)
 - iSCSI [131](#)
 - creating [65, 74](#)
 - defined [60](#)
 - deleting [121, 131, 166](#)
 - example [60](#)
 - FC [74, 166](#)
 - FC and FCoE zoning restrictions for Cisco switches [122](#)
 - iSCSI [65, 131](#)
 - LUN setup [54](#)
 - managing [118](#)
 - modifying [131, 166](#)
 - moving [119](#)
 - non-disruptive migrations [118](#)
 - required for volume moves [170](#)
 - SVM interaction [60](#)
 - troubleshooting [144](#)
- limiting
 - LUN availability [58](#)
- Linux
 - I/O alignment considerations [111](#)
- Logical
 - Block
 - Provisioning [25](#)
- logical interface [135](#)
- login
 - initiator
 - checks [180](#)
- LUN
 - space-reserved [27, 28](#)
- LUN commands [127](#)
- LUN configurations
 - FC [71](#)
 - iSCSI [64](#)
- LUN ID
 - displaying [127](#)
 - range [55](#)
- LUN not visible [144](#)
- LUN paths
 - displaying [127](#)
- LUN reservations
 - how they work [57](#)
- LUN setup
 - decisions [54](#)
 - FC [71](#)
 - iSCSI [64](#)
 - overview [63](#)
 - wizard [79](#)
- LUN type
 - deciding [11](#)
- LUNs
 - about restoring in SnapVault backup SAN environments [98](#)
 - accessing read-only copies from SnapVault backups [99](#)
 - adding [116](#)
 - available commands [116](#)
 - backups [106](#)
 - binding igroups [63](#)
 - binding port sets [63](#)
 - bringing online [116](#)
 - calculating rate of change [12](#)
 - changing serial number [116](#)
 - cloning [94, 95](#)
 - configured space [115](#)
 - configuring [46, 50](#)
 - controlling I/O performance [112](#)
 - copies [95](#)
 - copying [83](#)
 - creating [68, 116](#)
 - creating FC manually [76](#)
 - creating from files [55](#)
 - decreasing [114](#)
 - delayed availability after volume move [170](#)

- deleting [114](#), [116](#)
- disk space usage [29](#), [32](#)
- displaying
 - configurations [116](#)
 - statistics [116](#)
- displaying serial numbers [116](#)
- geometry [58](#)
- I/O alignment for Linux [111](#)
- I/O alignment for Solaris [111](#)
- increasing size [113](#)
- iSCSI manually [68](#)
- layout [57](#)
- limiting access to [63](#)
- managing [116](#)
- mapped [116](#)
- mapping guidelines [56](#)
- maximum size [58](#)
- maxsize [116](#)
- metadata [68](#), [76](#), [113](#)
- misaligned I/O [109](#)
- modifying [116](#)
- monitoring [18](#)
- offline
 - LUNs [18](#)
- pre-allocation [16](#)
- provisioning [43](#), [46](#), [50](#)
- rate of change [11](#)
- read-only [59](#)
- relocating [116](#)
- renaming [116](#)
- reserve
 - Snapshot [27](#)
- resizing [116](#)
- restoring all in a volume, from SnapVault backup [102](#)
- restoring single, from SnapVault backup [100](#)
- restoring Snapshot copies [85](#), [87](#)
- roadmap [9](#), [10](#)
- size limit [68](#), [76](#), [113](#)
- sizing [116](#)
- Snapshot
 - reserve [27](#)
- space consumption [14](#)
- space reclamation [22](#)
- space-reserved [27](#), [28](#), [46](#), [50](#)
- staying online [23](#)
- taking offline [116](#)
- thin [43](#)
- thin provisioning [14](#)
- thinly provisioned [16](#), [23](#)

- troubleshooting [144](#)
- truncating [114](#)
- unmapping [116](#)
- usable space [58](#)
- used space [115](#)
- volume copy [83](#)
- workflow [9](#), [10](#)

M

- management
 - disk space [169](#)
 - volume space [169](#)
- managing
 - volumes [17](#)
- manually
 - creating FC LUNs [71](#)
 - creating iSCSI LUNs [64](#), [68](#)
 - creating LUN [76](#)
 - map igroups [76](#)
- mapping
 - igroups [68](#), [69](#), [76](#), [78](#), [117](#), [127](#)
 - igroups manually [68](#)
 - LUNs [116](#)
 - manually [76](#)
- mapping conflicts [59](#)
- max_error_recovery_level option [130](#)
- methods
 - data protection [81](#)
- Microsoft Offloaded Data Transfer (ODX)
 - system requirements [174](#)
- modifying
 - iSCSI initiator security method [133](#)
 - LUNs
 - comment [116](#)
 - serial number [116](#)
 - security authentication method [133](#)
- moving
 - LUNs [116](#)
- MPIO [60](#), [63](#), [123](#), [126](#)

N

- network
 - configuring for best iSCSI performance [128](#)
- new
 - FC service [166](#)
 - igroup [127](#)
 - iSCSI initiator security [133](#)
 - port sets [117](#)

node

- name
 - iSCSI [125](#)
- node name
 - storage system [180](#)
- node type designator
 - eui [180](#)
 - iqn [179](#)
- nodes
 - FCP [182](#)
 - iSCSI [179](#)

O

ODX

- improving remote copy performance with [173](#)
- requirements [176](#)
- use cases for [175](#)

ODX (Microsoft Offloaded Data Transfer)

- system requirements [174](#)

Offloaded Data Transfer

- See* ODX

onboard adapters

- initiator [152](#)

optimized paths [60](#)

options

- max_error_recovery_level [130](#)

OS type

- setting [127](#)

ostypes

- for I/O alignment [110](#)

over-provisioning example [15](#)

over-subscribed storage [15](#), [16](#)

overview

- FC block access [71](#)
- iSCSI block access [64](#)
- LUN setup [54](#)
- SVM, LIF [60](#)

P

paths

- LIF considerations [63](#)

performance

- configuring network for best iSCSI [128](#)
- improving remote copy performance with ODX [173](#)

port set

- LIF [119](#)
- unbinding [127](#)

port sets

- adding [117](#)
- commands [117](#)
- creating [69](#), [78](#), [117](#)
- deleting [117](#)
- displaying [117](#)
- removing [117](#)

ports

- displaying [117](#)
- X1143A-R6 [163](#)

protection

- methods of data [81](#)

provisioning

- best practices [17](#)
- guidelines [56](#)
- options [11](#)
- thin [15](#)

Q

quotas [57](#)

R

rate of change

- LUN types [11](#)

Rate of change

- calculating [12](#)

reclaiming space [23](#)

relocating

- disk space [169](#)

removing

- initiator security list [133](#)
- LUNs [116](#)
- port set [117](#)

reservations

- how they work [57](#)

reserves

- considerations for setting fractional [39](#)

restoring

- all LUNs in a volume from SnapVault backup [102](#)
- LUNs [85](#)
- single LUN from SnapVault backup [100](#)

roadmap

- LUNs [9](#), [10](#)

S

SAN configuration limits

- volume and SVM requirements [38](#)

SAN environments

- about configuring and using SnapVault backups in [98](#)
- SAN systems
 - backing up [106](#)
- SCSI
 - SBC-3 standard [25](#)
 - thin provisioning [25](#)
- security authentication method
 - modifying [133](#)
- security method
 - defining [129](#)
 - iSCSI [129](#)
- security password
 - defining [133](#)
- session
 - checks [180](#)
- snap commands
 - SnapRestore [87](#)
- snap reserve [17](#)
- SnapDrive
 - about [178](#)
- SnapMirror destinations
 - mapping read-only LUNs [59](#)
- Snapshot
 - copy
 - duration [13](#)
 - room [13](#)
- Snapshot copies
 - cloning [95](#)
 - data protection [84](#)
 - deleting [89](#)
 - no pre-allocation [27](#)
 - not space-reserved [28](#)
 - pre-allocated [16, 28](#)
 - snap reserve [41](#)
- Snapshot reserve
 - pre-allocated [50](#)
 - without pre-allocated [43, 46](#)
- Snapshots
 - autodelete [40](#)
- SnapVault backups
 - about configuring and using in SAN environments [98](#)
 - accessing read-only LUN copies from [99](#)
 - restoring all LUNs in a volume from [102](#)
 - restoring single LUNs from [100](#)
- Solaris
 - I/O alignment considerations [111](#)
- space
 - how Data ONTAP can automatically add FlexVol volume [41](#)
 - reclamation [23, 25](#)
 - space allocation
 - guidelines [57](#)
 - LUN [11](#)
 - Snapshot copies [11](#)
 - space-reserved LUN [27, 28](#)
 - thin [16](#)
 - volume [11](#)
 - space consumption
 - LUN [14](#)
 - space reclamation
 - keeping LUNs online [22](#)
 - Space Reclamation [19](#)
 - space reservations
 - See* reservations
 - space usage
 - understanding [29](#)
 - with SnapDrive reserve [29](#)
 - within Snapshot reserve [29](#)
 - space-reserved
 - LUN [27, 28](#)
 - space
 - allocation [27](#)
 - storage administrator [15](#)
 - storage efficiency [11, 15, 16, 27](#)
 - storage system node name
 - defined [180](#)
- SVM
 - advantages [60](#)
 - commands
 - FC [166](#)
 - iSCSI [131](#)
 - deleting [131, 166](#)
 - example [60](#)
 - FC
 - service [166](#)
 - FC traffic [182](#)
 - iSCSI
 - service [131](#)
 - iSCSI service [131](#)
 - iSCSI traffic [181](#)
 - isns
 - SVM [141](#)
 - iSNS [138](#)
 - LIF interaction [60](#)
 - usage [60](#)
- SVM setup
 - wizard [79](#)

SVMs
 restrictions for SAN [38](#)

switch
 FC and FCoE zoning requirement [122](#)

sysconfig [159](#)

T

tape backup [105](#)

target adapter [150](#)

target adapters
 FCoE [149](#)

targets adapters
 configuring [154](#)

TCP port
 default for iSCSI [180](#)

thin provisioning
 about [15](#)
 best practices [17](#)
 LUNs [14](#), [43](#)
 LUNs offline [16](#)

troubleshooting
 delayed LUN availability after volume move [170](#)
 iSCSI error messages [145](#)
 keeping thinly provisioned LUNs online [14](#)
 LUN [144](#)
 LUNs offline [26](#)
 port sets [144](#)

troubleshooting iSCSI [144](#)

U

unbinding
 igroups [117](#)

unified Ethernet
 overview [149](#)

unified target adapters
 changing optical modules [162](#)
 configuration rules and supported settings [163](#)
 configuring ports [163](#)
 default configuration [163](#)
 managing [149](#), [161](#)
 managing as CNAs [162](#)
 viewing settings [162](#)
 x1143A-R6 [161](#)

usable space
 LUNs [58](#)

use cases
 for ODX [175](#)

UTA [149](#)

V

VAAI features
 copy offload [172](#)
 extended copy feature [172](#)
 VERIFY AND WRITE feature [172](#)
 WRITE SAME feature [172](#)

VERIFY AND WRITE feature
 environment [172](#)
 invoked automatically [172](#)
 see COMPARE WRITE feature [172](#)

verifying
 FC license [65](#), [73](#)
 iSCSI license [65](#), [73](#)

virtualization
 Microsoft
 virtualization [171](#)
 VMware
 virtualization [171](#)

volume
 copying
 LUNs [83](#)
 deleting
 Snapshot copies [89](#)
 management [169](#)
 restoring LUNs [87](#)
 scioning [95](#)
 Snapshot copies [95](#)
 space-reserved LUN [27](#), [28](#)
 thinly provisioned [28](#)

volume autosize [17](#)

volume copy command [83](#)

volume options
 Snapshot reserve [41](#)

volume size
 autodelete [21](#)

volume usage
 no Snapshot copies [29](#)
 Snapshot copies [32](#)

volumes
 configuration options [11](#)
 configuring [11](#), [43](#), [46](#), [50](#)
 move requirements [170](#)
 restoring LUNs [85](#)
 restrictions for SAN [38](#)
 Snapshot
 reserve [16](#)
 space
 allocation [43](#)
 thinly provisioned LUN [16](#)

W

wizard

LUN setup [79](#)

workflow

LUNs [9](#), [10](#)

WRITE SAME feature

environment [172](#)

invoked automatically [172](#)

VAAI feature [172](#)

WWNN [147](#)

WWPN

assigning [167](#)

assignment [147](#)

displaying [167](#)

target adapter

changing [165](#)

usage [182](#)

X

X1143A-R6 device [163](#)

Z

zero fat provisioning [15](#)

zoning

requirements for Cisco switches [122](#)