



Data ONTAP® 8.2

SAN Administration Guide

For 7-Mode



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: 215-07986_A0
May 2013

Contents

Preparing for LUN setup workflow	8
Deciding which LUN type to use in SAN environments	9
How the rate of change in your data determines LUN type	9
Calculating Rate of Change	10
Additional considerations for thinly provisioned LUNs	11
What LUN thin provisioning is	11
Why thinly provisioned LUNs go offline	12
How to keep your thinly provisioned LUNs online	12
When to use thinly provisioned LUNs	18
Additional considerations for space-reserved LUNs	18
What space-reserved LUNs are	18
When to use space-reserved LUNs with Snapshot reserve	18
When to use space-reserved LUNs without Snapshot reserve	19
Space-reserved LUNs with Snapshot reserve	20
Space-reserved LUNs in a thinly provisioned volume	20
Storage Provisioning for SAN	21
Storage units for managing disk space	21
Guidelines for provisioning storage in a SAN environment	22
Estimating storage in a SAN environment	23
How much room do you need for Snapshot copies	23
Determining the volume size when using Snapshot autodelete	24
Determining the volume size and fractional reserve setting when you need Snapshot copies	25
Determining the volume size when you do not need Snapshot copies	27
Creating LUNs on storage systems	28
Creating an aggregate	28
Creating a volume	29
Volume configuration options for a SAN environment	30
Configuring volumes in a SAN environment	34
Setting up LUNs and igroups	44
Creating LUNs on vFiler units	46
Displaying vFiler LUNs	48

LUN configuration	48
Information required to create a LUN	48
Guidelines for LUN layout and space allocation	51
LUN management	52
Displaying command-line Help for LUNs	52
Controlling LUN availability	53
Moving LUNs	54
Modifying LUN descriptions	54
How LUN reservations work	54
Enabling or disabling space reservations for LUNs	55
Accessing LUNs with NAS protocols	56
Checking LUN, igroup, and FC settings	56
Displaying LUN serial numbers	57
Displaying LUN statistics	58
Displaying LUN mapping information	59
Displaying detailed LUN information	59
Displaying hidden staging area LUNs	60
LUN alignment in virtual environments	60
Removing LUNs	61
Misaligned I/O can occur on properly aligned LUNs	61
igroup management	63
What igroups are	63
igroup example	63
Creating igroups	64
Required information for creating igroups	65
Creating FC protocol igroups on UNIX hosts using the sanlun command ...	66
Creating igroups for a non-default vFiler unit	67
igroup configuration	68
Enabling ALUA	68
Enabling report_scsi_name	69
Fibre Channel initiator request management	70
LUN and igroup mapping	75
What LUN mapping is	75
Required information for mapping a LUN to an igroup	75
Guidelines for mapping LUNs to igroups	76
SnapMirror destinations and read-only LUNs	76


How to make LUNs available on specific FC target ports	77
Unmapping LUNs from igroups	77
Deleting igroups	78
Adding initiators to an igroup	79
Removing initiators from an igroup	79
Displaying initiators	80
Renaming igroups	80
Setting the operating system type for an igroup	81
SAN Protocol Management	82
iSCSI network management	82
Enabling multi-connection sessions	82
Enabling error recovery levels 1 and 2	83
iSCSI service management	83
iSNS server registration	92
Displaying initiators connected to the storage system	95
iSCSI initiator security management	96
Target portal group management	105
Displaying iSCSI statistics	116
Displaying iSCSI session information	120
Displaying iSCSI connection information	121
Guidelines for using iSCSI with HA pairs	122
iSCSI troubleshooting tips	123
FC SAN management	127
How to manage FC with HA pairs	127
How to use port sets to make LUNs available on specific FC target ports ..	130
FC service management	135
Managing systems with Fibre Channel adapters	147
Unified Ethernet network management	161
Fibre Channel over Ethernet overview	161
What data center bridging is	162
Displaying DCB settings	163
Disk space management	165
Commands to display space information	165
Examples of disk space monitoring using the df command	166
Monitoring disk space on volumes with LUNs that do not use Snapshot copies	166

Monitoring disk space on volumes with LUNs that use Snapshot copies . .	168
Moving your volumes nondisruptively	170
Ways to use volume move	170
Requirements for performing a volume move	171
How the setup phase of volume move works	172
How the data copy phase of volume move works	173
How the cutover phase of volume move works	173
Performing the volume move operation	174
Pausing the volume move operation	175
Resuming the volume move operation	176
Monitoring the volume move status	176
Performing manual cutover of the volume move operation	177
Canceling the volume move operation	177
Working with VMware VAAI features for ESX hosts	177
Requirements for using the VAAI environment	178
Methods for determining whether VAAI features are supported	179
Statistics collected for VAAI counters	179
Viewing statistics for the VAAI features	181
Data protection with Data ONTAP	183
Data protection methods	183
LUN clones	185
Reasons for using FlexClone LUNs	185
Differences between FlexClone LUNs and LUN clones	185
Cloning LUNs	186
LUN clone splits	187
Displaying the progress of a clone-splitting operation	188
Stopping the clone-splitting process	188
Deleting Snapshot copies	188
Deleting backing Snapshot copies of deleted LUN clones	189
Deleting busy Snapshot copies	193
Restoring a Snapshot copy of a LUN in a volume	196
Restoring a single LUN	198
Backing up SAN systems to tape	199
Using volume copy to copy LUNs	202
Basic block access concepts	203
How hosts connect to storage systems	203

What Host Utilities are	203
What ALUA is	204
About SnapDrive for Windows and UNIX	204
How Data ONTAP implements an iSCSI network	205
What iSCSI is	205
What iSCSI nodes are	205
How iSCSI is implemented on the host	205
How iSCSI target nodes connect to the network	205
How iSCSI nodes are identified	206
How the storage system checks initiator node names	207
Default port for iSCSI	207
What target portal groups are	207
What iSNS is	208
What CHAP authentication is	208
How iSCSI communication sessions work	209
How iSCSI works with HA pairs	209
Setting up the iSCSI protocol on a host and storage system	209
How Data ONTAP implements an FC SAN	210
What FC is	210
What FC nodes are	210
How FC target nodes connect to the network	211
How FC nodes are identified	211
Copyright information	213
Trademark information	214
How to send your comments	215
Index	216

Preparing for LUN setup workflow

A prerequisite for creating LUNs is that you have an aggregate, 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 9

[*Configuring volumes in a SAN environment*](#) on page 34

[*Creating LUNs on storage systems*](#) on page 28

[*LUN configuration*](#) on page 48

Related tasks

[*Configuring volumes for thinly provisioned LUNs without Snapshot reserve*](#) on page 34

[*Configuring volumes for space-reserved LUNs with Snapshot reserve*](#) on page 36

[*Configuring volumes for spaced-reserved LUNs without Snapshot reserve*](#) on page 37

[*Creating LUNs on vFiler units*](#) on page 46

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

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)	• Use space-reserved LUNs • Need extra room for Snapshot copies • Set fractional reserve to 100%
Very low rate of change (low rate of change)	• Use thinly provisioned LUNs
Steady growth (low rate of change)	• Use thinly provisioned LUNs • Use volume autogrow

Rate of deletion or change of data	Notes
Inconsistent growth	• Use thinly provisioned LUNs • Use volume auto grow 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.

Additional considerations for thinly provisioned LUNs

When you use thinly provisioned LUNs, Data ONTAP has several strategies that can help keep your LUNs online, such as effective monitoring, space reclamation, and appropriate volume sizing.

Additionally, Data ONTAP has several strategies for making more space available on your volumes:

- volume auto grow
- Snapshot auto delete
- FlexClone LUN auto delete
- Automatic space reclamation

For more information about these strategies, see the *Data ONTAP Storage Management Guide for 7-Mode*.

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 controller than is actually available on the storage controller. 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 storage controller 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.

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 autosize command
- Enable Snapshot autodelete
- Enable FlexClone autodelete

For more information about the snap autodelete command, see the *Data ONTAP Storage Management Guide for 7-Mode*.

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.
- **Enable the LUN -e space_alloc option.**

Related concepts

[Volume option best practices for thinly provisioned LUNs](#) on page 13

[How the rate of change in your data determines LUN type](#) on page 9

[Space Reclamation](#) on page 14

Related tasks

[Calculating Rate of Change](#) on page 10

[Enabling the space_alloc option](#) on page 17

Volume option best practices for thinly provisioned LUNs

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

When you thinly provision a 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
- fractional reserve - determines how much space is for Snapshot copy overwrites
- 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
Fractional reserve	0	volume 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

[Volume Autosizing](#) on page 31

[What Snapshot autodelete is](#) on page 32

[Considerations for setting fractional reserve](#) on page 31

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)

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



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

Starting with Data ONTAP 8.2, you can use the `space_alloc` 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_alloc` 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_alloc` 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 autosize command
- Enable Snapshot autodelete
- Enable FlexClone autodelete

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 *Data ONTAP Storage Management Guide for 7-Mode*.

Related information

Host Utilities documentation: www.support.netapp.com

Enabling the `space_alloc` option

If a LUN runs out of space and the containing volume cannot automatically grow more, the LUN goes offline. To keep a LUN online, you should set the LUN option `space_alloc` to enable.

About this task

The LUN option `-space_alloc` is set to disable by default. If you leave this option set to disable, then the LUN goes offline when the volume runs out of space and is not permitted to grow. If you set this option to enable, Data ONTAP notifies the host that the LUN has run out of space. However, the LUN stays online. The host cannot write to the LUN, but the host can still read the LUN.

Steps

1. To enable `space_alloc`, enter the following command:

```
lun set space_alloc /vol/vol0/lun_name enable
```

Example

```
system1> lun set space_alloc /vol/vol0/lun1 enable
system1> lun set space_alloc /vol/vol0/lun1
Reporting of space allocation events is enabled
```

2. Rescan/Remount the LUN so that the change to `space_alloc` option is reflected on the host.

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

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 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

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 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

[Space-reserved LUNs with Snapshot reserve](#) on page 20

Related tasks

[Configuring volumes for space-reserved LUNs with Snapshot reserve](#) on page 36

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 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 tasks

[Configuring volumes for spaced-reserved LUNs without Snapshot reserve](#) on page 37

[Calculating Rate of Change](#) on page 10

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 with Snapshot reserve

Space-reserved LUNs and Snapshot copies have pre-allocated space that can be continually overwritten. This guaranteed space is not available to any other LUNs or Snapshot copies within the volume.

Related concepts

[When to use space-reserved LUNs with Snapshot reserve](#) on page 18

Related tasks

[Configuring volumes for space-reserved LUNs with Snapshot reserve](#) on page 36

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 tasks

[Calculating Rate of Change](#) on page 10

Storage Provisioning for SAN

Storage provisioning includes the process of creating LUNs, creating igroups, and mapping the LUNs to the igroups. There are various steps involved in this process.

Storage units for managing disk space

To properly provision storage, it is important to define and distinguish between the different units of storage.

The following list defines the various storage units:

- | | |
|--|---|
| Plexes | <p>A collection of one or more Redundant Array of Independent Disks (RAID) groups that together provide the storage for one or more Write Anywhere File Layout (WAFL) file system aggregates or traditional volumes.</p> <p>Data ONTAP uses plexes as the unit of RAID-level mirroring when the SyncMirror software is enabled.</p> |
| Aggregates | <p>The physical layer of storage that consists of the disks within the RAID groups and the plexes that contain the RAID groups.</p> <p>It is a collection of one or two plexes, depending on whether you want to take advantage of RAID-level mirroring. If the aggregate is unmirrored, it contains a single plex. Aggregates provide the underlying physical storage for traditional and FlexVol volumes.</p> |
| Traditional or flexible volumes | <p>A traditional volume is directly tied to the underlying aggregate and its properties. When you create a traditional volume, Data ONTAP creates the underlying aggregate based on the properties you assign with the <code>vol create</code> command, such as the disks assigned to the RAID group and RAID-level protection.</p> <p>A FlexVol volume is a volume that is loosely coupled to its containing aggregate. A FlexVol volume can share its containing aggregate with other FlexVol volumes. Thus, a single aggregate can be the shared source of all the storage used by all the FlexVol volumes contained by that aggregate.</p> <p>You can use either traditional or FlexVol volumes to organize and manage system and user data. A volume can hold qtrees and LUNs.</p> <p>After you set up the underlying aggregate, you can create, clone, or resize FlexVol volumes without regard to the underlying physical storage. You do not have to manipulate the aggregate frequently.</p> |

- Qtrees** A qtree is a subdirectory of the root directory of a volume. You can use qtrees to subdivide a volume in order to group LUNs.
- LUNs** A logical unit of storage that represents all or part of an underlying physical disk. You can create LUNs in the root of a volume (traditional or flexible) or in the root of a qtree.
- Note:** You should not create LUNs in the root volume because it is used by Data ONTAP for system administration. The default root volume is /vol/vol0.

For detailed information about storage units, see the *Data ONTAP Storage Management Guide for 7-Mode*.

Related information

[Data ONTAP documentation on the NetApp Support Site-support.netapp.com](http://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 traditional or FlexVol volumes that contain LUNs, regardless of which provisioning method you choose:

- You should not create any LUNs in the system's root volume.
Data ONTAP uses this volume to administer the storage system. The default root volume is /vol/vol0.
- 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 must ensure that the volume option `create_ucose` is set to on.
- You can make the required changes to the snap reserve default settings.
You can change the `snapsreserve` setting for the volume to 0, set the `snap schedule` so that no controller-based Snapshot copies are taken, and delete all Snapshot copies after you create the volume.
- 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 *Data ONTAP Storage Management Guide for 7-Mode*.

Related information

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

Estimating storage in a SAN environment

When provisioning storage, you need to estimate the size of your storage if you use autodelete, fractional reserve or you do not need Snapshot copies.

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 Snapshot autodelete, see the *Data ONTAP Data Protection Online Backup and Recovery Guide for 7-Mode*. For more information about volume autosize, see the *Data ONTAP Storage Management Guide for 7-Mode*.

How long do you need a Snapshot copy for?	Notes
Lots of Snapshot copies for a long duration	• 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	• Use Snapshot reserve • Use volume autogrow before snapshot autodelete
Some Snapshot copies for a short duration	• 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

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 ($400 \text{ GB ROC} \times 10 \text{ Snapshot copies}$).

- 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 Snapshot autodelete function, see the *Data ONTAP Storage Management Guide for 7-Mode*, and for more information about working with traditional and FlexVol volumes, see the *Data ONTAP Storage Management Guide for 7-Mode*.

Related concepts

[What Snapshot autodelete is](#) on page 32

Related tasks

[Calculating Rate of Change](#) on page 10

Related information

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

Determining the volume size and fractional reserve setting when you need Snapshot copies

Use the fractional reserve method to estimate the size of volumes on which you need to create Snapshot copies. Fractional reserve is not necessary for volumes when you do not need Snapshot copies.

About this task

The required volume size for a volume when you need Snapshot copies depends on several factors, including how much your data changes, how long you need to keep Snapshot copies, and how much data the volume is required to hold.

Steps

1. Add up all of the space-reserved LUNs.

Example

If you know your database needs 40 GB of space, you must create a 40 GB space-reserved LUN.

2. 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.

3. 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 40 GB LUN, and you estimate that your data changes at a rate of about 10 percent each day, which in this case is 4 GB per day. You want to take one Snapshot copy each day and want to keep three weeks of Snapshot copies, for a total of 21 days of Snapshot copies. The amount of space you need for Snapshot copies is 21×4 GB, or 84 GB.

4. Determine how much space you need for overwrites by multiplying your ROC by number of days you want to keep Snapshot copies before deleting.

Space required for overwrites = ROC $\times$ number of days you want to keep Snapshot copies before deleting

Example

You have a 40 GB LUN and your data changes at a rate of 4 GB each day. You want to retain daily snapshots for 3 days. You need $4 \text{ GB} \times 3$, or 12 GB of additional space in the volume reserved for overwrites to the LUN.

5. Calculate the required volume size by adding together the total data size and the space required for Snapshot copies.

Volume size = Total data size + space required for Snapshot copies

Example

You have a 40 GB LUN and 12 GB of Snapshot copies. The volume size needs to be 52 GB.

6. Calculate the minimum fractional reserve value for this volume by dividing the size of space required for Snapshots by the total size of the space-reserved LUNs in the volume. Setting this value will enable Data ONTAP to create Snapshots only when the minimum amount of space is available in the volume.

Fractional reserve = space required for overwrites $\div$ total data size.

Example

You have a 40 GB LUN. You require 12 GB of changes held in Snapshot copies. 12 GB is 30 percent of the total LUN size. Therefore the smallest volume size is 52 GB and you must set the Fractional Reserve value to 30 to enable Snapshot creation to succeed.

Volume size calculation example

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

- You need to create two 50 GB LUNs. The total LUN size is 100 GB.
- Your data changes at a rate of 10 percent of the total LUN size each day. Your ROC is 10 GB per day (10 percent of 100 GB).
- You take one Snapshot copy each day and you want to keep 10 days of Snapshot copies. You need 100 GB of space for Snapshot copies (10 GB ROC × 10 Snapshot copies).

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 200 GB (100 GB + 100 GB).

Related concepts

Considerations for setting fractional reserve on page 31

Related tasks

Calculating Rate of Change on page 10

Determining the volume size when you do not need Snapshot copies

If you are not using Snapshot copies, the size of your volume depends on the size of the LUNs and whether you are using traditional or FlexVol volumes.

Before you begin

Before you determine that you do not need Snapshot copies, you should verify the method for protecting data in your configuration. Most data protection methods, such as SnapRestore, SnapMirror, SnapManager for Microsoft Exchange or Microsoft SQL Server, SyncMirror, dump and restore, and ndmpcopy methods, rely on Snapshot copies. If you are using any of these methods, you cannot use this procedure to estimate volume size.

Note: Host-based backup methods do not require Snapshot copies.

Step

1. The FlexVol volume should be at least as large as the size of the data to be contained by the volume.

Example

If you need a FlexVol volume to contain two 200 GB LUNs, you must ensure that the aggregate containing the FlexVol has enough space to provide at least 400 GB of storage capacity.

Creating LUNs on storage systems

You can create LUNs on physical storage systems or on vFilers that have been partitioned using MultiStore. You must configure aggregates and volumes to contain your LUNs before your LUNs can be created on your storage system.

If your aggregates and volumes have already been setup, you can go directly to setting up LUNs and igroups. If your aggregates and volumes have not been setup, you must configure them before creating LUNs and igroups.

Related concepts

[Setting up LUNs and igroups](#) on page 44

Creating an aggregate

You create an aggregate to provide storage to one or more FlexVol volumes.

Before you begin

You should know what drives or array LUNs will be used in the new aggregate.

If you have multiple drive types in your system (heterogeneous storage), you should understand how you can ensure that the correct drive type is selected.

About this task

You can display a list of the available spares by using the `aggr status -s` command.

Aggregate names must conform to the following requirements:

- Begin with either a letter or an underscore (_).
- Contain only letters, digits, and underscores.
- Contain no more than 250 characters.

Steps

1. Create the aggregate by entering the following command:

```
aggr create aggr_name [-f] [-m] [-n] [-t {raid0 | raid4 | raid_dp}] [-r  
raidsizes] [-T disk-type] -R rpm] [-L] [-p] disk-list
```

aggr_name is the name for the new aggregate.

-f overrides the default behavior that does not permit drives in a plex to belong to different pools. This option also enables you to mix drives with different RPM speeds even if the appropriate `raid.rpm` option is not off.

-m specifies the optional creation of a SyncMirror-replicated volume if you want to supplement RAID protection with SyncMirror protection.

-n displays the results of the command but does not execute it. This is useful for displaying the drives that would be automatically selected prior to executing the command.

-t {raid0 | raid4 | raid_dp} specifies the level of RAID protection you want to provide for this aggregate. If no RAID level is specified for an aggregate composed of disks, the default value (raid_dp) is applied. raid0 is used only for array LUNs.

-r *raidsize* is the maximum size of the RAID groups for this aggregate. If no size is specified, the default is used.

-T *disk-type* specifies the Data ONTAP drive type. This option is needed when creating aggregates on systems that have mixed drive types or both drives and array LUNs.

-R *rpm* specifies the type of disk to use based on its speed. Valid values for *rpm* include 5400, 7200, 10000, and 15000.

-p specifies the pool from which the drives are selected.

disk-list is one of the following values:

- *ndisks[@disk-size]*
ndisks is the number of drives to use.
disk-size is the drive size to use, in gigabytes.
- -d *disk_name1 disk_name2... disk_nameN*
disk_name1, *disk_name2*, and *disk_nameN* are drive IDs of available drives; use a space to separate drive IDs.

2. Verify the RAID group and drives of your new aggregate by entering the following command:

aggr status -r *aggr_name*

Examples

The following command creates a 64-bit aggregate called newfastaggr, with 20 drives, the default RAID group size, and all drives with 15K RPM:

aggr create newfastaggr -R 15000 20

The following command creates a 64-bit aggregate called newFCALaggr.

aggr create newFCALaggr -T FCAL 15

Creating a volume

After determining the necessary size of your volume, you can create the volume. Volumes must be created before LUNs.

Details on how to create volumes can be found in the *Data ONTAP Storage Management Guide for 7-Mode*.

Volume configuration options for a SAN environment

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 allocate space ahead of time, or do you want to allocate space as you need the space?

You can pre-allocate space or add space as required for your LUNs and Snapshot copies in your volume. You must answer the following questions to determine the type 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 pre-allocate space on your volume for your LUNs and Snapshot copies?
- Do you want to pre-allocate 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 without Snapshot reserve
- Space-reserved LUNs without Snapshot reserve
- Space-reserved LUNs with Snapshot reserve
- Thinly provisioned LUNs with volume autosize enabled

Methods for managing volume size

Before estimating the necessary size of your volume, you must decide how you want to manage storage at the volume level.

In SAN environments, there are three methods to consider for managing your storage at the volume level: volume autosize, Snapshot autodelete and fractional reserve. The method you choose will determine how you later estimate the volume size. In Data ONTAP, by default, fractional reserve is set to 100 percent, and by default, volume autosize and Snapshot autodelete are both disabled. However, in a SAN environment, it usually makes more sense to use the Snapshot autodelete method, or sometimes, the autosize method, which are less complicated than using the fractional reserve method.

Volume autosize	Volume autosize allows you to automatically make more free space available for a FlexVol volume when that volume is nearly full by incrementally increasing the volume size.
Snapshot autodelete	Snapshot autodelete allows you to automatically reclaim space consumed by Snapshot copies when the volume is low in available space.

Fractional reserve Fractional reserve is a volume setting that enables you to configure how much space Data ONTAP reserves in the volume for overwrites in space-reserved LUNs and space-reserved files when Snapshot copies are created.

Related concepts

[Volume option best practices for thinly provisioned LUNs](#) on page 13

[Volume Autosizing](#) on page 31

[Considerations for setting fractional reserve](#) on page 31

[What Snapshot autodelete is](#) on page 32

Volume Autosizing

Volume autosize is useful if the volume's containing aggregate has enough space to support a larger volume. Volume autosize allows you to use the free space in the containing aggregate as a pool of available space shared between all the volumes on the aggregate.

Volumes can be configured to automatically grow as needed, as long as the aggregate has free space. When using the volume autosize method, you can increase the volume size incrementally and set a maximum size for the volume. You will need to monitor space usage of both the aggregate and the volumes within that aggregate to ensure volumes are not competing for available space.

Note: The autosize capability is disabled by default, so you must enable and configure it by using the `vol autosize` command. You can also use this command to view the current autosize settings for a volume.

For more information, see the *Data ONTAP Storage Management Guide for 7-Mode*.

Considerations for setting fractional reserve

Fractional reserve, also called *LUN overwrite reserve*, enables you to control the size of the overwrite reserve for reserved LUNs and files in a 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 `vol options` 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
- Autogrow feature enabled

In addition, you must monitor the free space in the associated aggregate. If the aggregate becomes full enough that the volume is prevented from growing, then data modification operations could fail even with all of the other configuration settings in place.

If you do not want to monitor aggregate free space, 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

For more information about using fractional reserve, see the following Technical Reports:

- *TR-3965: Thin Provisioning Deployment and Implementation Guide*
- *TR-3483: Thin Provisioning in a NetApp SAN or IP SAN Enterprise Environment*

Related information

Technical Report: Thin Provisioning Deployment and Implementation Guide: media.netapp.com/documents/tr-3965.pdf

Technical Report: Thin Provisioning in a NetApp SAN or IP SAN Enterprise Environment: media.netapp.com/documents/tr3483.pdf

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 can set that threshold, or *trigger*, to automatically delete Snapshot copies when:

- 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*.

When to use the autodelete configuration

Before implementing the autodelete configuration, it is important to consider the conditions under which this configuration works best.

The autodelete configuration is particularly useful under the following circumstances:

- You do not want your volumes to affect any other volumes in the aggregate.
For example, if you want to use the available space in an aggregate as a shared pool of storage for multiple volumes or applications, use the `autosize` option instead. Autosize is disabled under this configuration.
- Ensuring availability of your LUNs is more important to you than maintaining old Snapshot copies.

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 object). 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.
For more information about deleting Snapshot copies automatically, see the *Data ONTAP Data Protection Online Backup and Recovery Guide for 7-Mode*.

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.

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 *Data ONTAP Storage Management Guide for 7-Mode*.

Configuring volumes for thinly provisioned LUNs without Snapshot reserve

When you configure your volume for thinly provisioned LUNs without Snapshot copies, you get excellent storage utilization because you can add space to your volume, LUN, and Snapshot copies as needed. These volume configurations enable you to manage your volumes and LUNs more effectively by allowing your LUNs and volumes to grow automatically.

Before you begin

You have created a volume.

About this task

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

Steps

1. Use the `vol options` command to set space guarantee to none.

Example

```
vol options vol1 guarantee none
```

2. Use the `vol options` command to set fractional reserve to 0.

Example

```
vol options vol1 fractional_reserve 0
```

3. Use the `vol autosize` command to turn on volume autosize.

Example

```
vol autosize vol1 on
```

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

Example

```
vol autosize vol1 -m 40g -i 5g
```

5. Use the `vol options` command to set the space management first try option to `volume_grow` (autosize).

Example

```
vol options vol1 try_first volume_grow
```

6. Use the `snap reserve` command to change the snap reserve setting to 0.

Example

```
snap reserve vol1 0
```

7. Use the `snap reserve` command to verify the snap reserve setting has been changed to 0.

Example

```
snap reserve vol1
```

8. Use the `snap sched` command to disable the scheduled creation of Snapshot copies.

Example

```
snap sched vol1 0 0 0
```

9. Use the `snap sched` command to verify the scheduled creation of Snapshot copies has been disabled.

Example

```
snap sched vol1
```

10. Use the `vol status` command to verify the changes you made for volume autosize and snap reserve.

Example

```
vol status vol1 -v
```

11. Use the `snap autodelete` command to disable Snapshot autodelete.

Example

```
snap autodelete vol1 off
```

12. Use the `snap autodelete` command to verify the change you made for Snapshot autodelete.

Example

```
snap autodelete vol1
```

Result

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

Configuring volumes for space-reserved LUNs with Snapshot reserve

When you pre-allocate space for LUNs and Snapshot copies, you guarantee that the space is used just for those LUNs and Snapshot copies. The pre-allocated space for the LUNs and the Snapshot copies is not available to any other LUNs or Snapshot copies within that same volume.

Before you begin

You have created a volume.

About this task

You should use a 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 `vol options` command to set space guarantee to volume.

Example

```
vol options vol1 guarantee volume
```

2. Use the `vol options` command to set fractional reserve to 100.

Example

```
vol options vol1 fractional_reserve 100
```

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

Example

```
vol autosize vol1 off
```

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

Example

```
snap reserve vol1 0
```

5. Use the `snap reserve` command to verify the setting change.

Example

```
snap reserve vol1
```

6. Use the `snap sched` command to disable the scheduled creation of Snapshot copies.

Example

```
snap sched vol1 0 0 0
```

7. Use the `snap sched` command to verify scheduled creation of Snapshot copies has been disabled.

Example

```
snap sched vol1
```

8. Use the `vol status` command to verify changes.

Example

```
vol status vol1 -v
```

9. Use the `snap autodelete` command to disable autodelete.

Example

```
snap autodelete vol1 off
```

10. Use the `snap autodelete` command to verify your changes.

```
snap autodelete vol1
```

Result

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

Related concepts

[*When to use space-reserved LUNs with Snapshot reserve*](#) on page 18

[*Space-reserved LUNs with Snapshot reserve*](#) on page 20

Configuring volumes for spaced-reserved LUNs without Snapshot reserve

When you configure a space-reserved LUN, this space is pre-allocated and not available to other LUNs or Snapshot copies within the volume. However, when Snapshot reserve is not pre-allocated, 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 a 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 `vol options` command to set space guarantee to volume.

Example

```
vol options vol1 guarantee volume
```

2. Use the `vol options` command to set fractional reserve to 0.

Example

```
vol options vol1 fractional_reserve 0
```

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

Example

```
vol autosize vol1 on
```

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

Example

```
vol autosize vol1 -m 40g -i 5g
```

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

Example

```
vol options vol1 try_first volume_grow
```

6. Use the `snap reserve` command to set Snapshot reserve to 0.

Example

```
snap reserve vol1 0
```

7. Use the `snap reserve` command to verify the Snapshot reserve has been set to 0.

Example

```
snap reserve vol1
```

8. Use the `snap sched` command to disable the scheduled creation of Snapshot copies.

Example

```
snap sched vol1 0 0 0
```

9. Use the `snap sched` command to verify scheduled creation of Snapshot copies has been disabled.

```
snap sched vol1
```

10. Use the `vol status` command to verify your settings.

Example

```
vol status vol1 -v
```

11. Use the `snap autodelete` command to enable Snapshot autodelete.

Example

```
snap autodelete vol1 on
```

12. Use the `snap autodelete` command to set the autodelete trigger.

Example

```
snap autodelete vol1 trigger volume
```

13. Use the `snap autodelete` command to set the delete order to delete oldest Snapshot copy first.

Example

```
snap autodelete vol1 delete_order oldest_first
```

14. Use the `snap autodelete` command to verify your settings.

Example

```
snap autodelete vol1
```

Result

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

Related concepts

[When to use space-reserved LUNs without Snapshot reserve](#) on page 19

Volume Options and Settings

After you create your volume, you need to modify some of the default settings. If you are using Snapshot autodelete, you also need to set volume options related to that configuration such as space guarantee, autosize, fractional reserve, `try_first` and Snapshot copy.

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 20 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).

When finished, you must ensure that the `create_ucose` volume is enabled.

Turning off the automatic Snapshot copy schedule

When creating volumes that contain LUNs, you should turn off the automatic Snapshot copy schedule and verify that setting.

Steps

1. Turn off the automatic Snapshot copy schedule by entering the following command:

```
snap sched volname 0 0 0
```

Example

```
snap sched vol1 0 0 0
```

This command turns off the Snapshot copy schedule because there are no weekly, nightly, or hourly Snapshot copies scheduled. You can still take Snapshot copies manually by using the `snap` command.

2. Verify that the automatic Snapshot copy schedule is off by entering the following command:

```
snap sched [volname]
```

Example

```
snap sched vol1
```

The following output is a sample of what is displayed:

```
Volume vol1: 0 0 0
```

Deleting all existing Snapshot copies in a volume

If there is no space reservation for LUNs, then you must delete the existing Snapshot copies in the volume.

Step

1. Delete the existing Snapshot copies by entering the following command:

```
snap delete -a volname
```

Setting the percentage of snap reserve space to zero

When creating volumes that contain LUNs, you should set the percentage of space reserved for Snapshot copies to zero. Setting space reserve to zero ensures that there are no Snapshot copies for the volume containing LUNs.

Steps

1. Set the percentage by entering the following command:

```
snap reserve volname percent
```

Example

```
snap reserve vol1 0
```

2. Verify the percentage that is set by entering the following command:

```
snap reserve [volname]
```

Example

```
snap reserve vol1
```

The following output is a sample of what is displayed:

```
Volume vol1: current snapshot reserve is 0% or 0 k-bytes.
```

Enabling the create_unicode volume option

Data ONTAP requires that the path of a volume or qtree containing a LUN is in the Unicode format. This option is off by default when you create a volume. It is important to enable this option for volumes that contain LUNs.

Step

1. Enable the create_unicode option by entering the following command:

```
vol options volname create_unicode on
```

Example

```
vol options vol1 create_unicode on
```

Verifying the create_ucose volume option

You can use the `vol status` command to verify that the `create_ucose` volume option is enabled to avoid directory conversion.

Step

1. Verify that the `create_ucose` option is enabled (on) by entering the following command:

```
vol status [volname] -v
```

Example

```
vol status vol1 -v
```

Note: If you do not specify a volume, the status of all the volumes is displayed.

Result

The following output example shows that the `create_ucose` option is on:

```
Volume  State  Status  Options
vol1    online  normal  nosnap=off, nosnapdir=off,
                                minra=off, no_atime_update=off,
                                raidsize=8, nvfail=off,
snapmirrored=off,
                                resyncsnaptime=60, create_ucose=on
                                convert_ucose=off,
                                maxdirsize=10240,
                                fs_size_fixed=off,
                                create_reserved=on
                                raid_type=RAID4

Plex /vol/vol1/plex0: online, normal, active
RAID group /vol/vol1/plex0/rg0: normal
```

After you finish

If necessary, you should enable the `create_ucose` volume option.

Setting volume options for the Snapshot autodelete configuration

When implementing the Snapshot autodelete configuration, you need to set the required volume space guarantee, autosize, fractional reserve, `try_first`, and Snapshot copy options.

Before you begin

Volumes must be created according to the guidelines in the *Data ONTAP Storage Management Guide for 7-Mode*. For information about options related to Snapshot copies, see the *Data ONTAP*

Data Protection Online Backup and Recovery Guide for 7-Mode and for information about volume options, see the *Data ONTAP Storage Management Guide for 7-Mode*.

Steps

1. Set the space guarantee on the volumes by entering the following command:

```
vol options vol_name guarantee volume
```

2. Ensure that autosize is disabled by entering the following command:

```
vol autosize vol_name off
```

Note: This option is disabled by default.

3. Set fractional reserve to zero percent, if it is not already set to that, by entering the following command:

```
vol options vol_name fractional_reserve 0
```

4. Set the Snapshot copy reserve to zero percent by entering the following command:

```
snap reserve vol_name 0
```

The Snapshot copy space and application data are now combined into one large storage pool.

5. Configure Snapshot copies to begin being automatically deleted when the volume reaches the capacity threshold percentage by entering the following command:

```
snap autodelete vol_name trigger volume
```

Note: The capacity threshold percentage is based on the size of the volume. For more details, see the *Data ONTAP Data Protection Online Backup and Recovery Guide for 7-Mode*.

6. Set the try_first option to snap_delete by entering the following command:

```
vol options vol_name try_first snap_delete
```

This enables Data ONTAP to begin deleting Snapshot copies, starting with the oldest first, to free up space for application data.

7. Activate the snap autodelete settings by entering the following command:

```
snap autodelete vol_name on
```

8. Create your space-reserved LUNs.

Related tasks

[Setting up LUNs and igroups using individual commands](#) on page 45

Related information

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

Setting up LUNs and igroups

There are three high-level steps involved in the storage provisioning process: creating LUNs, creating igroups, and mapping the LUNs to the igroups. Several methods are available for completing this process.

lun setup command	This method prompts you through the process of creating a LUN, creating an igroup, and mapping the LUN to the igroup.
System Manager Application	System Manager provides a LUN Wizard that steps you through the process of creating and mapping new LUNs. You can use this method to create one or more LUNs and igroups in any order.
Individual commands	Entering a series of individual commands (such as <code>lun create</code> , <code>igroup create</code> , and <code>lun map</code>).

Related tasks

[Setting up LUNs and igroups using the LUN setup program](#) on page 44

[Setting up LUNs and igroups using individual commands](#) on page 45

Setting up LUNs and igroups using the LUN setup program

LUN setup is a guided program that prompts you for the information needed to create a LUN and an igroup, and to map the LUN to the igroup. When a default is provided in brackets in the prompt, you should press Enter to accept it.

Before you begin

- The volumes for storing LUNs must be created.
- qtrees must be created if you want to use them.
- The LUN type must be specified.

About this task

After the LUN is created, you cannot modify the LUN host operating system type.

Steps

1. On the storage system command line, enter the following command:

```
lun setup
```

The lun setup program is started.

2. Follow the prompts to complete the setup process.

Setting up LUNs and igroups using individual commands

Instead of using LUN setup, you can use individual commands to create LUNs, create igroups, and map the LUNs to the appropriate igroups.

Before you begin

The LUN type must be specified.

About this task

After the LUN is created, you cannot modify the LUN host operating system type.

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

Steps

1. Create a space-reserved LUN by entering the following command on the storage system command line:

```
lun create -s size -t ostype lun_path
```

-s *size* indicates the size of the LUN to be created, in bytes by default.

-t *ostype* indicates the LUN type. The LUN type refers to the operating system type, which determines the geometry used to store data on the LUN.

lun_path is the LUN's path name that includes the volume and qtree.

Example

The following example command creates a 5-GB LUN called /vol/vol2/qtree1/lun3 that is accessible by a Windows host. Space reservation is enabled for the LUN.

```
lun create -s 5g -t windows_2008 /vol/vol2/qtree1/lun3
```

2. Create an igroup by entering the following command on the storage system command line:

```
igroup create {-i | -f} -t ostype initiator_group [node ...]
```

-i specifies that the igroup contains iSCSI node names.

-f specifies that the igroup contains FCP WWPNS.

-t *ostype* indicates the operating system type of the initiator. The values are *solaris*, *Solaris_efi*, *windows*, *windows_gpt*, *windows_2008*, *hpux*, *aix*, *linux*, *netware*, *vmware*, *xen*, and *hyper_v*.

initiator_group is the name you specify as the name of the igroup.

node is a list of iSCSI node names or FCP WWPNS, separated by spaces.

Example

iSCSI example:

```
igroup create -i -t windows_2008 win_host5_group2 iqn.  
1991-05.com.microsoft:host5.domain.com
```

FCP example:

```
igroup create -f -t aix aix-igroup3 10:00:00:00:0c:2b:cc:92
```

3. Map the LUN to an igroup by entering the following command on the storage system command line:

```
lun map lun_path initiator_group [lun_id]
```

lun_path is the path name of the LUN you created.

initiator_group is the name of the igroup you created.

lun_id is the identification number that the initiator uses when the LUN is mapped to it. If you do not enter a number, Data ONTAP generates the next available LUN ID number.

Example

The following command maps /vol/vol1/qtrees1/lun3 to the igroup win_host5_group2 at LUN ID 0:

```
lun map /vol/vol2/qtrees1/lun3 win_host5_group2 0
```

Related concepts

[LUN size](#) on page 50

[ostype \(LUN multiprotocol type\) guidelines](#) on page 49

[What igroups are](#) on page 63

[ostype \(LUN multiprotocol type\) guidelines](#) on page 49

Creating LUNs on vFiler units

Except when using SnapDrive, the process for creating LUNs on vFiler units is slightly different from the process of creating LUNs on storage systems. SnapDrive can create, connect to, and manage LUNs on the vFiler units in the same way it does on the physical storage system.

Before you begin

- The vFiler units must be created. To use vFiler units, you must have MultiStore. For more information about MultiStore and creating vFiler units, see the *Data ONTAP MultiStore Management Guide for 7-Mode*.
- The iSCSI license must be enabled in order for each vFiler unit to manage LUNs on a per vFiler unit basis.

Note: vFilers only work with iSCSI. vFilers do not work with FCP.

About this task

You should use the following guidelines when creating LUNs on vFiler units:

- The vFiler unit access rights are enforced when the storage system processes iSCSI host requests.
- LUNs inherit vFiler unit ownership from the storage unit on which they are created. For example, if /vol/vfstore/vf1_0 is a qtree owned by vFiler unit vf1, all LUNs created in this qtree are owned by vf1.
- As vFiler unit ownership of storage changes, so does ownership of the storage's LUNs.

You can issue LUN subcommands using the following methods:

- From the default vFiler unit (vfiler0) on the hosting storage system, you can do the following:
 - You can enter the `vfiler run * lun` subcommand, which runs the `lun` subcommand on all vFiler units.
 - You can run a LUN subcommand on a specific vFiler unit. To access a specific vFiler unit, you change the vFiler unit context by entering the following commands:

```
filer> vfiler context vfiler_name
```

```
vfiler_name@filer> lun subcommand
```

- From non-default vFiler units, you can enter `vfiler run * lun` command.

Step

1. Enter the `lun create` command in the vFiler unit context that owns the storage, as follows:

```
vfiler run vfiler_name lun create -s 2g -t os_type /vol/vfstore/vf1_0/  
lun0
```

Example

The following command creates a LUN on a vFiler unit at /vol/vfstore/vf1_0:

```
vfiler run vf1 lun create -s 2g -t windows_2008 /vol/vfstore/vf1_0/  
lun0
```

Related information

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

Displaying vFiler LUNs

You might need to display all LUNs owned by a vFiler context. The command for displaying vFiler LUNs is slightly different from the command used on other storage systems.

Step

1. Enter the following command from the vFiler unit that contains the LUNs:

```
vfiler run * lun show
```

Result

The following information shows sample output:

```
system1> vfiler run * lun show
==== vfiler0

/vol/vfstore/vf0_0/vf0_lun0    2g    (21437483648)    (r/w, online)
/vol/vfstore/vf0_0/vf0_lun1    2g    (21437483648)    (r/w, online)

==== vfiler1

/vol/vfstore/vf0_0/vf1_lun0    2g    (21437483648)    (r/w, online)
/vol/vfstore/vf0_0/vf1_lun1    2g    (21437483648)    (r/w, online)
```

LUN configuration

After configuring your volume, you can configure your LUNs. You will need to follow certain guidelines and gather specific information to configure your LUNs.

Information required to create a LUN

When you create a LUN, you must specify the path name of the LUN, the name of the LUN, the LUN Multiprotocol Type (also called ostype), the LUN size, the LUN description, the LUN identification number, and the space reservation setting.

Path name of the LUN

The path name of a LUN must be at the root level of the qtree or volume in which the LUN is located.

You should not create LUNs in the root volume. The default root volume is `/vol/vol0`.

For HA configurations, you should distribute LUNs across the HA pairs.

Note: You might find it useful to provide a meaningful path name for the LUN. For example, you might choose a name that describes how the LUN is used, such as the name of the application, the

type of data that it stores, or the name of the user accessing the data. Examples are /vol/database/lun0, /vol/finance/lun1, and /vol/bill/lun2.

Name of the LUN

The name of the LUN is case-sensitive and can contain 1 to 255 characters. You cannot use spaces. LUN names must use only specific letters and characters.

LUN names can contain only the letters A through Z, a through z, numbers 0 through 9, hyphen (“-”), underscore (“_”), left brace (“{”), right brace (“}”), and period (“.”).

ostype (LUN multiprotocol type) guidelines

The ostype (sometimes called LUN multiprotocol type) specifies the OS of the host accessing the LUN. It also determines the layout of data on the LUN, the geometry used to access that data, and the minimum and maximum size of the LUN.

Not all Data ONTAP versions support all LUN multiprotocol types. You should consult the Interoperability Matrix to get the most up-to-date information. The ostype (LUN multiprotocol type) options and when each should be used are listed below:

Note: If you are using SnapDrive for Windows, the LUN multiprotocol type is automatically set.

solaris	Use if your host operating system is Solaris and you are not using Solaris EFI labels.
Solaris_efi	Use if you are using Solaris EFI labels. Note: Using any other LUN multiprotocol type with Solaris EFI labels might result in LUN (mis)alignment problems. For more information, see your Solaris Host Utilities documentation and release notes.
windows	Use if your host operating system is Windows 2000 Server, Windows XP, or Windows Server 2003 using the MBR partitioning method.
windows_gpt	Use if you want to use the GPT partitioning method and your host is capable of using it. Windows Server 2003, Service Pack 1 and later are capable of using the GPT partitioning method, and all 64-bit versions of Windows support it.
windows_2008	Use if your host operating system is Windows Server 2008 or Windows Server 2012; both MBR and GPT partitioning methods are supported.
hpux	Use if your host operating system is HP-UX.
aix	Use if your host operating system is AIX.
linux	Use if your host operating system is Linux.
netware	Use if your host operating system is Netware.
vmware	Use if you are using ESX Server and your LUNs will be configured with VMFS.

Note: If you configure the LUNs with RDM, you can use the guest operating system as the LUN multiprotocol type.

xen Use if you are using Xen and your LUNs will be configured with Linux LVM with Dom0.

Note: For raw LUNs, you can use the type of guest operating system as the LUN multiprotocol type.

hyper_v Use if you are using Windows Server 2008 or Windows Server 2012 Hyper-V and your LUNs contain virtual hard disks (VHDs). If you are using hyper_v for your LUN type, you should also use hyper_v for your igroup os type.

Note: For raw LUNs, you can use the type of the child operating system as the LUN multiprotocol type.

For information about supported hosts, see the Interoperability Matrix.

Related tasks

[Setting the operating system type for an igroup](#) on page 81

Related information

[NetApp Interoperability Matrix-support.netapp.com/NOW/products/interoperability](http://support.netapp.com/NOW/products/interoperability)

LUN size

You specify the size of a LUN in bytes or by using specific multiplier suffixes.

Multiplier suffixes that can be used are:

Multiplier suffix	Size
c	bytes
w	words or double bytes
b	512-byte blocks
k	kilobytes
m	megabytes
g	gigabytes
t	terabytes

The usable space in the LUN depends on host or application requirements for overhead. For example, 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.

The disk geometry used by the operating system determines the minimum and maximum size values of LUNs. 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, you should consult the vendor's documentation for more information about how disk geometry affects LUN size.

LUN description

The LUN description is an optional attribute you can use to specify additional information about the LUN.

You can edit this description at the command line.

Space reservation setting

When you create a LUN by using the `lun setup` command, you specify whether you want to enable space reservations. When you create a LUN using the `lun create` command, space reservation is automatically turned on.

Note: You should keep space reservation on.

Guidelines for LUN layout and space allocation

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

- You should 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 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 you 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.
The data in a backup LUN changes 100 percent for each backup period. 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.
- You can use quotas to allocate space.
For example, you might want to assign volume space to various database administrators and allow them to create and manage their own LUNs. You can organize the volume into qtrees with quotas and enable the individual database administrators to manage the space they have been allocated.
If you organize your LUNs in qtrees with quotas, ensure that the quota limit can accommodate the sizes of the LUNs you want to create. Data ONTAP does not allow you to create a LUN in a qtree with a quota if the LUN size exceeds the quota.

LUN management

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 the command-line interface (CLI) to manage LUNs.

Displaying command-line Help for LUNs

You can use the `lun help` command to display online Help for all LUN commands and sub-commands.

Steps

1. On the storage system's command line, enter the following command:

```
lun help
```

A list of all the LUN subcommands is displayed:

```
lun help          - List LUN (logical unit of block storage) commands
lun config_check  - Check all lun/igroup/fcp settings for correctness
lun clone         - Manage LUN cloning
lun comment       - Display/Change descriptive comment string
lun create        - Create a LUN
lun destroy       - Destroy a LUN
lun map           - Map a LUN to an initiator group
lun maxsize       - Show the maximum possible size of a LUN on a given
volume or qtree
lun move          - Move (rename) LUN
lun offline       - Stop block protocol access to LUN
lun online        - Restart block protocol access to LUN
lun resize        - Resize LUN
lun serial        - Display/change LUN serial number
lun set           - Manage LUN properties
lun setup         - Initialize/Configure LUNs, mapping
lun share         - Configure NAS file-sharing properties
lun show          - Display LUNs
lun snap          - Manage LUN and snapshot interactions
lun stats         - Displays or zeros read/write statistics for LUN
lun unmap         - Remove LUN mapping
```

2. Display the syntax for any of the subcommands by entering the following command:

```
lun help subcommand
```

Example

```
lun help show
```

Controlling LUN availability

You can use the `lun online` and `lun offline` commands to control the availability of LUNs while preserving the LUN mappings.

Bringing LUNs online

You can use the `lun online` command to bring one or more LUNs back online.

About this task

Note: The `lun online` command fails when the cluster interconnect is down to avoid possible LUN mapping conflicts.

Step

1. Enter the following command:

```
lun online lun_path [lun_path ...]
```

Example

```
lun online /vol/vol1/lun0
```

Taking LUNs offline

Taking a LUN offline makes it unavailable for block protocol access. You can use the `lun offline` command to take the LUN offline.

Before you begin

Any host application that is accessing the LUN must be quiesced or synchronized.

About this task

Taking a LUN offline makes it unavailable for block protocol access.

Step

1. Take a LUN offline by entering the following command:

```
lun offline lun_path [lun_path ...]
```

Example

```
lun offline /vol/vol1/lun0
```

Moving LUNs

You can use the `lun move` command to rename or move a LUN.

About this task

If you are organizing LUNs in qtrees, the existing path (`lun_path`) and the new path (`new_lun_path`) must be either in the same qtree or in another qtree in that same volume.

Note: This process is completely nondisruptive; it can be performed while the LUN is online and serving data.

Step

1. Enter the following command:

```
lun move lun_path new_lun_path
```

Example

```
lun move /vol/vol1/mylun /vol/vol1/mynewlun
```

Modifying LUN descriptions

You may have added a LUN description when creating the LUN. You can use the `lun comment` command to modify that description or add a new one.

About this task

If you use spaces in the comment, you must enclose the comment in quotation marks.

Step

1. Enter the following command:

```
lun comment lun_path [comment]
```

Example

```
lun comment /vol/vol1/lun2 "10 GB for payroll records"
```

How LUN reservations work

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 set reservation` 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.

Enabling or disabling space reservations for LUNs

You can use the `lun set reservation` command to enable or disable space reservations for a LUN.

About this task

Attention: If you disable space reservations, write operations to a LUN might fail due to insufficient disk space, and the host application or operating system might crash. When write operations fail, Data ONTAP displays system messages on the console, or sends these messages to log files and other remote systems, as specified by its `/etc/syslog.conf` configuration file.

Steps

1. Display the status of space reservations for LUNs in a volume by entering the following command:

```
lun set reservation lun_path
```

Example

```
lun set reservation /vol/lunvol/hpux/lun0
```

```
Space Reservation for LUN /vol/lunvol/hpux/lun0 (inode 3903199): enabled
```

2. Enter the following command:

```
lun set reservation lun_path [enable | disable]
```

lun_path is the LUN in which space reservations are to be set. This must be an existing LUN.

Note: Enabling space reservation on a LUN fails if there is not enough free space in the volume for the new reservation.

Accessing LUNs with NAS protocols

When you create a LUN, you can only access it with the iSCSI, FC, or FCoE protocol by default. However, you can use NAS protocols to make a LUN available to a host if the NAS protocols are licensed and enabled on the storage system.

About this task

The usefulness of accessing a LUN over NAS protocols depends on the host application. For example, the application must be equipped to understand the format of the data within the LUN and be able to traverse any file system the LUN may contain. Access is provided to the LUN's raw data, but not to any particular piece of data within the LUN.

If you want to write to a LUN using a NAS protocol, you must take the LUN offline or unmap it to prevent an iSCSI or FCP host from overwriting data in the LUN.

Note: A LUN cannot be extended or truncated using NFS or CIFS protocols.

Steps

1. Determine whether you want to read, write, or do both to the LUN using the NAS protocol and take the appropriate action:
 - If you want read access, the LUN can remain online.
 - If you want write access, ensure that the LUN is offline or unmapped.

2. Enter the following command:

```
lun share lun_path {none|read|write|all}
```

Example

```
lun share /vol/vol1/qtree1/lun2 read
```

The LUN is now readable over NAS.

Checking LUN, igroup, and FC settings

You can use the `lun config_check` command to verify a number of LUN, igroup, and FC settings.

About this task

The command performs the following actions:

- Checks whether any FC target interfaces are down.
- Verifies that the ALUA igroup settings are valid.
- Checks for nodename conflicts.
- Checks for igroup and LUN map conflicts.
- Checks for initiators with mixed/incompatible settings.
- Checks for duplicate WWPNs.

- Checks for igroup ALUA conflicts.

Step

1. Enter the following command:

```
lun config_check [-S] [-w] [-s] [-v]
```

- You can use the -S option to only check the single_image cfmode settings.
- You can use the -w option to only check for WWPN conflicts.
- You can use the -s option for silent mode, which only provides output if there are errors.
- You can use the -v option for verbose mode, which provides detailed information about each check.

Related concepts

[What ALUA is](#) on page 204

[igroup ostype](#) on page 66

[How Data ONTAP avoids igroup mapping conflicts during cluster failover](#) on page 128

Displaying LUN serial numbers

A LUN serial number is a unique, 12-byte, ASCII string generated by the storage system. Many multipathing software packages use this serial number to identify redundant paths to the same LUN.

About this task

Although the storage system displays the LUN serial number in ASCII format by default, you can display the serial number in hexadecimal format as well.

Step

1. Enter one of the following commands:

- **lun show [-v] *lun_path***
- **lun serial [-x]*lun_path new_lun_serial***

The -v option displays the serial numbers in ASCII format.

The -x option displays the serial numbers in hexadecimal format.

The *new_lun_serial* changes the existing LUN serial number to the specified serial number.

Note: Under normal circumstances, you should not change the LUN serial number. However, if you do need to change it, ensure that the LUN is offline before issuing the command. Also, you cannot use the -x option when changing the serial number; the new serial number must be in ASCII format.

Example

```
lun serial -x /vol/blocks_fvt/ncmds_lun2
```

```
Serial (hex)#: 0x4334656f476f424f594d2d6b
```

Displaying LUN statistics

You can use the `lun stats` command to display the number of read and write operations and the number of operations per second for LUNs.

Step

1. Enter the following command:

```
lun stats -z -i interval -c count -o [-a | lun_path]
```

-z resets the statistics on all LUNs or the LUN specified in the *lun_path* option.

-i *interval* is the interval, in seconds, at which the statistics are displayed.

-c *count* is the number of intervals. For example, the `lun stats -i 10 -c 5` command displays statistics in ten-second intervals, for five intervals.

-o displays additional statistics, including the number of QFULL messages the storage system sends when its SCSI command queue is full and the amount of traffic received from the partner storage system.

-a shows statistics for all LUNs.

lun_path displays statistics for a specific LUN.

Example

```
system1>lun stats -o -i 1
Read   Write Other QFull Read  Write Average Queue Partner  Lun
Ops    Ops  Ops   Ops  kB    kB   Latency Length  Ops  kB
0      351   0     0    0   44992   11.35   3.00   0    0  /vol/tpcc/
log_22
0      233   0     0    0   29888   14.85   2.05   0    0  /vol/tpcc/
log_22
0      411   0     0    0   52672   8.93    2.08   0    0  /vol/tpcc/
log_22
2       1    0     0   16     8    1.00    1.00   0    0  /vol/tpcc/
ctrl_0
1       1    0     0    8     8    1.50    1.00   0    0  /vol/tpcc/
ctrl_1
0      326   0     0    0   41600   11.93   3.00   0    0  /vol/tpcc/
log_22
0      353   0     0    0   45056   10.57   2.09   0    0  /vol/tpcc/
log_22
0      282   0     0    0   36160   12.81   2.07   0    0  /vol/tpcc/
log_22
```

Displaying LUN mapping information

You can use the `lun show -m` command to display a list of LUNs and the hosts to which they are mapped.

Step

1. On the storage system's command line, enter the following command:

```
lun show -m
```

Example

```
system1>lun show -m
LUN path                Mapped to      LUN ID  Protocol
-----
/vol/tpcc/ctrl_0        host5          0       iSCSI
/vol/tpcc/ctrl_1        host5          1       iSCSI
/vol/tpcc/crash1        host5          2       iSCSI
/vol/tpcc/crash2        host5          3       iSCSI
/vol/tpcc/cust_0        host6          4       iSCSI
/vol/tpcc/cust_1        host6          5       iSCSI
/vol/tpcc/cust_2        host6          6       iSCSI
```

Displaying detailed LUN information

You can use the `lun show -v` command to show additional LUN details, such as the serial number, ostype (multiprotocol type), and maps.

Step

1. On the storage system's command line, enter the following command to display LUN status and characteristics:

```
lun show -v
```

Example

```
system1>lun show -v
/vol/vol1/lun1      4m (4194304)  (r/w, online)
Serial#: BYjB3?-iq3hU
Share: none
Space Reservation: enabled
Multiprotocol Type: linux
Occupied Size:      0 (0)
Creation Time: Tue Aug 30 09:58:48 GMT 2011
Cluster Shared Volume Information: 0x0
```

Displaying hidden staging area LUNs

You can use the `lun show staging` command to obtain a list of all the hidden staging area LUNs. If you want to destroy an igroup to which the staging LUN is mapped, the `lun show staging` command indicates the reason for not being able to destroy an igroup.

About this task

The staging area LUNs are temporarily stored in `/vol/volnam/Staging_XXXX/lun_name` path when a nondisruptive restore is in progress and are automatically cleared when the restore completes successfully. If the nondisruptive restore fails, you should destroy the temporary LUNs manually using the `lun destroy` command.

Step

1. Obtain the list of hidden staging area LUNs by entering the following command:

```
lun show staging
```

Example: Hidden staging area LUNs

```
system1> lun show -v staging
/vol/volz/Staging_123/lun0      10m (10485760)  (r/w, online, mapped)
  Comment: "staging lun"
  Serial#: 1BbFb+8rmk/f
  Share: none
  Space Reservation: enabled
  Multiprotocol Type: linux
  Maps: gaston=1
```

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 information

Best Practices for File System Alignment in Virtual Environments: media.netapp.com/documents/tr-3747.pdf

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

Removing LUNs

You can use the `lun destroy` command to remove one or more LUNs.

Before you begin

Without the `-f` parameter, the LUN must be taken offline and unmapped, and then the `lun destroy` command must be entered.

Step

1. Remove one or more LUNs by entering the following command:

```
lun destroy [-f] lun_path [lun_path ...]
```

`-f` forces the `lun destroy` command to execute even if the LUNs specified by one or more `lun_paths` are mapped or are online.

Misaligned I/O can occur on properly aligned LUNs

Data ONTAP can report alignment problems with properly aligned LUNs. A certain amount of misaligned I/O occurs on properly aligned LUNs and can be harmless.

For proper performance, it is important that LUNs be correctly aligned with the underlying WAFL block boundaries. For physical hosts, selecting the correct `ostype` value when creating the LUN is usually sufficient. For virtualized server environments, an additional alignment step is sometimes needed.

However, the alignment detection features of Data ONTAP sometimes report misalignment for properly aligned LUNs because Data ONTAP detects misaligned I/O.

The presence of misaligned I/O does not necessarily mean the LUN or datastore itself is misaligned. Certain types of applications, primarily databases, generate misaligned write I/O that does not harm performance.

The performance problem with misaligned I/O is that additional reads are required to complete a write. For example, a 4 KB block that straddles a WAFL block boundary requires an update to parts

of two WAFL blocks. Before the misaligned write is committed to disk, two partial reads are required to retrieve the data which is not being changed and to allow Data ONTAP to complete the write. These additional reads increase the I/O load on the storage system and can increase latency as seen by the host system. An aligned 4 KB block requires only an overwrite.

Databases and other applications which incrementally update a log file with variable sized blocks sometimes trigger misalignment warnings. Although these applications sometimes perform a 4 KB misaligned write, the next write performed completes the partial block. As the sequential write continues, each new block written eliminates the need for any partial reads caused by prior write operations.

In general, misalignment warnings on LUNs or datastores can be disregarded as long you are confident that the LUN or datastore is properly provisioned and any partitioning table is correct.

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

igroup management

To manage your initiator groups (igroups), you can perform a range of tasks, including creating igroups, destroying them, and renaming them.

Related concepts

[What igroups are](#) on page 63

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 otypes. Also, a given igroup can be used for FC protocol or iSCSI, but not both.

Related concepts

[igroup management](#) on page 63

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

Creating igroups

Initiator groups, or igroups, are tables of host identifiers such as Fibre Channel WWPNs and iSCSI node names. You can use igroups to control which hosts can access specific LUNs.

Step

1. Create an igroup by entering the following command:

```
igroup create [-i | -f] -t ostype initiator_group [nodename ... |  
WWPN ...] [wwpn alias ...] [-a portset]
```

-i indicates that it is an iSCSI igroup.

-f indicates that it is an FC igroup.

-t *ostype* indicates the operating system of the host. The values are *solaris*, *Solaris_efi*, *windows*, *windows_gpt*, *windows_2008*, *hpux*, *aix*, *linux*, *netware*, *vmware*, *xen*, and *hyper_v*.

initiator_group is the name you give to the igroup.

nodename is an iSCSI node name. You can specify more than one node name.

WWPN is the FC worldwide port name. You can specify more than one WWPN.

wwpn alias is the name of the alias you created for a WWPN. You can specify more than one alias.

-a *portset* applies only to FC igroups. This binds the igroup to a port set. A port set is a group of target FC ports. When you 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.

Example

```
igroup create -i -t windows_2008 win-group0 iqn.1991-05.com.microsoft:eng1
```

Creates an iSCSI igroup called win-group0 that contains the node name of the Windows host associated with that node name.

Related concepts

[How to use port sets to make LUNs available on specific FC target ports](#) on page 130

[What igroups are](#) on page 63

[ostype \(LUN multiprotocol type\) guidelines](#) on page 49

Required information for creating igroups

There are a number of attributes required when creating igroups, including the name of the igroup, type of igroup, ostype, iSCSI node name for iSCSI igroups, and WWPN for FCP igroups.

igroup name

The igroup name is a case-sensitive name that must satisfy several requirements.

The igroup name:

- Contains 1 to 96 characters. Spaces are not allowed.
- Can contain the letters A through Z, a through z, numbers 0 through 9, hyphen (“-”), underscore (“_”), colon (“:”), and period (“.”).
- Must start with a letter or number.

The name you assign to an igroup is independent of the name of the host that is used by the host operating system, host files, or Domain Name Service (DNS). If you name an igroup aix1, for example, it is not mapped to the actual IP host name (DNS name) of the host.

Note: You might find it useful to provide meaningful names for igroups, ones that describe the hosts that can access the LUNs mapped to them.

igroup type

The igroup type can be either `-i` for iSCSI or `-f` for FC.

igroup ostype

The ostype indicates the type of host operating system used by all of the initiators in the igroup. All initiators in an igroup must be of the same ostype. The ostyles of initiators are `solaris`, `windows`, `hpux`, `aix`, `netware`, `xen`, `hyper_v`, `vmware`, and `linux`.

You must select an ostype for the igroup.

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 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 tasks

[Creating FC protocol igroups on UNIX hosts using the sanlun command](#) on page 66

Related information

Host Utilities documentation: www.support.netapp.com

Creating FC protocol igroups on UNIX hosts using the sanlun command

If you have a UNIX host, you can use the `sanlun` command to create FC protocol igroups. The command obtains the host's WWPNs and prints out the `igroup create` command with the correct arguments. Then you can copy and paste this command into the storage system command line.

Steps

1. Ensure that you are logged in as root on the host.

2. Change to the `/opt/netapp/santools/bin` directory.
3. Enter the following command to print a command to be run on the storage system that creates an igroup containing all the HBAs on your host:

```
./sanlun fcp show adapter -c
```

`-c` prints the full `igroup create` command on the screen.

The relevant `igroup create` command is displayed:

```
Enter this controller command to create an initiator group for this
system:
igroup create -f -t solaris "ssan-280r-15" 21000003ba14a568
10000000c9580da0
```

In this example, the name of the host is **hostA**, so the name of the igroup with the two WWPNS is **hostA**.

4. Create a new session on the host and use the `telnet` command to access the storage system.
5. Copy the `igroup create` command from Step 3, paste the command on the storage system's command line, and press Enter to run the `igroup` command on the storage system.

An igroup is created on the storage system.

6. On the storage system's command line, enter the following command to verify the newly created igroup:

```
igroup show
```

Example

```
systemX> igroup show
hostA (FCP) (ostype: aix):
    10:00:00:00:AA:11:BB:22
    10:00:00:00:AA:11:EE:33
```

The newly created igroup with the host's WWPNS is displayed.

Creating igroups for a non-default vFiler unit

You can create iSCSI igroups for non-default vFiler units. With vFiler units, igroups are owned by vFiler contexts. The vFiler ownership of igroups is determined by the vFiler context in which the igroup is created.

Steps

1. Change the context to the desired vFiler unit by entering the following command:

```
vfiler context vf1
```

The vFiler unit's prompt is displayed.

2. Create the igroup on the vFiler unit determined in step 1 by entering the following command:
igroup create -i vf1_iscsi_group iqn.1991-05.com.microsoft:server1
3. Display the igroup by entering the following command:
igroup show

The following information is displayed:

```
vf1_iscsi_group (iSCSI) (ostype: windows_2008):
    iqn.1991-05.com.microsoft:server1
```

After you finish

You must map LUNs to igroups that are in the same vFiler unit.

igroup configuration

igroups can be configured for various settings such as ALUA and report_scsi_name. You can also configure throttles for your igroups to limit and control other parameters of the igroup.

Enabling ALUA

You can enable ALUA for your igroups, as long as the host supports the ALUA standard.

About this task

Only FCP igroups support ALUA.

If ALUA is not enabled for your igroup, you can manually enable it by setting the alua option to yes. If you map multiple igroups to a LUN and you enable one of the igroups for ALUA, you must enable all the igroups for ALUA.

Steps

1. Check whether ALUA is enabled by entering the following command:
igroup show -v *igroup_name*

Example

igroup show -v igroup1

```
system1> igroup show -v igroup1
igroup1:
OS Type: solaris
Member: 10:00:00:00:c9:2b:cc:39 (logged in on: vtic, 5a, 5b)
Member: 10:00:00:00:c9:2b:cb:7e
ALUA: Yes
Report SCSI Name in Inquiry Descriptor: No
```

Note: The output of `igroup show -v` displays the FCP initiator logged in on physical ports as well as a port called “vtic”. VTIC is an abbreviation for “virtual target interconnect”. VTIC provides a connection between the two nodes in an HA pair, enabling LUNs to be served through target ports on both nodes. It is normal to see VTIC as one of the ports in the output of `igroup show -v`.

2. Enter the following command to enable ALUA if it has not already been enabled:

```
igroup set igroup alua yes
```

Related concepts

[What ALUA is](#) on page 204

Related tasks

[Checking LUN, igroup, and FC settings](#) on page 56

Enabling report_scsi_name

You can enable `report_scsi_name` for your igroups to control reporting or hiding the new inquiry descriptor to the initiators.

When report_scsi_name is automatically enabled

The newly implemented inquiry descriptor should not be reported to Windows based initiator groups by default. For all other otypes, such as Linux, HP-UX, and AIX the newly implemented descriptor is reported by default. This behavior of the descriptor is controlled by `report_scsi_name` attribute.

The default value of the attribute `report_scsi_name` is NO for all initiator groups with otype as windows. Otherwise, for all initiator groups with otype AIX, HP-UX, or Linux, the default value of the attribute `report_scsi_name` is YES.

You can modify the `report_scsi_name` attribute to YES or NO manually too.

Related tasks

[Manually setting the report_scsi_name option to yes](#) on page 69

Manually setting the report_scsi_name option to yes

You can set or unset the `report_scsi_name` attribute to control reporting or hiding the new inquiry descriptor to the initiators.

Steps

1. Check whether `report_scsi_name` is enabled by entering the following command:

```
igroup show -v igroup_name
```

Example

```
system1> igroup show -v
fcplnx (FCP):
OS Type: linux
Member: 21:00:00:24:ff:17:d7:11 (not logged in)
Member: 10:00:00:00:d9:e6:c1:b1 (logged in on: 0a)
UUID: ab7b40ac-917c-17e0-b240-123478563412
ALUA: Yes
Report SCSI Name in Inquiry Descriptor: NO
```

Note: The output of `igroup show -v` displays the FCP initiator logged in on physical ports as well as a port called “vtic”. VTIC is an abbreviation for “virtual target interconnect.” VTIC provides a connection between the two nodes in an HA pair, enabling LUNs to be served through target ports on both nodes. It is normal to see VTIC as one of the ports in the output of `igroup show -v`.

2. Enable `report_scsi_name` by entering the following command:

```
igroup set igroup_name report_scsi_name yes
```

Related information

What does the order of listed target ports mean in the output of `igroup show`? - <https://kb.netapp.com/support/index?page=content&id=3011742>

Fibre Channel initiator request management

Data ONTAP implements a mechanism called `igroup` throttles, which you can use to ensure that critical initiators are guaranteed access to the queue resources and that less-critical initiators are not flooding the queue resources.

How Data ONTAP manages Fibre Channel initiator requests

When you use `igroup` throttles, Data ONTAP calculates the total amount of command blocks available and allocates the appropriate number to reserve for an `igroup`, based on the percentage you specify when you create a throttle for that `igroup`.

Data ONTAP does not allow you to reserve more than 99 percent of all the resources. The remaining command blocks are always unreserved and are available for use by `igroups` without throttles.

How to use `igroup` throttles

You can use `igroup` throttles to specify what percentage of the queue resources they can reserve for their use.

For example, if you set an `igroup`’s throttle to be 20 percent, then 20 percent of the queue resources available at the storage system’s ports are reserved for the initiators in that `igroup`. The remaining 80 percent of the queue resources are unreserved. In another example, if you have four hosts and they are in separate `igroups`, you might set the `igroup` throttle of the most critical host at 30 percent, the

least critical at 10 percent, and the remaining two at 20 percent, leaving 20 percent of the resources unreserved.

You can use igroup throttles to perform the following tasks:

- You can create one igroup throttle per igroup, if desired.
- Note:** Any igroups without a throttle share all the unreserved queue resources.
- You can assign a specific percentage of the queue resources on each physical port to the igroup.
 - You can reserve a minimum percentage of queue resources for a specific igroup.
 - You can restrict an igroup to a maximum percentage of use.
 - You can allow an igroup throttle to exceed its limit by borrowing from these resources:
 - The pool of unreserved resources to handle unexpected I/O requests
 - The pool of unused reserved resources, if those resources are available

How failover affects igroup throttles

Throttles manage physical ports, so during a takeover, their behavior is important to understand. Throttles apply to all ports and are divided by two when the HA pair is in takeover mode.

Creating igroup throttles

You can use igroup throttles to limit the number of concurrent I/O requests an initiator can send to the storage system, prevent initiators from flooding a port, and ensure that specific initiators have guaranteed access to the queue resources.

Step

1. Enter the following command:

```
igroup set igroup_name throttle_reserve percentage
```

Example

```
igroup set aix-igroup1 throttle_reserve 20
```

The igroup throttle is created for aix-igroup1, and it persists through reboots.

Destroying igroup throttles

You can destroy an igroup throttle by setting the throttle reserve to zero.

Step

1. Enter the following command:

```
igroup set igroup_name throttle_reserve 0
```

Borrowing queue resources from the unreserved pool

If queue resources are available in the unreserved pool, you can borrow resources from the pool for a particular igroup.

Step

1. To define whether an igroup can borrow queue resources from the unreserved pool, enter the following command:

```
igroup set igroup_name throttle_borrow [yes|no]
```

Note: The default when you create an igroup throttle is no.

Example

```
igroup set aix-igroup1 throttle_borrow yes
```

When you set the `throttle_borrow` setting to yes, the percentage of queue resources used by the initiators in the igroup might be exceeded if resources are available.

Displaying throttle information

You can use the `igroup show -t` command to display important information about the throttles assigned to igroups.

Step

1. Enter the following command:

```
igroup show -t
```

Example

```
system1>igroup show -t
      name      reserved    exceeds    borrows
aix-igroup1      20%         0         N/A
aix-igroup2      10%         0         0
```

The *exceeds* column displays the number of times the initiator sends more requests than the throttle allows. The *borrows* column displays the number of times the throttle is exceeded and the storage system uses queue resources from the unreserved pool. In the borrows column, *N/A* indicates that the igroup `throttle_borrow` option is set to no.

Displaying igroup throttle usage

You can display real-time information about how many command blocks the initiator in the igroup is using, as well as the number of command blocks reserved for the igroup on the specified port.

Step

1. Enter the following command:

```
igroup show -t -i interval -c count [igroup] -a]
```

-t displays information on igroup throttles.

-i *interval* displays statistics for the throttles over an interval in seconds.

-c *count* determines how many intervals are shown.

igroup is the name of a specific igroup for which you want to show statistics.

-a displays statistics for all igroups, including idle igroups.

Example

```
system1> igroup show -t -i 1
```

name	reserved	4a	4b	5a	5b
igroup1	20%	45/98	0/98	0/98	0/98
igroup2	10%	0/49	0/49	17/49	0/49
unreserved		87/344	0/344	112/344	0/344

The first number under the port name indicates the number of command blocks the initiator is using. The second number under the port name indicates the number of command blocks reserved for the igroup on that port.

In this example, the display indicates that igroup1 is using 45 of the 98 reserved command blocks on adapter 4a, and igroup2 is using 17 of the 49 reserved command blocks on adapter 5a.

igroups without throttles are counted as unreserved.

Displaying LUN statistics on exceeding throttles

Statistics are available about I/O requests for LUNs that exceed the igroup throttle. These statistics can be useful for troubleshooting and monitoring performance.

Steps

1. Enter the following command:

```
lun stats -o -i time_in_seconds
```

-i *time_in_seconds* is the interval over which performance statistics are reported. For example, -i 1 reports statistics each second.

-o displays additional statistics, including the number of QFULL messages, or "QFULLS".

Example

lun stats -o -i 1 /vol/vol0/lun1

```
system1> lun stats -o -i 1 /vol/vol0/lun1
Read  Write Other QFull Read  Write Average Queue Partner  Lun
Ops   Ops   Ops    Full  kB   kB   Latency Length  Ops  kB   /vol/
0     5108   0      0     0    20432  0.62    6.00    0   0   /vol/
vol0/lun1
---
0     7555   0      0     0    30220  0.00    5.05    0   0   /vol/
vol0/lun1
---
0     7535   0      0     0    30144  0.01    5.05    0   0   /vol/
vol0/lun1
---
0     5599   0      0     0    22396  0.38    5.08    0   0   /vol/
vol0/lun1
---
0     6847   0      0     0    27384  0.16    5.07    0   0   /vol/
vol0/lun1
---
0     7460   0      0     0    29836  0.01    5.05    0   0   /vol/
vol0/lun1
---
0     7461   0      0     0    29844  0.01    5.05    0   0   /vol/
vol0/lun1
---
0     4962   0      0     0    19848  0.64    6.00    0   0   /vol/
vol0/lun1
---
0     7379   0      0     0    29516  0.05    5.05    0   0   /vol/
vol0/lun1
---
0     7482   0      0     0    29924  0.01    5.05    0   0   /vol/
vol0/lun1
---
0     7416   0      0     0    29664  0.02    5.05    0   0   /vol/
vol0/lun1
```

The output displays performance statistics, including the QFULL column. This column indicates the number of initiator requests that exceeded the number allowed by the igroup throttle, and as a result, received the SCSI Queue Full response.

2. Display the total count of QFULL messages sent for each LUN by entering the following command:

lun stats -o *lun_path*

Example

```
system1> lun stats -o /vol/vol0/lun1
/vol/vol0/lun1 (11 hours, 19 minutes, 0 seconds)
Read(kbytes) Write(kbytes) Read Ops Write Ops Other Ops QFulls
Partner Ops Partner KBytes
      488      4875956      60      1218939      84      0
83      448
```

LUN and igroup mapping

Before you can use your LUN it must be mapped to an igroup.

What LUN mapping is

LUN mapping is the process of associating a LUN with an igroup. When you map the LUN to the igroup, you grant the initiators in the igroup access to the LUN.

Required information for mapping a LUN to an igroup

You must map a LUN to an igroup to make the LUN accessible to the host. Data ONTAP maintains a separate LUN map for each igroup to support a large number of hosts and to enforce access control.

LUN name

Specify the path name of the LUN to be mapped.

igroup name

Specify the name of the igroup that contains the hosts that will access the LUN.

LUN identification number

A LUN must have a unique identification number (ID) so that the host can identify and access the LUN. You map the LUN ID to an igroup so that all the hosts in that igroup can access the LUN.

If you do not specify a LUN ID, Data ONTAP automatically assigns one.

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 have additional considerations for LUN identification numbers.

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

If you are attempting to map a LUN when the cluster interconnect is down, you must not include a LUN ID, because the partner system will have no way of verifying that the LUN ID is unique. Data ONTAP reserves a range of LUN IDs for this purpose and automatically assigns the first available LUN ID in this range.

- If you are mapping the LUN from the primary system, Data ONTAP assigns a LUN in the range of 193 to 224.
- If you are mapping the LUN from the secondary system, Data ONTAP assigns a LUN in the range of 225 to 255.

Related information

Host Utilities documentation: www.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 two different LUNs with the same LUN ID to two different igroups without having a conflict, provided that the igroups do not share any initiators or only one of the LUNs is online at a given time.
- You should ensure that the LUNs are online before mapping them to an igroup. You should not map LUNs that are in the offline state.
- 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.

SnapMirror destinations and read-only LUNs

When a qtree or volume containing LUNs is used as a SnapMirror source, the LUNs copied to the SnapMirror destination appear as read-only LUNs to the destination storage system. However, in prior versions of Data ONTAP, you could not manage these LUNs as long as the SnapMirror relationship was intact. In addition, you can manage LUN maps for LUNs on mirrored qtrees and volumes.

In prior versions of Data ONTAP, LUN maps created at the source location were copied to the destination storage system.

As a result, the LUNs appear as unmapped and read-only. Therefore, you must explicitly map these read-only LUNs to the hosts at the destination. Once you map the LUNs to the host, the LUNs remain online, even after the SnapMirror relationship is broken.

You map these LUNs to the host in the same way that you map any other LUNs to a host.

The destination LUN is also assigned a new serial number. The online/offline status is inherited from the source LUN and cannot be changed on the destination LUN. The only operations allowed on

read-only LUNs are `lun map`, `lun unmap`, `lun show`, `lun stats`, and changes to SCSI-2 reservations and SCSI-3 persistent reservations.

You can create new igroups on the destination, map the destination LUN to those igroups, or use any existing igroups. After you set up the LUN maps for the destination LUN, you can continue to use the LUN, regardless of the current mirror relationship.

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.

Also note that you cannot create LUNs on read-only qtrees or volumes. The LUNs that display in a mirrored destination inherit the read-only property from the container.

For more information about read-only LUNs and SnapMirror, see the *Data ONTAP Data Protection Online Backup and Recovery Guide for 7-Mode*.

How to make LUNs available on specific FC target ports

When you map a LUN to a FC igroup, the LUN is available on all of the storage system's FC target ports if the igroup is not bound to a port set. A port set consists of a group of FC target ports.

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

You can define port sets for FC target ports only. You should not use port sets for iSCSI target ports.

Related concepts

[How to use port sets to make LUNs available on specific FC target ports](#) on page 130

Unmapping LUNs from igroups

You might need to occasionally unmap a LUN from an igroup. After you take the LUN offline, you can use the `lun unmap` command to unmap the LUN.

About this task

You need to unmap the LUN and bring the LUN back online to map it to a different host. This prevents any data corruption if the host tries to do some I/O.

Steps

1. Enter the following command:

```
lun offline lun_path
```

Example

```
lun offline /vol/vol1/lun1
```

- 2. Enter the following command:

```
lun unmap lun_path igroup
```

Example

```
lun unmap /vol/vol1/lun1 solaris-igroup0
```

- 3. Bring the LUN back online:

```
lun online lun_path [lun_path ...]
```

Example

```
lun online /vol/vol1/lun1
```

Deleting igroups

When deleting igroups, you can use a single command to simultaneously remove the LUN mapping and delete the igroup. You can also use two separate commands to unmap the LUNs and delete the igroup.

Step

- 1. Delete one or more igroups by completing one of the following steps.

If you want to...	Then enter this command...
Remove LUN mappings before deleting the igroup	<pre>lun unmap lun-path igroup</pre> then <pre>igroup destroy igroup1 [igroup2, igroup3...]</pre>
Remove all LUN maps for an igroup and delete the igroup with one command	<pre>igroup destroy -f igroup1 [igroup2, igroup3...]</pre>

Example

```
lun unmap /vol/vol12/qtree/LUN10 win-group5
then
igroup destroy win-group5
```

Example

```
igroup destroy -f win-group5
```

Adding initiators to an igroup

You can use the `igroup add` command to add initiators to an igroup.

About this task

An initiator cannot be a member of two igroups of differing types. For example, if you have an initiator that belongs to a Solaris igroup, Data ONTAP does not allow you to add this initiator to an AIX igroup.

Step

1. Enter the following command:

```
igroup add igroup_name [nodename|WWPN|WWPN alias]
```

Example

For Windows:

```
igroup add win-group2 iqn.1991-05.com.microsoft:eng2
```

For AIX:

```
igroup add aix-group2 10:00:00:00:c9:2b:02:1f
```

Removing initiators from an igroup

You can use the `igroup remove` command to remove an initiator from an igroup.

Step

1. Enter the following command:

```
igroup remove igroup_name [nodename|WWPN|WWPN alias]
```

Example

For Windows:

```
igroup remove win-group1 iqn.1991-05.com.microsoft:eng1
```

For AIX:

```
igroup remove aix-group1 10:00:00:00:c9:2b:7c:0f
```

Displaying initiators

You can use the `igroup show` command to display all initiators belonging to a particular `igroup`.

Step

1. Enter the following command:

```
igroup show igroup_name
```

Example

```
igroup show -v igroup1
```

```
system1> igroup show -v igroup1
igroup1:
OS Type: solaris
Member: 10:00:00:00:c9:2b:cc:39 (logged in on: vtic, 5a, 5b)
Member: 10:00:00:00:c9:2b:cb:7e
ALUA: Yes
Report SCSI Name in Inquiry Descriptor: No
```

Note: The output of `igroup show -v` displays the FCP initiator logged in on physical ports as well as a port called "vtic". VTIC is an abbreviation for "virtual target interconnect." VTIC provides a connection between the two nodes in an HA pair, enabling LUNs to be served through target ports on both nodes. It is normal to see VTIC as one of the ports in the output of `igroup show -v`.

Related information

What does the order of listed target ports mean in the output of `igroup show`? - <https://kb.netapp.com/support/index?page=content&id=3011742>

Renaming igroups

You can use the `igroup rename` command to rename an `igroup`.

Step

1. Enter the following command:

```
igroup rename current_igroup_name new_igroup_name
```

Example

```
igroup rename win-group3 win-group4
```

Setting the operating system type for an igroup

When creating an igroup, you must set the operating system type, or ostype, to one of the supported ostype values.

About this task

The supported ostyles are: solaris, Solaris_efi, windows, windows_gpt, windows_2008, hpux, aix, linux, netware, vmware, xen, and hyper_v.

Step

1. Enter the following command:

```
igroup set [-f]igroup ostyle value
```

-f overrides all warnings.

igroup is the name of the igroup.

value is the operating system type of the igroup.

Example

```
igroup set aix-group3 ostyle aix
```

The ostyle for igroup aix-group3 is set to aix.

Related concepts

[*ostype \(LUN multiprotocol type\) guidelines*](#) on page 49

SAN Protocol Management

SAN supports iSCSI networks, Fibre Channel fabrics and Fibre Channel over Ethernet. You have various options in the management of each protocol type.

iSCSI network management

You can understand how to manage the iSCSI service, as well as manage the storage system as a target in the iSCSI network.

Enabling multi-connection sessions

By default, Data ONTAP is now configured to use a single TCP/IP connection for each iSCSI session. If you are using an initiator that has been qualified for multi-connection sessions, you can specify the maximum number of connections allowed for each session on the storage system.

About this task

The `iscsi.max_connections_per_session` option specifies the number of connections per session allowed by the storage system. You can specify between 1 and 32 connections, or you can accept the default value.

Note that this option specifies the maximum number of connections per session supported by the storage system. The initiator and storage system negotiate the actual number allowed for a session when the session is created; this is the smaller of the initiator's maximum and the storage system's maximum. The number of connections actually used also depends on how many connections the initiator establishes.

Steps

1. Verify the current option setting by entering the following command on the system console:

```
options iscsi.max_connections_per_session
```

The current setting is displayed.

2. If needed, change the number of connections allowed by entering the following command:

```
options iscsi.max_connections_per_session [connections |  
use_system_default]
```

connections is the maximum number of connections allowed for each session, from 1 to 32.

`use_system_default` equals 1 for Data ONTAP 7.1, 16 for Data ONTAP 7.2 and subsequent maintenance releases, and 32 starting with Data ONTAP 7.3. The meaning of this default might change in later releases.

Enabling 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 `iscsi.max_error_recovery_level` option 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.

Note: This option 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: You can only change the session error recovery level for newly created sessions. This change does not affect the level for existing sessions.

Steps

1. Verify the current option setting by entering the following command on the system console:

```
options iscsi.max_error_recovery_level
```

The current setting is displayed.

2. If needed, change the error recovery levels allowed by entering the following command:

```
options iscsi.max_error_recovery_level [level | use_system_default]
```

level is the maximum error recovery level allowed, 0, 1, or 2.

`use_system_default` equals 0 for Data ONTAP 7.1 and 7.2. The value of this default might change in later releases.

iSCSI service management

You need to ensure the iSCSI service is licensed and running on your system, as well as properly manage the target node name and target alias.

Verifying that the iSCSI service is running

You can use the `iscsi status` command to verify that the iSCSI service is running.

Step

1. On the storage system console, enter the following command:

```
iscsi status
```

A message is displayed indicating whether iSCSI service is running.

Verifying that iSCSI is licensed

You can use the `license` command to verify that iSCSI is licensed on the storage system.

Step

1. On the storage system console, enter the following command:

license

Displays the list of all services that are licensed and the details about the license package in Type, Description, and Expiration columns. This command does not display the services that are not licensed.

Enabling the iSCSI license

Before you can use the iSCSI target service, you must enable the iSCSI license by entering the iSCSI license key and turning on the `iscsi` option.

Steps

1. Use the following command to enter your license key:

license add iscsi_license_code

Example

```
system1> license add XXXXXXXXXXXXXXXXXXXXXXXXXXXX
license add: successfully added license key "
XXXXXXXXXXXXXXXXXXXXXXXXXXXXX".
```

2. Enter the following command to enable the `iscsi` option:

options licensed_feature.iscsi.enable on

Example

```
system1> options licensed_feature.iscsi.enable on
Tue Sep 11 05:24:44 GMT [f3170-SAN-235-12:kern.cli.cmd:debug]:
Command line input: the command is 'options'.
The full command line is 'options licensed_feature.iscsi.enable on'.
Run 'iscsi start' to start the iSCSI service.
Also run 'lun setup' if necessary to configure LUNs.
```

Starting the iSCSI service

You can use the `iscsi start` command to start the iSCSI service on the storage system.

Step

1. On the storage system console, enter the following command:

```
iscsi start
```

Disabling the iSCSI license

If you do not want to use the iSCSI service, disable the iSCSI license by turning off the `iscsi` option and removing the `iscsi` license key.

About this task

Note: If you disable the iSCSI license, you cannot access the iSCSI service and iSCSI target connectivity is lost. Any LUNs being served to the initiators are terminated.

Steps

1. Enter the following command to remove your iSCSI license key:

```
license delete iscsi
```

Example

```
f3170-SAN-235-12> license delete iSCSI
license delete: successfully deleted "iSCSI"
```

2. Enter the following command to disable the `iscsi` option:

```
options licensed_feature.iscsi.enable off
```

Example

```
f3170-SAN-235-12> options licensed_feature.iscsi.enable off
Tue Sep 11 07:43:55 GMT [f3170-SAN-235-12:kern.cli.cmd:debug]:
Command line input: the command is 'options'.
The full command line is 'options licensed_feature.iscsi.enable off'.
Tue Sep 11 07:43:55 GMT [f3170-
SAN-235-12:iscsi.service.shutdown:info]: iSCSI service shutdown
```

Stopping the iSCSI service

You can use the `iscsi stop` command to stop the iSCSI service on the storage system.

Step

1. On the storage system console, enter the following command:

```
iscsi stop
```

Displaying the target node name

You can use the `iscsi nodename` command to display the storage system's target node name.

Step

1. On the storage system console, enter the following command:

```
iscsi nodename
```

Example

```
system1> iscsi nodename  
iSCSI target nodename: iqn.1992-08.com.netapp:sn.12345678
```

Changing the target node name

You might have to change the storage system's target node name.

About this task

Changing the storage system's node name while iSCSI sessions are in progress does not disrupt the existing sessions. However, when you change the storage system's node name, you must reconfigure the initiator so that it recognizes the new target node name. If you do not reconfigure the initiator, subsequent initiator attempts to log in to the target fail.

When you change the storage system's target node name, be sure the new name follows all of these rules:

- A node name can be up to 223 bytes.
- Uppercase characters are always mapped to lowercase characters.
- A node name can contain alphabetic characters (a to z), numbers (0 to 9) and three special characters:
 - Period (“.”)
 - Hyphen (“-”)
 - Colon (“:”)

- The underscore character (“_”) is *not* supported.

Step

1. On the storage system console, enter the following command:

```
iscsi nodename iqn.1992-08.com.netapp:unique_device_name
```

Example

```
iscsi nodename iqn.1992-08.com.netapp:filerhq
```

Displaying the iSCSI target alias

The target alias is an optional name for the iSCSI target consisting of a text string with a maximum of 128 characters. It is displayed by an initiator's user interface to make it easier for someone to identify the desired target in a list of targets.

About this task

Depending on your initiator, the user interface of the initiator might display the alias name.

Step

1. On the storage system console, enter the following command:

```
iscsi alias
```

Example

```
system1> iscsi alias  
iSCSI target alias: Filer_1
```

Adding or changing the iSCSI target alias

You can change the target alias or clear the alias at any time without disrupting existing sessions. The new alias is sent to the initiators the next time they log in to the target.

Step

1. On the storage system console, enter the following command:

```
iscsi alias [-c | string]
```

-c clears the existing alias value

string is the new alias value, maximum 128 characters

Examples

```
system1> iscsi alias Storage-System_2
New iSCSI target alias: Storage-System_2
```

```
system1> iscsi alias -c
Clearing iSCSI target alias
```

iSCSI service management on storage system interfaces

You can manage the iSCSI service on the storage system's Ethernet interfaces by using the `iscsi interface` command.

You can control which network interfaces are used for iSCSI communication. For example, you can enable iSCSI communication over specific gigabit Ethernet (GbE) interfaces.

By default, the iSCSI service is enabled on all Ethernet interfaces after you enable the license. The e0M management interface on storage systems is a 10/100 interface.

Note: iSCSI communication cannot be enabled in all the private ports and management ports. If you attempt to enable these ports, you will get an error message indicating the interface is not usable for iSCSI.

Displaying iSCSI interface status

You can use the `iscsi interface show` command to display the status of the iSCSI service on a storage system interface.

Step

1. On the storage system console, enter the following command:

```
iscsi interface show [-a | interface]
```

-a specifies all interfaces. This is the default.

interface is a list of specific Ethernet interfaces, separated by spaces.

Example

The following example shows the iSCSI service enabled on two storage system Ethernet interfaces:

```
system1> iscsi interface show
Interface e0a disabled
Interface e9a enabled
Interface e9b enabled
```

Enabling iSCSI on a storage system interface

You can use the `iscsi interface enable` command to enable the iSCSI service on an interface.

Step

1. On the storage system console, enter the following command:

```
iscsi interface enable [-a | interface ...]
```

`-a` specifies all interfaces.

interface is a list of specific Ethernet interfaces, separated by spaces.

Example

The following example enables the iSCSI service on interfaces e9a and e9b:

```
iscsi interface enable e9a e9b
```

Disabling iSCSI on a storage system interface

You can use the `iscsi interface disable` command to disable the iSCSI service on an interface.

Step

1. On the storage system console, enter the following command:

```
iscsi interface disable [-f] {-a | interface ...}
```

`-f` forces the termination of any outstanding iSCSI sessions without prompting you for confirmation. If you do not use this option, the command displays a message notifying you that active sessions are in progress on the interface and requests confirmation before terminating these sessions and disabling the interface.

`-a` specifies all interfaces.

interface is a list of specific Ethernet interfaces, separated by spaces.

Displaying the storage system's target IP addresses

You can use the `iscsi portal show` command to display the target IP addresses of the storage system. The storage system's target IP addresses are the addresses of the interfaces used for the iSCSI protocol.

Step

1. On the storage system console, enter the following command:

iscsi portal show

Result

The IP address, TCP port number, target portal group tag, and interface identifier are displayed for each interface.

Example

```
system1> iscsi portal show
Network portals:
IP address                TCP Port  TPGroup  Interface
10.60.155.105             3260     1000     e0b
fe80::2a0:98ff:fe00:fd81  3260     1000     e0b
10.1.1.10                 3260     1003     e10a
fe80::200:c9ff:fe44:212b  3260     1003     e10a
```

iSCSI interface access management

Although you can use the `iscsi interface enable` command to enable the iSCSI service on an iSCSI interface, this command enables access for all initiators. As of Data ONTAP 7.3, you can use access lists to control the interfaces over which an initiator can access the storage system.

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.
- **Controller failover:** Instead of contacting all interfaces advertised by the storage system during giveback, the host attempts to contact the interfaces to which it has access, thereby improving failover times.

By default, all initiators have access to all interfaces, so 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.

Creating iSCSI interface access lists

You can use iSCSI interface access lists to control which interfaces an initiator can access. An access list ensures that an initiator only logs in with IP addresses associated with the interfaces defined in the access list.

About this task

Access list policies are based on the interface name, and can include physical interfaces, interface groups, and VLANs.

Note: For vFiler contexts, all interfaces can be added to the vFiler unit's access list, but the initiator can only access the interfaces that are bound to the vFiler unit's IP addresses.

Step

1. On the storage system console, enter the following command:

```
iscsi interface accesslist add initiator name [-a | interface...]
```

-a specifies all interfaces. This is the default.

interface lists specific Ethernet interfaces, separated by spaces.

Example

```
iscsi interface accesslist add iqn.1991-05.com.microsoft:ms e0b
```

Related concepts

[*Guidelines for using iSCSI with HA pairs*](#) on page 122

Removing interfaces from iSCSI interface access lists

If you created an access list, you can remove one or more interfaces from the access list.

Step

1. On the storage system console, enter the following command:

```
iscsi interface accesslist remove initiator name [-a | interface...]
```

-a specifies all interfaces. This is the default.

interface lists specific Ethernet interfaces, separated by spaces.

Example

```
iscsi interface accesslist remove iqn.1991-05.com.microsoft:ms e0b
```

Displaying iSCSI interface access lists

If you created one or more access lists, you can display the initiators and the interfaces to which they have access.

Step

1. On the storage system console, enter the following command:

```
iscsi interface accesslist show
```

Example

```
system1> iscsi interface accesslist show
Initiator Nodename          Access List
iqn.1987-05.com.cisco:redhat    e0a, e0b
iqn.1991-05.com.microsoft:ms    e9b
```

Only initiators defined as part of an access list are displayed.

iSNS server registration

If you decide to use an iSNS service, you must ensure that your storage systems are properly registered with an Internet Storage Name Service server.

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 the storage system interacts with an iSNS server

The storage system automatically registers its IP address, node name, and portal groups with the iSNS server when the iSCSI service is started and iSNS is enabled. After iSNS is initially configured, Data ONTAP automatically updates the iSNS server any time the storage system's configuration settings change.

There can be a delay of a few minutes between the time of the configuration change and the update being sent; you can use the `iscsi isns update` command to send an update immediately.

About iSNS service version incompatibility

The specification for the iSNS service is still in draft form. Some draft versions are different enough to prevent the storage system from registering with the iSNS server. Because the protocol does not provide version information to the draft level, iSNS servers and storage systems cannot negotiate the draft level being used.

In Data ONTAP 7.1 and after, the default iSNS version is draft 22. This draft is also used by Microsoft iSNS server 3.0.

Note: When you upgrade to a new version of Data ONTAP, the existing value for the `iscsi.isns.rev` option is maintained. This reduces the risk of a draft version problem when upgrading.

Setting the iSNS service revision

You can configure Data ONTAP to use a different iSNS draft version by changing the `iscsi.isns.rev` option on the storage system.

Steps

1. Verify the current iSNS revision value by entering the following command on the system console:

options iscsi.isns.rev

The current draft revision used by the storage system is displayed.

2. If needed, change the iSNS revision value by entering the following command:

options iscsi.isns.rev *draft*

draft is the iSNS standard draft revision, either 18 or 22.

Registering the storage system with an iSNS server

You can use the `iscsi isns` command to configure the storage system to register with an iSNS server. This command specifies the information the storage system sends to the iSNS server.

About this task

The `iscsi isns` command only configures the storage system to register with the iSNS server. The storage system 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 of the iSNS server.

Steps

1. Ensure that the iSCSI service is running by entering the following command on the storage system console:

iscsi status

2. If the iSCSI service is not running, enter the following command:

iscsi start

3. On the storage system console, enter the following command to identify the iSNS server that the storage system registers with:

iscsi isns config [*ip_addr*|*hostname*]

ip_addr is the IP address of the iSNS server.

hostname is the hostname associated with the iSNS server.

Note: As of Data ONTAP 7.3.1, you can configure iSNS with an IPv6 address.

4. Enter the following command:

```
iscsi isns start
```

The iSNS service is started and the storage system registers with the iSNS server.

Note: iSNS registration is persistent across reboots if the iSCSI service is running and iSNS is started.

Updating the iSNS server immediately

Data ONTAP checks for iSCSI configuration changes on the storage system every few minutes and automatically sends any changes to the iSNS server. If you do not want to wait for an automatic update, you can immediately update the iSNS server.

Step

1. On the storage system console, enter the following command:

```
iscsi isns update
```

Disabling iSNS

When you stop the iSNS service, the storage system stops registering its iSCSI information with the iSNS server.

Step

1. On the storage system console, enter the following command:

```
iscsi isns stop
```

Setting up vFiler units with the iSNS service

You can use the `iscsi isns` command on each vFiler unit to configure which iSNS server to use and to turn iSNS registration on or off.

About this task

For information about managing vFiler units, see the sections on iSCSI service on vFiler units in the *Data ONTAP MultiStore Management Guide for 7-Mode*.

Steps

1. Register the vFiler unit with the iSNS service by entering the following command:

```
iscsi isns config ip_addr
```

ip_addr is the IP address of the iSNS server.

2. Enable the iSNS service by entering the following command:

iscsi isns start**Examples for vFiler units**

The following example defines the iSNS server for the default vFiler unit (vfiler0) on the hosting storage system:

```
iscsi isns config 10.10.122.101
```

The following example defines the iSNS server for a specific vFiler unit (vf1). The **vfiler** context command switches to the command line for a specific vFiler unit.

```
vfiler context vf1
vf1> iscsi isns config 10.10.122.101
```

Related information

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

Displaying initiators connected to the storage system

You can display a list of initiators currently connected to the storage system. The information displayed for each initiator includes the target session identifier handle (TSIH) assigned to the session, the target portal group tag of the group to which the initiator is connected, the iSCSI initiator alias (if provided by the initiator), the initiator's iSCSI node name and initiator session identifier (ISID), and the igroup.

Step

1. On the storage system console, enter the following command:

```
iscsi initiator show
```

The initiators currently connected to the storage system are displayed.

Example

```
system1> iscsi initiator show
Initiators connected:
  TSIH  TPGGroup  Initiator/ISID/IGroup
    1    1000    iqn.1991-05.com.microsoft:hual-lxp.hq.netapp.com /
40:00:01:37:00:00 / windows_ig2; windows_ig
    2    1000    vanclibern (iqn.1987-05.com.cisco:vanclibern /
00:02:3d:00:00:01 / linux_ig)
    4    1000    iqn.1991-05.com.microsoft:cox / 40:00:01:37:00:00 /
```

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.

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 random 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.

If you use iSCSI with vFiler units, the CHAP authentication settings are configured separately for each vFiler unit. Each vFiler unit has its own default authentication mode and list of initiators and passwords.

To configure CHAP settings for vFiler units, you must use the command line.

For information about managing vFiler units, see the sections on iSCSI service on vFiler units in the *Data ONTAP MultiStore Management Guide for 7-Mode*.

Related information

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

Guidelines for using CHAP authentication

You should follow certain guidelines when using CHAP authentication.

- If you are not using RADIUS and 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.

Defining an authentication method for an 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.

About this task

You can generate a random password or you can specify the password that you want to use.

Steps

1. Generate a random password by entering the following command:

```
iscsi security generate
```

The storage system generates a 128-bit random password.

2. For each initiator, enter the following command:

```
iscsi security add -i initiator -s [chap | deny | none] [-f radius | -p  
inpassword -n inname] [-o outpassword -m outname]
```

initiator is the initiator name in the iSCSI nodename format.

The -s option takes one of several values:

chap—Authenticate using a CHAP user name and password.

none—The initiator can access the storage system without authentication.

deny—The initiator cannot access the storage system.

radius indicates that RADIUS is used for authentication. You can use the *-f* option to ensure that initiator only uses RADIUS as the authentication method. If you do not use the *-f* option, the initiator only attempts to authenticate via RADIUS if the local CHAP authentication fails.

inpassword is the inbound password for CHAP authentication. The storage system uses the inbound password to authenticate the initiator. An inbound password is required if you are using CHAP authentication and you are not using RADIUS.

iname is a user name for inbound CHAP authentication. The storage system uses the inbound user name to authenticate the initiator.

outpassword is a password for outbound CHAP authentication. It is stored locally on the storage system, which uses this password for authentication by the initiator.

outname is a user name for outbound CHAP authentication. The storage system uses this user name for authentication by the initiator.

Note: If you generated a random password, you can use this string for either *inpassword* or *outpassword*. If you enter a string, the storage system interprets an ASCII string as an ASCII value and a hexadecimal string, such as 0x1345, as a binary value.

Defining a default authentication method for initiators

You can use the `iscsi security default` command to define a default authentication method for all initiators not specified with the `iscsi security add` command.

Step

1. On the storage system console, enter the following command:

```
iscsi security default -s [chap | none | deny] [-f radius | -p  
inpassword -n inname] [-o outpassword -m outname]
```

The *-s* option takes one of three values:

chap—Authenticate using a CHAP user name and password.

none—The initiator can access the storage system without authentication.

deny—The initiator cannot access the storage system.

radius indicates that RADIUS authentication is used. You can use the *-f* option to ensure that initiator only uses RADIUS as the authentication method. If you do not use the *-f* option, the initiator only attempts to authenticate via RADIUS if the local CHAP authentication fails.

inpassword is the inbound password for CHAP authentication. The storage system uses the inbound password to authenticate the initiator.

iname is a user name for inbound CHAP authentication. The storage system uses the inbound user name to authenticate the initiator.

outpassword is a password for outbound CHAP authentication. The storage system uses this password for authentication by the initiator.

outname is a user name for outbound CHAP authentication. The storage system uses this user name for authentication by the initiator.

Displaying initiator authentication methods

You can use the `iscsi security show` command to view a list of initiators and their authentication methods.

Step

1. On the storage system console, enter the following command:

```
iscsi security show
```

Removing authentication settings for an initiator

You can use the `iscsi security delete` command to remove the authentication settings for an initiator and use the default authentication method.

Step

1. On the storage system console, enter the following command:

```
iscsi security delete -i initiator
```

-i initiator is the initiator name in the iSCSI node name format.

The initiator is removed from the authentication list and logs in to the storage system using the default authentication method.

iSCSI RADIUS configuration

You can configure your storage systems to use RADIUS for centrally managing iSCSI initiator authentication.

RADIUS uses CHAP to authenticate iSCSI initiators, but it enables you to manage the authentication process from a central RADIUS server, rather than manage it manually on each storage system. In larger SAN environments, this can greatly simplify iSCSI initiator management, CHAP password management, and provide added security.

RADIUS also reduces the load on your storage system because most of the authentication processing is handled by the RADIUS server.

Defining RADIUS as the authentication method for initiators

You can define RADIUS as the authentication method for one or more initiators, as well as make it the default authentication method that applies to initiators that are not on this list.

About this task

You can generate a random password, or you can specify the password you want to use. Inbound passwords are saved on the RADIUS server and outbound passwords are saved on the storage system.

Steps

1. To generate a random password, enter the following command:

```
iscsi security generate
```

The storage system generates a 128-bit random password.

Note: If you generate a random inbound password, you must add this password to the RADIUS server.

2. For each initiator, enter the following command:

```
iscsi security add -i initiator -s chap -f radius [-o outpassword -m outname]
```

initiator is the initiator name in the iSCSI nodename format.

Use the **-f** option to ensure that initiator only uses RADIUS as the authentication method. If you do not use the **-f** option, the initiator only attempts to authenticate via RADIUS if the local CHAP authentication fails.

outpassword is a password for outbound CHAP authentication. It is stored locally on the storage system, which uses this password for authentication by the initiator.

outname is a user name for outbound CHAP authentication. The storage system uses this user name for authentication by the initiator.

Note: If you generated a random password, you can use this string for *outpassword*. If you enter a string, the storage system interprets an ASCII string as an ASCII value and a hexadecimal string, such as 0x1345, as a binary value.

3. To define RADIUS as the default authentication method for all initiators not previously specified, enter the following command:

```
iscsi security default -s chap -f radius [-o outpassword -m outname]
```

Examples

```
system1> iscsi security add -i iqn.1992-08.com.microsoft:system1
-s chap -f radius
```

```
system1> iscsi security show
Default sec is CHAP RADIUS Outbound password: **** Outbound
username:
init: iqn.1994-05.com.redhat:10ca21e21b75 auth: CHAP RADIUS
Outbound password: **** Outbound username: icroto
```

```
system1> iscsi security default -s chap -f radius
```

After you finish

After enabling RADIUS authentication for the initiators, start the RADIUS client service on the storage system.

Starting the RADIUS client service

After you enable RADIUS authentication for the appropriate initiators, you must start the RADIUS client.

Step

1. Enter the following command:

```
radius start
```

Example

```
system1> radius start
RADIUS client service started
```

After you finish

After the RADIUS service is started, ensure that you add one or more RADIUS servers with which the storage system can communicate.

Adding a RADIUS server

After you start the RADIUS client service, add a RADIUS server with which the storage system can communicate. You can add up to three RADIUS servers.

Step

1. Enter the following command:

```
radius add [-d] RADIUS_server_hostname or ip_address [-p port_number]
```

You can use the `-d` option to make the RADIUS server you are adding the default server. If there is no default server, the one you add becomes the default.

You can use the `-p` option to specify a port number on the RADIUS server. The default port number is 1812.

Example

```
system1> radius add 10.60.155.58 -p 1812
system1> radius show
RADIUS client service is running

Default RADIUS server : IP_Addr=10.60.155.58  UDPPort=1812
```

After you finish

After adding the necessary servers, you must enable the storage system to use the RADIUS server for CHAP authentication.

Enabling the storage system to use RADIUS for CHAP authentication

After RADIUS authentication is enabled for the initiators and the RADIUS client service is started, you must set the `iscsi.auth.radius.enable` option to on. This ensures that the storage system uses RADIUS for CHAP authentication.

About this task

This option is set to off by default, and you must set it to on, regardless of whether you used the `-f` option when enabling RADIUS for the initiators.

Step

1. Enter the following command:

`options iscsi.auth.radius.enable on`

Your system is now enabled for RADIUS authentication.

```
system1> options iscsi.auth.radius.enable on
system1> options iscsi
iscsi.auth.radius.enable      on
iscsi.enable                  on
iscsi.isns.rev                22
iscsi.max_connections_per_session use_system_default
iscsi.max_error_recovery_level use_system_default
iscsi.max_ios_per_session     128
iscsi.tcp_window_size         131400
```

Displaying the RADIUS service status

You can use the `radius show` command to display important RADIUS information, including whether the service is running and the default RADIUS server.

Step

1. Enter the following command:

```
radius show
```

Example

```
system1> radius show
RADIUS client service is running

Default RADIUS server : IP_Addr=10.60.155.58  UDPPort=1812
```

You can also run the `radius status` command to see if the client service is running.

Example

```
system1> radius status
RADIUS client service is running
```

Stopping the RADIUS client service

You can use the `radius stop` command to stop the RADIUS client service.

Step

1. Enter the following command:

```
radius stop
```

```
system1> radius stop
RADIUS client service stopped
```

Removing a RADIUS server

You can use the `radius remove` command to ensure that a RADIUS server is no longer used for RADIUS authentication.

Step

1. Enter the following command:

```
radius remove RADIUS_server_hostname or ip_address [-p port_number]
```

If the server is using a port other than 1812, use the `-p` option to specify the port number.

```
system1> radius show
RADIUS client service is running

Default RADIUS server : IP_Addr=10.60.155.58 UDPPort=1812

system1> radius remove 10.60.155.58
system1> radius show
RADIUS client service is running
```

Displaying and clearing RADIUS statistics

You can use the `radius stats` command to display important details about the RADIUS service, including packets accepted, packets rejected, and the number of authentication requests. You can also clear the existing statistics.

Step

1. Enter the following command:

```
radius stats [-z]
```

You can use the `-z` option to clear the statistics.

```
system1> radius stats
RADIUS client statistics
  RADIUS access-accepted-packets: 121
  RADIUS access-challenged-packets: 3
  RADIUS access-rejected-packets: 0
  RADIUS authentication-requests: 124
  RADIUS denied-packets: 0
  RADIUS late-packets: 0
  RADIUS retransmitted-packets: 14
  RADIUS short-packets: 0
  RADIUS timed-out-packets: 0
```

```
RADIUS unknown-packets:      0
RADIUS unknown-server-packets: 0
```

```
system1> radius stats -z
```

```
system1> radius stats
RADIUS client statistics
  RADIUS access-accepted-packets:      0
  RADIUS access-challenged-packets:    0
  RADIUS access-rejected-packets:      0
  RADIUS authentication-requests:      0
  RADIUS denied-packets:               0
  RADIUS late-packets:                 0
  RADIUS retransmitted-packets:        0
  RADIUS short-packets:                0
  RADIUS timed-out-packets:            0
  RADIUS unknown-packets:              0
  RADIUS unknown-server-packets:       0
```

Target portal group management

A target portal group is a set of one or more storage system network interfaces that can be used for an iSCSI session between an initiator and a target. A target portal group is identified by a name and a numeric tag. If you want to have multiple connections per session across more than one interface for performance and reliability reasons, then you must use target portal groups.

Note: If you are using MultiStore, you can also configure non-default vFiler units for target portal group management based on IP address.

For iSCSI sessions that use multiple connections, all of the connections must use interfaces in the same target portal group. Each interface belongs to one and only one target portal group. Interfaces can be physical interfaces or logical interfaces (VLANs and interface groups).

Prior to Data ONTAP 7.1, each interface was automatically assigned to its own target portal group when the interface was added. The target portal group tag was assigned based on the interface location and could not be modified. This works fine for single-connection sessions.

You can explicitly create target portal groups and assign tag values. If you want to increase performance and reliability by using multi-connections per session across more than one interface, you must create one or more target portal groups.

Because a session can use interfaces in only one target portal group, you might want to put all of your interfaces in one large group. However, some initiators are also limited to one session with a given target portal group. To support multipath I/O (MPIO), you need to have one session per path, and therefore more than one target portal group.

When a new network interface is added to the storage system, that interface is automatically assigned to its own target portal group.

Range of values for target portal group tags

If you create target portal groups, the valid values you can assign to target portal group tags vary depending on the type of interface.

The following table shows the range of values for target portal group tags:

Target portal group type	Allowed values
User defined groups	1 - 256
Default groups for physical devices	1,000 - 1,511
Default groups for VLANs and interface groups	2,000 - 2,511
Default groups for IP-based vFiler units	4,000 - 65,535

Important cautions for using target portal groups

You must be aware of some important cautions when using target portal groups.

- Some initiators, including those used with Windows, HP-UX, Solaris, and Linux, create a persistent association between the target portal group tag value and the target. If the target portal group tag changes, the LUNs from that target are unavailable.
- Adding or removing a NIC might change the target portal group assignments. You should ensure that you verify the target portal group settings are correct after adding or removing hardware, especially in HA pairs.
- When used with multi-connection sessions, the Windows iSCSI software initiator creates a persistent association between the target portal group tag value and the target interfaces. If the tag value changes while an iSCSI session is active, the initiator recovers only one connection for a session. To recover the remaining connections, you must refresh the initiator's target information.

Displaying target portal groups

You can use the `iscsi tpgroup show` command to display a list of existing target portal groups.

Step

1. On the storage system console, enter the following command:

`iscsi tpgroup show`

Example

```
system1> iscsi tpgroup show
TPGTag  Name           Member Interfaces
1000    e0_default      e0
```

1001	e5a_default	e5a
1002	e5b_default	e5b
1003	e9a_default	e9a
1004	e9b_default	e9b

Creating target portal groups

If you want to employ multiconnection sessions to improve performance and reliability, you must use target portal groups to define the interfaces available for each iSCSI session.

About this task

You must create a target portal group that contains all of the interfaces you want to use for one iSCSI session. However, note that you cannot combine iSCSI hardware-accelerated interfaces with standard iSCSI storage system interfaces in the same target portal group.

When you create a target portal group, the specified interfaces are removed from their current groups and added to the new group. Any iSCSI sessions using the specified interfaces are terminated, but the initiator should automatically reconnect. However, initiators that create a persistent association between the IP address and the target portal group cannot reconnect.

Step

1. On the storage system console, enter the following command:

```
iscsi tpgroup create [-f] tpgroup_name [-t tag] [interface ...]
```

-f forces the new group to be created, even if that terminates an existing session using one of the interfaces being added to the group.

tpgroup_name is the name of the group being created (1 to 60 characters, no spaces or non-printing characters).

-t *tag* sets the target portal group tag to the specified value. In general you should accept the default tag value. User-specified tags must be in the range 1 to 256.

interface ... is the list of interfaces to include in the group, separated by spaces.

Example

The following command creates a target portal group named `server_group` that includes interfaces `e8a` and `e9a`:

```
iscsi tpgroup create server_group e8a e9a
```

Destroying target portal groups

Destroying a target portal group removes the group from the storage system. Any interfaces that belonged to the group are returned to their individual default target portal groups. Any iSCSI sessions with the interfaces in the group being destroyed are terminated.

Step

1. On the storage system console, enter the following command:

```
iscsi tpgroup destroy [-f] tpgroup_name
```

-f forces the group to be destroyed, even if that terminates an existing session using one of the interfaces in the group.

tpgroup_name is the name of the group being destroyed.

Adding interfaces to target portal groups

You can add interfaces to an existing target portal group. The specified interfaces are removed from their current groups and added to the new group.

About this task

Any iSCSI sessions using the specified interfaces are terminated, but the initiator should reconnect automatically. However, initiators that create a persistent association between the IP address and the target portal group are not able to reconnect.

Step

1. On the storage system console, enter the following command:

```
iscsi tpgroup add [-f] tpgroup_name [interface ...]
```

-f forces the interfaces to be added, even if that terminates an existing session using one of the interfaces being added to the group.

tpgroup_name is the name of the group.

interface ... is the list of interfaces to add to the group, separated by spaces.

Example

The following command adds interfaces e8a and e9a to the portal group named server_group:

```
iscsi tpgroup add server_group e8a e9a
```

Removing interfaces from target portal groups

You can remove interfaces from an existing target portal group. The specified interfaces are removed from the group and returned to their individual default target portal groups.

About this task

Any iSCSI sessions with the interfaces being removed are terminated, but the initiator should reconnect automatically. However, initiators that create a persistent association between the IP address and the target portal group are not able to reconnect.

Step

1. On the storage system console, enter the following command:

```
iscsi tpgroup remove [-f] tpgroup_name [interface ...]
```

-f forces the interfaces to be removed, even if that terminates an existing session using one of the interfaces being removed from the group.

tpgroup_name is the name of the group.

interface ... is the list of interfaces to remove from the group, separated by spaces.

Example

The following command removes interfaces e8a and e9a from the portal group named `server_group`, even though there is an iSCSI session currently using e8a:

```
iscsi tpgroup remove -f server_group e8a e9a
```

Target portal group management for online migration of vFiler units

Target portal groups enable you to efficiently manage iSCSI sessions between initiators and targets. Although Data ONTAP manages target portal groups using network interfaces by default, you can also manage these groups using IP addresses, starting with Data ONTAP 7.3.3. This is required if you want to perform an online migration of vFiler units, which allows you to nondisruptively migrate data from one storage system to another.

Note: The NetApp Management Console provisioning capability is required for performing online migrations of vFiler units.

When you migrate data, the target portal group tag on the destination network interface must be identical to the target portal group tag on the source network interface. This is problematic in a MultiStore environment because the source and destination storage systems might be of different hardware platforms. Changing the target portal group tags after migration is not sufficient because some hosts, such as HP-UX and Solaris, do not support dynamic iSCSI target discovery, which results in a disruption of service to those hosts in the change process.

If offline (disruptive) migrations are not problematic in your environment, or if all of your hosts support dynamic iSCSI target discovery, then IP-based target portal group management is unnecessary.

If you choose to implement IP-based target portal groups by enabling the `iscsi.ip_based_tpgroup` option, interface-based target portal groups are automatically converted to IP-based target portal groups, and any future target portal group assignments are IP-based as well. However, note that if you are migrating between a system with IP-based target portal groups and a system with interface-based target portal groups, the target portal group information is lost and the iSCSI service might be disrupted.

Note: ALUA is not supported with IP-based target portal groups.

For more information about the NetApp Management Console provisioning capability, see the *Provisioning Manager and Protection Manager Guide to Common Workflows for Administrators*.

Related information

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

Upgrade and revert implications for IP-based target portal group management

Before implementing IP-based target portal groups for online migrations, it is important to understand the limitations under various upgrade and revert scenarios.

The following table describes the impact to your target portal group assignments when upgrading to or reverting from Data ONTAP 7.3.3 or later.

Scenario	Impact to target portal groups
Upgrade to Data ONTAP 7.3.3	No change—existing interface-based target portal groups are not converted to IP-based target portal groups.

Scenario	Impact to target portal groups
Revert from Data ONTAP 7.3.3 or later	- For the default vFiler unit (vfiler0), there is no impact. vfiler0 must always use interface-based target portal groups. - For non-default vFiler units: - If you implement interface-based target portal groups, then there is no impact; the existing assignments remain intact. - If you implement IP-based target portal groups, those assignments are lost, potentially disrupting the iSCSI service. Attention: Before reverting, make sure you turn off IP-based target portal groups by entering the following command: <pre>options iscsi.ip_based_tpgroup off</pre> Failure to do so might disrupt subsequent upgrades.

Enabling IP-based target portal group management

If you want to perform online migrations of vFiler unit, you must enable IP-based target portal groups on your vFiler units.

About this task

When you enable IP-based target portal groups, the existing interface-based target portal groups are automatically converted to IP-based target portal groups. However, note that the interface-based target portal groups remain intact for the default vFiler unit.

Step

1. Enter the following command:

```
vfiler run vFiler_unit options iscsi.ip_based_tpgroup on
```

The existing interface-based target portal groups are converted to IP-based target portal groups with no disruption in service to the host.

Example

Before enabling IP-based target portal groups, the target port group information for vFiler unit 2 (vf2) looks like this:

```
system1>vfiler run vf2 iscsi tpgroup show
TPGTag  Name                      Member Interfaces
```

```

32      user_defined32      (none)
1000    e0_default          e0
1002    e11b_default        e11b
1003    e11c_default        e11c
1004    e11d_default        e11d
1005    e9a_default         e9a
1006    e9b_default         e9b
1007    e10a_default        e10a
1008    e10b_default        e10b
2000    vif_e0-1_default    vif_e0-1
2001    vif_e0-2_default    vif_e0-2
2002    vif_e0-3_default    vif_e0-3
2003    vif_e11a-1_default  vif_e11a-1
2004    vif_e11a-2_default  vif_e11a-2
2005    vif_e11a-3_default  vif_e11a-3

```

Each interface is associated with various IP addresses, and some of those are assigned to vFiler unit vf2. For example:

```

system1> vfiler run vf2 iscsi portal show
Network portals:
IP address      TCP Port  TPGroup  Interface
10.60.155.104   3260     1000     e0
192.168.11.100  3260     2003     vif_e11a-1
192.168.11.101  3260     2003     vif_e11a-1
192.168.13.100  3260     2005     vif_e11a-3
192.168.13.101  3260     2005     vif_e11a-3

```

After enabling IP-based target portal groups for vf2, the relevant interface-based target portal groups for vf2 are nondisruptively converted to IP-based target portal groups.

```

system1> vfiler run vf2 options iscsi.ip_based_tpgroup on

system1> vfiler run -q vf2 iscsi ip_tpgroup show
TPGTag  Name                               Member IP Addresses
1000    e0_default                        10.60.155.104
2003    vif_e11a-1_default                192.168.11.100, 192.168.11.101
2005    vif_e11a-3_default                192.168.13.100, 192.168.13.101

system1> vfiler run -q vf2 iscsi portal show
Network portals:
IP address      TCP Port  TPGroup  Interface
10.60.155.104   3260     1000     e0
192.168.11.100  3260     2003     vif_e11a-1
192.168.11.101  3260     2003     vif_e11a-1
192.168.13.100  3260     2005     vif_e11a-3
192.168.13.101  3260     2005     vif_e11a-3

```

If you configure another IP address for vf2, then a new default IP-based target portal group (4000) is automatically created. For example:

```

system1> vfiler add vf2 -i 192.168.13.102

system1> ifconfig vif_e11a-3 alias 192.168.13.102

system1> vfiler run vf2 iscsi ip_tpgroup show
TPGTag  Name                               Member IP Addresses
1000    e0_default                           10.60.155.104
2003    vif_e11a-1_default                  192.168.11.100, 192.168.11.101
2005    vif_e11a-3_default                  192.168.13.100, 192.168.13.101
4000    192.168.13.102_default              192.168.13.102

system1> vfiler run vf2 iscsi portal show
Network portals:
IP address      TCP Port  TPGroup  Interface
10.60.155.104   3260     1000     e0
192.168.11.100  3260     2003     vif_e11a-1
192.168.11.101  3260     2003     vif_e11a-1
192.168.13.100  3260     2005     vif_e11a-3
192.168.13.101  3260     2005     vif_e11a-3
192.168.13.102  3260     4000     vif_e11a-3

```

After you finish

After you enable IP-based target portal group management, it is recommended to leave it enabled. However, if you must disable IP-based target portal groups for some reason, enter the following command:

options iscsi.ip_based_tpgroup off

As a result, any IP-based target portal group information is discarded, and the interface-based target portal group information is reenabled. Note that this process might disrupt the iSCSI service to the hosts.

Also note that if an IP address is unassigned from a vFiler unit or unconfigured from the network interface, that IP address is no longer a valid iSCSI portal. However, the IP-based target portal group to which that IP address belonged remains intact so that if you add the IP address back later, it is automatically assigned back to the original target portal group.

Displaying IP-based target portal group information

You can use the `iscsi ip_tpgroup show` command to display important information about your IP-based target portal groups, including target portal group tags, target portal group names, and the IP addresses that belong to each group.

Step

1. Enter the following command:

vfiler run vFiler_unit iscsi ip_tpgroup show

Example

```
system1> vfiler run vfiler2 iscsi ip_tpgroup show
TPGTag  Name                               Member IP Addresses
  1     vfiler2_migrate_test0         (none)
  2     vfiler2_migrate_test1         (none)
  3     vfiler2_migrate_test3         (none)
100     user_defined_tp1              (none)
128     vfiler2_ui_review             1.1.1.1
1007    e10a_default                  10.1.1.8
1008    e10b_default                  1.1.1.2
4000    10.1.1.5_default              10.1.1.5
4001    10.60.155.104_default          10.60.155.104
4002    192.168.1.1_default           192.168.1.1
```

Creating IP-based target portal groups

You can create new IP-based target portal groups in which to add and remove existing IP addresses.

Before you begin

IP-based target portal group management must be enabled by entering the following command:

options iscsi.ip_based_tpgroup on

Step

1. Enter the following command:

```
vfiler run vFiler_unit ip_tpgroup create [-f] [-t | tag] tpgroup_name
IP_address...
```

-f forces the new group to be created, even if that terminates an existing session using one of the IP addresses being added to the group.

-t sets the target portal group tag to the specified value. In general, you should accept the default tag value.

tpgroup_name is the target portal group name.

IP_address is the list of IP addresses to include in the group, separated by spaces.

Example

```
vfiler run vfiler2 iscsi ip_tpgroup create -t 233 vfiler2_tpg1
10.1.3.5
```

After you finish

You can add and remove IP addresses from the new group.

Related tasks

[Enabling IP-based target portal group management](#) on page 111

Destroying IP-based target portal groups

If necessary, you can destroy IP-based target portal groups.

Before you begin

No active sessions must be in progress.

Step

1. Enter the following command:

```
vfiler run vFiler unit iscsi ip_tpgroup destroy [-f] tpgroup_name
```

-f forces the group to be destroyed, even if that terminates an existing session using one of the IP addresses in the group.

tpgroup_name is the target portal group name.

The target portal group is destroyed, and if there are active iSCSI sessions, a warning message indicates that those connections are lost.

Example

```
vfiler run vfiler2 iscsi ip_tpgroup destroy vfiler2_tpg1
```

Adding IP addresses to IP-based target portal groups

You can use the `iscsi ip_tpgroup add` command to add an IP address to an existing IP-based target portal group.

Before you begin

- IP-based target portal group management must be enabled.
- There must be at least one existing IP-based target portal group.

Step

1. Enter the following command:

```
vfiler run vFiler_unit iscsi ip_tpgroup add [-f] tpgroup_name  
IP_address ...
```

-f forces the new group to be created, even if that terminates an existing session using one of the IP addresses being added to the group.

tpgroup_name is the target portal group name.

IP_address is the list of IP addresses to include in the group, separated by spaces.

Example

```
vfiler run vfiler2 iscsi ip_tpgroup add vfiler2_tpg1 192.168.2.1
192.112.2.1
```

Removing IP addresses from IP-based target portal groups

In the course of reconfiguring your network, you might need to remove one or more IP addresses from an IP-based target portal group.

Step

1. Enter the following command:

```
vfiler run vFiler_unit iscsi ip_tpgroup remove [-f] tpgroup_name
IP_address ...
```

-f forces the new group to be created, even if that terminates an existing session using one of the IP addresses being added to the group.

tpgroup_name is the target portal group name.

IP_address is the list of IP addresses to remove from the group, separated by spaces.

Example

```
vfiler run vfiler2 iscsi ip_tpgroup remove vfiler2_tpg1 192.112.2.1
```

Displaying iSCSI statistics

You can use the `iscsi stats` command to display important iSCSI statistics.

Step

1. On the storage system console, enter the following command:

```
iscsi stats [-a | -z | ipv4 | ipv6]
```

-a displays the combined IPv4 and IPv6 statistics followed by the individual statistics for IPv4 and IPv6.

-z resets the iSCSI statistics.

ipv4 displays only the IPv4 statistics.

ipv6 displays only the IPv6 statistics.

Entering the `iscsi stats` command without any options displays only the combined IPv4 and IPv6 statistics.

```

system1> iscsi stats -a

iSCSI stats(total)
iSCSI PDUs Received
  SCSI-Cmd:      1465619 | Nop-Out:      4 | SCSI
  TaskMgtCmd:    0      | LogoutReq:    1 | Text
  LoginReq:      6      | SNACK:        0 |
  Req:           1      |
  DataOut:       0      |
  Unknown:       0
  Total: 1465631
iSCSI PDUs Transmitted
  SCSI-Rsp:      733684 | Nop-In:      4 | SCSI
  TaskMgtRsp:    0      | LogoutRsp:    1 |
  LoginRsp:      6      |
  TextRsp:       1      |
  Data_In:       790518 | R2T:         0 |
  Asyncmsg:      0
  Reject:        0
  Total: 1524214
iSCSI CDBs
  DataIn Blocks:  5855367 | DataOut Blocks: 0
  Error Status:   1      | Success Status: 1465618
  Total CDBs: 1465619
iSCSI ERRORS
  Failed Logins:  0      | Failed TaskMgt: 0
  Failed Logouts: 0      | Failed TextCmd: 0
  Protocol:       0
  Digest:         0
  PDU discards (outside CmdSN window): 0
  PDU discards (invalid header): 0
  Total: 0

iSCSI Stats(ipv4)
iSCSI PDUs Received
  SCSI-Cmd:      732789 | Nop-Out:      1 | SCSI
  TaskMgtCmd:    0      | LogoutReq:    0 | Text
  LoginReq:      2      | SNACK:        0 |
  Req:           0      |
  DataOut:       0      |
  Unknown:       0
  Total: 732792
iSCSI PDUs Transmitted
  SCSI-Rsp:      366488 | Nop-In:      1 | SCSI
  TaskMgtRsp:    0      | LogoutRsp:    0 |
  LoginRsp:      2      |
  TextRsp:       0      |
  Data_In:       395558 | R2T:         0 |
  Asyncmsg:      0
  Reject:        0
  Total: 762049
iSCSI CDBs
  DataIn Blocks:  2930408 | DataOut Blocks: 0
  Error Status:   0      | Success Status: 732789
  Total CDBs: 732789

```

```

iSCSI ERRORS
  Failed Logins:          0 | Failed TaskMgt:          0
  Failed Logouts:         0 | Failed TextCmd:          0
  Protocol:               0
  Digest:                 0
  PDU discards (outside CmdSN window): 0
  PDU discards (invalid header):    0
  Total: 0

iSCSI Stats(ipv6)
iSCSI PDUs Received
  SCSI-Cmd:      732830 | Nop-Out:          3 | SCSI
  TaskMgtCmd:      0
  LoginReq:       4 | LogoutReq:        1 | Text
  Req:            1
  DataOut:        0 | SNACK:            0 |
  Unknown:        0
  Total: 732839
iSCSI PDUs Transmitted
  SCSI-Rsp:      367196 | Nop-In:           3 | SCSI
  TaskMgtRsp:      0
  LoginRsp:       4 | LogoutRsp:        1 |
  TextRsp:         1
  Data_In:       394960 | R2T:              0 |
  Asyncmsg:       0
  Reject:         0
  Total: 762165
iSCSI CDBs
  DataIn Blocks:    2924959 | DataOut Blocks:      0
  Error Status:      1 | Success Status:    732829
  Total CDBs: 732830
iSCSI ERRORS
  Failed Logins:          0 | Failed TaskMgt:          0
  Failed Logouts:         0 | Failed TextCmd:          0
  Protocol:               0
  Digest:                 0
  PDU discards (outside CmdSN window): 0
  PDU discards (invalid header):    0
  Total: 0

```

Definitions for iSCSI statistics

You can obtain the iSCSI statistics that are displayed when you run the `iscsi stats` command. For vFiler contexts, the statistics displayed refer to the entire storage system, not the individual vFiler units.

iSCSI PDUs received

The iSCSI Protocol Data Units (PDUs) sent by the initiator include the following statistics:

SCSI-CMD	SCSI-level command descriptor blocks.
LoginReq	Login request PDUs sent by initiators during session setup.

DataOut	PDU's containing write operation data that did not fit within the PDU of the SCSI command. The PDU maximum size is set by the storage system during the operation negotiation phase of the iSCSI login sequence.
Nop-Out	A message sent by initiators to check whether the target is still responding.
Logout-Req	Request sent by initiators to terminate active iSCSI sessions or to terminate one connection of a multi-connection session.
SNACK	A PDU sent by the initiator to acknowledge receipt of a set of DATA_IN PDU's or to request retransmission of specific PDU's.
SCSI TaskMgtCmd	SCSI-level task management messages, such as ABORT_TASK and RESET_LUN.
Text-Req	Text request PDU's that initiators send to request target information and renegotiate session parameters.

iSCSI PDU's transmitted

The iSCSI PDU's sent by the storage system include the following statistics:

SCSI-Rsp	SCSI response messages.
LoginRsp	Responses to login requests during session setup.
DataIn	Messages containing data requested by SCSI read operations.
Nop-In	Responses to initiator Nop-Out messages.
Logout-Rsp	Responses to Logout-Req messages.
R2T	Ready to transfer messages indicating that the target is ready to receive data during a SCSI write operation.
SCSI TaskMgtRsp	Responses to task management requests.
TextRsp	Responses to Text-Req messages.
Asynmsg	Messages the target sends to asynchronously notify the initiator of an event, such as the termination of a session.
Reject	Messages the target sends to report an error condition to the initiator, for example: • Data Digest Error (checksum failed) • Target does not support command sent by the initiator • Initiator sent a command PDU with an invalid PDU field

iSCSI CDBs

You can obtain statistics associated with handling iSCSI Command Descriptor Blocks, including the number of blocks of data transferred, and the number of SCSI-level errors and successful completions.

iSCSI Errors

You can obtain a list of login failures and other SCSI protocol errors.

Displaying iSCSI session information

You can use the `iscsi session show` command to display iSCSI session information, such as TCP connection information and iSCSI session parameters.

About this task

An iSCSI session can have zero or more connections. Typically a session has at least one connection. Connections can be added and removed during the life of the iSCSI session.

You can display information about all sessions or connections, or only specified sessions or connections. The `iscsi session show` command displays session information, and the `iscsi connection show` command displays connection information. The session information is also available through System Manager.

The command line options for these commands control the type of information displayed. For troubleshooting performance problems, the session parameters (especially HeaderDigest and DataDigest) are particularly important. The `-v` option displays all available information. In System Manager, the iSCSI Session Information page has buttons that control which information is displayed.

Step

1. On the storage system console, enter the following command:

```
iscsi session show [-v | -t | -p | -c] [session_tsih ...]
```

`-v` displays all information and is equivalent to `-t -p -c`.

`-t` displays the TCP connection information for each session.

`-p` displays the iSCSI session parameters for each session.

`-c` displays the iSCSI commands in progress for each session.

`session_tsih` is a list of session identifiers, separated by spaces.

```
system1> iscsi session show -t
Session 2
Initiator Information
Initiator Name: iqn.1991-05.com.microsoft:legbreak
```

```

ISID: 40:00:01:37:00:00
Connection Information
Connection 1
  Remote Endpoint: fe80::211:43ff:fece:ccce:1135
  Local Endpoint: fe80::2a0:98ff:fe00:fd81:3260
  Local Interface: e0
  TCP recv window size: 132480
Connection 2
  Remote Endpoint: 10.60.155.31:2280
  Local Endpoint: 10.60.155.105:3260
  Local Interface: e0
  TCP recv window size: 131400

```

Displaying iSCSI connection information

You can use the `iscsi connection show` command to display iSCSI connection parameters.

Step

1. On the storage system console, enter the following command:

```
iscsi connection show [-v] [{new | session_tsih} conn_id]
```

`-v` displays all connection information.

`newconn_id` displays information about a single connection that is not yet associated with a session identifier. You must specify both the keyword `new` and the connection identifier.

`session_tsih conn_id` displays information about a single connection. You must specify both the session identifier and the connection identifier.

Example

The following example shows the `-v` option.

```

system1> iscsi connection show -v
No new connections
Session connections
Connection 2/1:
  State: Full_Feature_Phase
  Remote Endpoint: fe80::211:43ff:fece:ccce:1135
  Local Endpoint: fe80::2a0:98ff:fe00:fd81:3260
  Local Interface: e0
Connection 2/2:
  State: Full_Feature_Phase
  Remote Endpoint: 10.60.155.31:2280
  Local Endpoint: 10.60.155.105:3260
  Local Interface: e0

```

Guidelines for using iSCSI with HA pairs

To ensure that the partner storage system successfully takes over during a failure, you need to make sure that the two systems and the TCP/IP network are correctly configured.

Of special concern are the target portal group tags configured on the two storage systems.


The best practice is to configure the two partners of the HA pair identically:

- You should use the same network cards in the same slots.
- You should create the same networking configuration with the matching pairs of ports connected to the same subnets.
- You should put the matching pairs of interfaces into the matching target portal groups and assign the same tag values to both groups.

Simple HA pairs with iSCSI

The following example describes how to implement the best practices for using iSCSI with HA pairs.

Consider the following simplified example. Storage System A has a two-port Ethernet card in slot 9. Interface e9a has the IP address 10.1.2.5, and interface e9b has the IP address 10.1.3.5. The two interfaces belong to a user-defined target portal group with tag value 2.



Storage System B has the same Ethernet card in slot 9. Interface e9a is assigned 10.1.2.6, and e9b is assigned 10.1.3.6. The two interfaces are in a user-defined target portal group with tag value 2.

In the HA pair, interface e9a on Storage System A is the partner of e9a on Storage System B. Likewise, e9b on System A is the partner of e9b on system B. For more information on configuring interfaces for an HA pair, see the *Data ONTAP High Availability and MetroCluster Configuration Guide for 7-Mode*.

Now assume that Storage System B fails and its iSCSI sessions are dropped. Storage System A assumes the identity of Storage System B. Interface e9a now has two IP addresses: its original address of 10.1.2.5, and the 10.1.2.6 address from Storage System B. The iSCSI host that was using Storage System B reestablishes its iSCSI session with the target on Storage System A.

If the e9a interface on Storage System A was in a target portal group with a different tag value than the interface on Storage System B, the host might not be able to continue its iSCSI session from Storage System B. This behavior varies depending on the specific host and initiator.

To ensure correct CFO behavior, both the IP address and the tag value must be the same as on the failed system. And because the target portal group tag is a property of the interface and not the IP address, the surviving interface cannot change the tag value during a CFO.

Related information

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

Complex HA pairs with iSCSI

If your cluster has a more complex networking configuration, including interface groups and VLANs, follow the best practice of making the configurations identical.

For example, if you have an interface group on storage system A, create the same interface group on storage system B. Ensure that the target portal group tag assigned to each interface group is the same. The name of the target portal group does not have to be the same; only the tag value matters.

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	You should verify that there is TCP/IP connectivity between the host and the storage system. - 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.

Configuration setting	What to do
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/NOW/products/interoperability .
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.
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: • <code>lun show -m</code> Displays all LUNs and the igroups to which they are mapped. • <code>lun show -g igroup-name</code> Displays the LUNs mapped to a specific igroup.
iSCSI ports enable	You should check if iSCSI ports are enabled or disabled.

Related concepts

[igroup management](#) on page 63

[Setting up LUNs and igroups](#) on page 44

Related tasks

[Verifying that the iSCSI service is running](#) on page 83

[Displaying initiators connected to the storage system](#) on page 95

Related information

[Interoperability Matrix: support.netapp.com/NOW/products/interoperability](http://support.netapp.com/NOW/products/interoperability)

System cannot register with iSNS server

Different iSNS server versions follow different draft levels of the iSNS specification.

If there is a mismatch between the iSNS draft version used by the storage system and by the iSNS server, the storage system cannot register.

Related concepts

[About iSNS service version incompatibility](#) on page 92

No multi-connection session

All of the connections in a multi-connection iSCSI session must go to interfaces on the storage system that are in the same target portal group.

If an initiator is unable to establish a multi-connection session, check the portal group assignments of the initiator.

If an initiator can establish a multi-connection session, but not during a cluster failover (CFO), the target portal group assignment on the partner storage system is probably different from the target portal group assignment on the primary storage system.

Related concepts

[Target portal group management](#) on page 105

[Guidelines for using iSCSI with HA pairs](#) on page 122

Sessions constantly connecting and disconnecting during takeover

An iSCSI initiator that uses multipath I/O constantly connects and disconnect from the target during cluster failover if the target portal group is not correctly configured.

The interfaces on the partner storage system must have the same target portal group tags as the interfaces on the primary storage system.

Related concepts

[Guidelines for using iSCSI with HA pairs](#) on page 122

Resolving iSCSI error messages on the storage system

There are a number of common iSCSI-related error messages that might display on your storage system console.

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

Message	Explanation	What to do
ISCSI: network interface <i>identifier</i> disabled for use; incoming connection discarded	The iSCSI service is not enabled on the interface.	You can use the <code>iscsi interface enable</code> command to enable the iSCSI service on the interface. For example: <code>iscsi interface enable e9b</code>
ISCSI: Authentication failed for initiator <i>nodename</i>	CHAP is not configured correctly for the specified initiator.	Check CHAP settings. • 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.
ifconfig: <i>interface</i> cannot be configured: Address does not match any partner interface. or Cluster monitor: takeover during ifconfig_2 failed; takeover continuing...	A single-mode interface group can be a partner interface to a standalone, physical interface on a cluster partner. However, the partner statement in the <code>ifconfig</code> command must use the name of the partner interface, not the partner's IP address. If the IP address of the partner's physical interface is used, the interface is not successfully taken over by the storage system's interface group.	1. Add the partner's interface using the <code>ifconfig</code> command on each system in the HA pair. For example: <pre>system1> ifconfig vif0 partner e0a system2> ifconfig e0a partner vif0</pre> 2. Modify the <code>/etc/rc</code> file on both systems to contain the same interface information.

Related concepts

Guidelines for using CHAP authentication on page 97

FC SAN management

You need to know some critical information that is required to successfully manage your FC SAN.

How to manage FC with HA pairs

Data ONTAP provides important functionality that allows your system to continue running smoothly when one of the devices in your HA pairs fails.

For additional configuration details, see the *Data ONTAP SAN Configuration Guide for 7-Mode*.

Related information

[Configuration and hardware guides on the NetApp Support Site-support.netapp.com](http://support.netapp.com)


How controller failover works

Data ONTAP is equipped with functionality called controller failover that allows the partner system to assume responsibility for the failed system's LUNs without interruption.

A storage system with an HA pair has a single global WWNN, and both systems in the HA pair function as a single FC node. Each node in the HA pair shares the partner node's LUN map information.

All LUNs in the HA pair are available on all ports in the HA pair by default. As a result, there are more paths to LUNs stored on the HA pair because any port on each node can provide access to both local and partner LUNs. You can specify the LUNs available on a subset of ports by defining port sets and binding them to an igroup. Any host in the igroup can access the LUNs only by connecting to the target ports in the port set.

The following figure shows an example configuration with a multi-attached host. If the host accesses lun_1 through ports 4a, 4b, 5a, or 5b on storage system X, then storage system X recognizes that lun_1 is a local LUN. If the host accesses lun_1 through any of the ports on storage system Y, lun_1 is recognized as a partner LUN and storage system Y sends the SCSI requests to storage system X over the HA pair interconnect.



How Data ONTAP avoids igroup mapping conflicts during cluster failover

Each node in the HA pair shares its partner's igroup and LUN mapping information. Data ONTAP uses the cluster interconnect to share igroup and LUN mapping information and also provides the mechanisms for avoiding mapping conflicts.

Related tasks

[Checking LUN, igroup, and FC settings](#) on page 56

[Bringing LUNs online](#) on page 53

igroup ostype conflicts

When you add an initiator WWPN to an igroup, Data ONTAP verifies that there are no igroup ostype conflicts.

An ostype conflict occurs, for example, when an initiator with the WWPN 10:00:00:00:c9:2b:cc:39 is a member of an AIX igroup on one node in the HA pair and the same WWPN is also a member of an group with the default ostype on the partner.

Reserved LUN ID ranges

By reserving LUN ID ranges on each storage system, Data ONTAP provides a mechanism for avoiding mapping conflicts.

If the cluster interconnect is down, and you try to map a LUN to a specific ID, the `lun map` command fails. If you do not specify an ID in the `lun map` command, Data ONTAP automatically assigns one from a reserved range.

The LUN ID range on each storage system is divided into three areas:

- IDs 0 to 192 are shared between the nodes. You can map a LUN to an ID in this range on either node in the HA pair.
- IDs 193 to 224 are reserved for one storage system.
- IDs 225 to 255 are reserved for the other storage system in the HA pair.

When to override possible mapping conflicts

When the cluster interconnect is down, Data ONTAP cannot check for LUN mapping or igroup ostype conflicts.

The following commands fail unless you use the `-f` option to force these commands. The `-f` option is only available with these commands when the cluster interconnect is down.

- `lun map`
- `lun online`
- `igroup add`
- `igroup set`

You might want to override possible mapping conflicts in disaster recovery situations or situations in which the partner in the HA pair cannot be reached and you want to regain access to LUNs. For example, the following command maps a LUN to an AIX igroup and assigns a LUN ID of 5, regardless of any possible mapping conflicts:

```
lun map -f /vol/vol2/qtree1/lun3 aix_host5_group2 5
```

Multipathing requirements for controller failover

Multipathing software is required on the host so that SCSI commands fail over to alternate paths when links go down due to switch failures or controller failovers. In the event of a failover, none of the adapters on the takeover storage system assume the WWPNs of the failed storage system.

How to use port sets to make LUNs available on specific FC target ports

A port set consists of a group of FC target ports. You bind a port set to an igroup, to make the LUN available only on a subset of the storage system's target ports. Any host in the igroup can access the LUNs only by connecting to the target ports in the port set.

If an igroup is not bound to a port set, the LUNs mapped to the igroup are available on all of the storage system's FC target ports. The igroup controls which initiators LUNs are exported to. The port set limits the target ports on which those initiators have access.

You use port sets for LUNs that are accessed by FC hosts only. You cannot use port sets for LUNs accessed by iSCSI hosts.

How port sets work in HA pairs

All ports on both systems in the HA pairs are visible to the hosts. You can use port sets to fine-tune which ports are available to specific hosts and limit the amount of paths to the LUNs to comply with the limitations of your multipathing software.

When using port sets, ensure that your port set definitions and igroup bindings align with the cabling and zoning requirements of your configuration. For additional configuration details, see the *Data ONTAP SAN Configuration Guide for 7-Mode*.

Related concepts

[How controller failover works](#) on page 127

Related information

[Configuration and hardware guides on the NetApp Support Site-support.netapp.com](#)

How upgrades affect port sets and igroups

When you upgrade to Data ONTAP 7.1 and later, all ports are visible to all initiators in the igroups until you create port sets and bind them to the igroups.

How port sets affect igroup throttles

Port sets enable you to control queue resources on a per-port basis.

If you assign a throttle reserve of 40 percent to an igroup that is not bound to a port set, then the initiators in the igroup are guaranteed 40 percent of the queue resources on every target port. If you bind the same igroup to a port set, then the initiators in the igroup have 40 percent of the queue

resources only on the target ports in the port set. This means that you can free up resources on other target ports for other igroups and initiators.

Before you bind new port sets to an igroup, verify the igroup's throttle reserve setting by using the `igroup show -t` command. It is important to check existing throttle reserves because you cannot assign more than 99 percent of a target port's queue resources to an igroup. When you bind more than one igroup to a port set, the combined throttle reserve settings might exceed 100 percent.

Example: port sets and igroup throttles

`igroup_1` is bound to `portset_1`, which includes ports 4a and 4b on each system in the HA pair (SystemA:4a, SystemA:4b, SystemB:4a, SystemB:4b). The throttle setting of `igroup` is 40 percent.

You create a new igroup (`igroup_2`) with a throttle setting of 70 percent. You bind `igroup_2` to `portset_2`, which includes ports 4b on each system in the HA pair (SystemA:4b, SystemB:4b). The throttle setting of the igroup is 70 percent. In this case, ports 4b on each system are overcommitted. Data ONTAP prevents you from binding the port set and displays a warning message prompting you to change the igroup throttle settings.

It is also important to check throttle reserves before you unbind a port set from an igroup. In this case, you make the ports visible to all igroups that are mapped to LUNs. The throttle reserve settings of multiple igroups might exceed the available resources on a port.

Creating port sets

You can use the `portset create` command to create port sets for FCP.

About this task

For HA pairs, when you add local ports to a port set, also add the partner system's corresponding target ports to the same port set.

For example, if you have local systems' target port 4a port in the port set, then ensure that you include the partner system's port 4a in the port set as well. This ensures that the takeover and giveback occurs without connectivity problems.

Step

1. Enter the following command:

```
portset create portset_name filename:slotletter
```

Example

```
portset create portset4 filerA:4b
```

portset_name is the name you specify for the port set. You can specify a string of up to 95 characters.

You should specify a port by using one of the following formats:

- *slotletter* is the slot and letter of the port—for example, 4b. If you use the slotletter format and the system is in an HA pair, the port from both the local and partner storage system is added to the port set.
- *filename:slotletter* adds only a specific port on a storage system—for example, SystemA:4b.

Binding igroups to port sets

After you create a port set, you must bind the port set to an igroup so the host knows which FC ports to access.

About this task

If you do not bind an igroup to a port set, and you map a LUN to the igroup, then the initiators in the igroup can access the LUN on any port on the storage system.

Note: You cannot bind an igroup to an empty port set, as the initiators in the igroup would have no ports by which to access the LUN.

Step

1. Enter the following command:

```
igroup bind igroup_name portset_name
```

Example

```
igroup bind aix-igroup1 portset4
```

Unbinding igroups from port sets

You can use the `igroup unbind` command to unbind an igroup from a port set.

About this task

If you unbind or unmap an igroup from a port set, then all the host initiator ports in the igroup can access LUNs on all target ports.

Step

1. Enter the following command:

```
igroup unbind igroup_name
```

Example

```
igroup unbind aix-igroup1
```

Adding ports to port sets

After you create a port set, you can use the `portset add` command to add ports to the port set.

Step

1. Enter the following command:

```
portset add portset_name [port...]
```

portset_name is the name you specify for the port set. You can specify a string of up to 95 characters.

port is the target FCP port. You can specify a list of ports. If you do not specify any ports, then you create an empty port set. You can add as many as 18 target FCP ports.

You specify a port by using the following formats:

- *slotletter* is the slot and letter of the port—for example, 4b. If you use the slotletter format and the system is in an HA pair, the port from both the local and partner storage system is added to the port set.
- *filename:slotletter* adds only a specific port on a storage system—for example, SystemA:4b.

Removing ports from port sets

After you create a port set, you can use the `portset remove` command to remove ports from the port set.

About this task

Note that you cannot remove the last port in the port set if the port set is bound to an igroup. To remove the last port, you must first unbind the port set from the igroup, then remove the port.

Step

1. Enter the following command:

```
portset remove portset_name [port...]
```

portset_name is the name you specify for the port set. You can specify a string of up to 95 characters.

port is the target FCP port. You can specify a list of ports. If you do not specify any ports, then you create an empty port set. You can add as many as 18 target FCP ports.

You can specify a port by using the following format:

- *slotletter* is the slot and letter of the port—for example, 4b. If you use the *slotletter* format and the system is in an HA pair, the port from both the local and partner storage system is added to the port set.

Destroying port sets

You can use the `portset destroy` command to delete a port set.

Steps

1. Unbind the port set from any igroups by entering the following command:

```
igroup unbind igroup_name portset_name
```

2. Enter the following command:

```
portset destroy [-f] portset_name...
```

You can specify a list of port sets.

If you use the `-f` option, you can destroy the port set even if it is still bound to an igroup.

If you do not use the `-f` option and the port set is still bound to an igroup, the `portset destroy` command fails.

Example

```
portset destroy portset1 portset2 portset3
```

Displaying the ports in a port set

You can use the `portset show` command to display all ports belonging to a particular port set.

Step

1. Enter the following command:

```
portset show portset_name
```

If you do not supply *portset_name*, all port sets and their respective ports are listed. If you supply *portset_name*, only the ports in the port set are listed.

Example

```
portset show portset1
```

Displaying igroup-to-port-set bindings

You can use the `igroup show` command to display which igroups are bound to port sets.

Step

1. Enter the following command:

```
igroup show igroup_name
```

Example

```
igroup show aix-igroup1
```

FC service management

You can use the `fc` commands for most of the tasks involved in managing the FC service and the target and initiator adapters.

You should enter `fc help` at the command line to display the list of available commands.

Verifying that the FC service is running

If the FC service is not running, target expansion adapters are automatically taken offline. They cannot be brought online until the FC service is started.

Step

1. Enter the following command:

```
fc status
```

A message is displayed indicating whether FC service is running.

Note: If the FC service is not running, you must verify that FC is licensed on the system.

Verifying that the FC service is licensed

If you cannot start the FC service, you should verify that the service is licensed on the system.

Step

1. Enter the following command:

```
license
```

Example

```
filer> license
Serial Number: 8000022008
Owner: ssan-6240-4b
Package          Type      Description          Expiration
-----
iSCSI            license  iSCSI License       -
FCP              license  FCP License         -
```

Displays the list of all services that are licensed and the details about the license package in Type, Description, and Expiration columns. This command does not display the services that are not licensed.

Enabling the FC license

Before you can use the FC target service, you must enable the FC license by entering the FC license key and turning on the `fc` option.

Steps

1. Enter the following command to add your FC license key:

```
license add fcp_license_code
```

Example

```
filer> license add QVCPRDBBVAAAAAAAAAAAAAAAAAAAA
license add: successfully added license key
"QVCPRDBBVAAAAAAAAAAAAAAAAAAAA".
```

2. Enter the following command to enable the `fc` option:

```
options licensed_feature.fcp.enable on
```

Example

```
filer> options licensed_feature.fcp.enable on
Thu Feb 14 16:09:50 EST [filer:kern.cli.cmd:debug]: Command line
input: the command is "options".
The full command line is "options licensed_feature.fcp.enable on".
cf.takeover.on_panic is already on
cf.takeover.on_reboot is changed to off
Run 'fcp start' to start the FCP service.
Also run "lun setup" if necessary to configure LUNs.
filer > fcp start
Thu Feb 14 16:09:54 EST [filer:fcp.service.startup:info]: FCP service
startup
```

Disabling the FC license

To disable the FC license, you must remove the FC license key and turn off the `fc` option.

About this task

Note: If you disable the FC license, you cannot access the FC service and FC target connectivity is lost. Therefore, any LUNs being served to the initiators are terminated.

Steps

1. Enter the following command to remove your FC license key:

```
license delete fcp
```

Example

```
filer> license delete fcp
license delete: successfully deleted "fcp".
```

2. Enter the following command to disable the fcp option:

```
options licensed_feature.fcp.enable off
```

Example

```
filer > options licensed_feature.fcp.enable off
Thu Feb 14 16:11:09 EST [filer:kern.cli.cmd:debug]: Command line
input: the command is 'options'.
The full command line is 'options licensed_feature.fcp.enable off'.
Thu Feb 14 16:11:09 EST [filer:fcp.service.shutdown:info]: FCP
service shutdown
```

Starting and stopping the FC service

After the FC service is licensed, you can start and stop the service.

About this task

Stopping the FC service disables all FC ports on the system, which has important ramifications for HA pairs during cluster failover. For example, if you stop the FC service on System1, and System2 fails over, System1 is unable to service System2's LUNs.

On the other hand, if System2 fails over, and you stop the FC service on System2 and start the FC service on System1, System1 successfully services System2's LUNs.

You can use the `partner fcp stop` command to disable the FC ports on the failed system during takeover, and use the `partner fcp start` command to re-enable the FC service after the giveback is complete.

Step

1. Enter the following command:

```
fcp [start|stop]
```

Example**fcp start**

The FC service is enabled on all FC ports on the system. If you enter `fcp stop`, the FC service is disabled on all FC ports on the system.

Taking target expansion adapters offline and bringing them online

You can use the `fcv config` command to take a target expansion adapter offline and to bring it back online.

Step

1. Enter the following command:

```
fcv config adapter [up|down]
```

Example

```
fcv config 4a down
```

The target adapter 4a is offline. If you enter `fcv config 4a up`, the adapter is brought online.

Changing the adapter speed

You can use the `fcv config` command to change the FC adapter speed.

About this task

The available speeds depend on the HBA being used. The following is a list of the supported speeds available to the controllers:

- Autonegotiate (default)
- 1 Gb
- 2 Gb
- 4 Gb
- 8 Gb

Steps

1. Set the adapter to down by using the following command:

```
fcv config adapter down
```

Example

```
: system1> fcv config 2a down
: Wed Jun 15 14:04:47 GMT [device1:
: scsitarget.ispfct.offlineStart:notice]:
: Offlining Fibre Channel target adapter 2a.
: Wed Jun 15 14:04:47 GMT [device1:
: scsitarget.ispfct.offlineComplete:notice]: Fibre Channel
: target adapter
: 2a offlined.
```

Adapter 2a is taken down, and the FC service might be temporarily interrupted on the adapter.

2. Enter the following command:

```
fcv config adapter speed [auto|1|2|4|8|10]
```

Example

```
: system1> fcv config 2a speed 2
```

The speed for adapter 2a is changed to 2 Gb per second.

3. Enter the following command:

```
fcv config adapter up
```

Example

```
: device1> fcv config 2a up
: Wed Jun 15 14:05:04 GMT [device1:
scsitarget.ispfct.onlining:notice]:
: Onlining Fibre Channel target adapter 2a.

: device1> fcv config
: 2a:  ONLINE [ADAPTER UP]  Loop  No Fabric
:      host address 0000da
:      portname 50:0a:09:81:96:97:a7:f3  nodename
: 50:0a:09:80:86:97:a7:f3
mediatype auto speed 2Gb
```

Adapter 2a is brought back up and the speed is 2 Gb per second.

After you finish

Although the `fcv config` command displays the current adapter speed setting, it does not necessarily display the actual speed at which the adapter is running. For example, if the speed is set to auto, the actual speed might be 1 Gb, 2 Gb, 4 Gb, and so on.

You can use the `show adapter -v` command to view the following:

- Actual speed at which the adapter is running and examine the Data Link Rate value
- Switchname and port number

```
system1> fcv show adapter -v 4a
Slot: 4a
Description: Fibre Channel Target Adapter 4a (Dual-channel,
QLogic CNA 8112 (8152) rev. 2)
Status: ONLINE
Host Port Address: 0x98d601
Firmware Rev: 5.3.4
MPI Firmware Rev: 1.38.0
PHY Firmware Rev: 1.7.0
FC VLAN ID: 5
FC Nodename: 50:0a:09:80:87:69:68:5a (500a09808769685a)
```

```

FC Portname:          50:0a:09:81:87:69:68:5a (500a09818769685a)
Cacheline Size:       16
FC Packet Size:       2048
SRAM Parity:          Yes
External GBIC:        No
Data Link Rate:    10 GBit
Adapter Type:         Local
Fabric Established:    Yes
Connection Established: PTP
Mediatype:            auto
Partner Adapter:      None
Standby:              No
Target Port ID:        0x1
Switch Port:       brcddcx_rtp02:214
Physical Link Rate:    10 GBit
Physical Link Status:  LINK UP

```

How WWPN assignments work with FC target expansion adapters

It is important to understand how WWPN assignments work with FC target expansion adapters so that your systems continue to run smoothly in the event of head swaps and upgrades, new adapter installations, and slot changes for existing adapters.

When the FC service is initially licensed and enabled on your storage system, the FC target expansion adapters are assigned WWPNs, which persist through head upgrades and replacements. The assignment information is stored in the system's root volume.

The WWPN is associated with the interface name. For example, a target expansion adapter installed in slot 2 might have the interface name of 2a and a WWPN of 50:0a:09:81:96:97:c3:ac. Since the WWPN assignments are persistent, a WWPN is not automatically re-used, even if the port is disabled or removed. However, there are some circumstances under which you might have to manually change the WWPN assignments.

The following examples explain how WWPN assignments work under the most common circumstances:

- Swapping or upgrading a head
- Adding a new FC target expansion adapter
- Moving an existing adapter to a different slot

Swapping or upgrading a head

As long as the existing root volume is used in the head swap or upgrade, the same port-to-WWPN mapping applies. For example, port 0a on the replacement head has the same WWPN as the original head. If the new head has different adapter ports, the new ports are assigned new WWPNs.

Adding new FC target expansion adapters

If you add a new adapter, the new ports are assigned new WWPNs. If you replace an existing adapter, the existing WWPNs are assigned to the replacement adapter.

For example, the following table shows the WWPN assignments if you replace a dual-port adapter with a quad-port adapter.

Original configuration	New configuration	WWPN assignments
2a - 50:0a:09:81:96:97:c3:ac	2a - 50:0a:09:81:96:97:c3:ac	No change
2b - 50:0a:09:83:96:97:c3:ac	2b - 50:0a:09:83:96:97:c3:ac	No change
	2c - 50:0a:09:82:96:97:c3:ac	New
	2d - 50:0a:09:84:96:97:c3:ac	New

Moving a target expansion adapter to a different slot

If you move an adapter to a new slot, then adapter is assigned new WWPNs.

Original configuration	New configuration	WWPN assignments
2a - 50:0a:09:81:96:97:c3:ac	4a - 50:0a:09:85:96:97:c3:ac	New
2b - 50:0a:09:83:96:97:c3:ac	4b - 50:0a:09:86:96:97:c3:ac	New

Related tasks

[Changing the WWPN for a target adapter](#) on page 141

Changing the WWPN for a target adapter

Data ONTAP automatically sets the WWPNs on your target adapters during initialization. However, there are some circumstances in which you might need to change the WWPN assignments on your target expansion adapters or your onboard adapters.

About this task

There are two scenarios that might require you to change the WWPN assignments:

- **Head swap:** after performing a head swap, you might not be able to place the target adapters in their original slots, resulting in different WWPN assignments. In this situation it is important to change the WWPN assignments because many of the hosts bind to these WWPNs. In addition, the fabric might be zoned by WWPN.

- Fabric reorganization: you might want to reorganize the fabric connections without having to physically move the target adapters or modify your cabling.

Sometimes, you might need to set the new WWPN on a single adapter. In other cases, it is easier to swap the WWPNs between two adapters, rather than individually set the WWPNs on both adapters.

Steps

1. Take the adapter offline by entering the following command:

```
fcpl config adapter down
```

Example

```
fcpl config 4a down
```

Note: If you are swapping WWPNs between two adapters, ensure that you take both adapters offline first.

2. Display the existing WWPNs by entering the following command:

```
fcpl portname show [-v]
```

If you do not use the `-v` option, all currently used WWPNs and their associated adapters are displayed. If you use the `-v` option, all other valid WWPNs that are not being used are also shown.

3. Set the new WWPN for a single adapter or swap WWPNs between two adapters.

Note: If you do not use the `-f` option, initiators might fail to reconnect to this adapter if the WWPN is changed. If you use the `-f` option, it overrides the warning message of changing the WWPNs.

If you want to...	Then...
Set the WWPN on a single adapter	Enter the following command: fcpl portname set [-f] adapter wwpn
Swap WWPNs between two adapters.	Enter the following command: fcpl portname swap [-f] adapter1 adapter2

Example

```
fcpl portname set -f 1b 50:0a:09:85:87:09:68:ad
```

Example

```
fcpl portname swap -f 1a 1b
```

4. Bring the adapter back online by entering the following command:

```
fcpl config adapter up
```

Example

```
fcp config 4a up
```

Related concepts

How WWPN assignments work with FC target expansion adapters on page 140

Changing the system's WWNN

The WWNN of a storage system is generated by a serial number in its NVRAM, but it is stored on the disk. If you ever replace a storage system chassis and reuse it in the same FC SAN, it is possible, although extremely rare, that the WWNN of the replaced storage system is duplicated. In this unlikely event, you can change the WWNN of the storage system.

About this task

Attention: You must change the WWNN on both systems. If both systems do not have the same WWNN, hosts cannot access LUNs on the same HA pair.

Step

1. Enter the following command:

```
fcp nodename [-f]nodename
```

nodename is a 64-bit WWNN address in the following format: 50:0a:09:80:8X:XX:XX:XX, where X is a valid hexadecimal value.

You can use `-f` to force the system to use an invalid nodename. You should not, under normal circumstances, use an invalid nodename.

Example

```
fcp nodename 50:0a:09:80:82:02:8d:ff
```

WWPN aliases

A WWPN is a unique, 64-bit identifier displayed as a 16-character hexadecimal value in Data ONTAP. However, SAN Administrators may find it easier to identify FC ports using an alias instead, especially in larger SANs.

You can use the `wwpn-alias` sub-command to create, remove, and display WWPN aliases.

Creating WWPN aliases

You can use the `fcp wwpn-alias set` command to create a new WWPN alias.

About this task

You can create multiple aliases for a WWPN, but you cannot use the same alias for multiple WWPNs. The alias can consist of up to 32 characters and can contain only the letters A through Z, a

through z, numbers 0 through 9, hyphen ("-"), underscore ("_"), left brace ("{"), right brace ("}"), and period (".").

Step

1. Enter the following command:

```
fcpx wwpn-alias set [-f] alias wwpn
```

-f allows you to override a WWPN associated with an existing alias with the newly specified WWPN.

Example

```
fcpx wwpn-alias set my_alias_1 10:00:00:00:c9:30:80:2f
```

Example

```
fcpx wwpn-alias set -f my_alias_1 11:11:00:00:c9:30:80:2e
```

Removing WWPN aliases

You can use the `fcpx wwpn-alias remove` command to remove an alias for a WWPN.

Step

1. Enter the following command:

```
fcpx wwpn-alias remove [-a alias ... | -w wwpn]
```

-a *alias* removes the specified aliases.

-w *wwpn* removes all aliases associated with the WWPN.

Example

```
fcpx wwpn-alias remove -a my_alias_1
```

Example

```
fcpx wwpn-alias remove -w 10:00:00:00:c9:30:80:2
```

Displaying WWPN alias information

You can use the `fcpx wwpn-alias show` command to display the aliases associated with a WWPN or the WWPN associated with an alias.

Step

1. Enter the following command:

```
fcpx wwpn-alias show [-a alias | -w wwpn]
```

- a *alias* displays the WWPN associated with the alias.
- w *wwpn* displays all aliases associated with the WWPN.

Example

```
fcp wwpn-alias show -a my_alias_1
```

Example

```
fcp wwpn-alias show -w 10:00:00:00:c9:30:80:2
```

Example

```
fcp wwpn-alias show
```

WWPN	Alias
----	-----
10:00:00:00:c9:2b:cb:7f	temp
10:00:00:00:c9:2b:cc:39	lrrr_1
10:00:00:00:c9:4c:be:ec	alias_0
10:00:00:00:c9:4c:be:ec	alias_0_temp
10:00:00:00:c9:2b:cc:39	lrrr_1_temp

Note: You can also use the `igroup show`, `igroup create`, `igroup add`, `igroup remove`, and `fcp show initiator` commands to display WWPN aliases.

Obtaining fabric zone server data

You can use the zone server to access zone membership as well as port information. The `fcp zone show` command enables you to view the active zone set on the fabric connected to the target port and to verify the zoning information on the fabric zone server.

About this task

Note: You should understand that not all FC switch vendors support the necessary fabric commands that are used to obtain zoning information.

Step

1. Obtain the fabric zone server data by entering the following command:

```
fcp zone show
```

Example: Fabric zone server data

```
system1> fcp zone show 4a
Active Zone Set on adapter 4a:
Zone Set Name: sanset (1 zones)
Zone Name: testzone
```

```
Member Port Name: 10:00:00:00:c9:2d:60:dc
Member Port Name: 50:0a:09:82:87:09:2b:7d
Member Port ID: 0x650003
Member Fabric Port Name: 20:07:00:0d:ec:00:22:80
```

Obtaining a physical topology of the FC fabric

The fabric configuration server provides information about the switches and their ports. This information can be used to generate a physical topology of the fabric.

Step

1. Obtain the physical topology of the fabric by entering the following command:

```
fcv topology show [adapter]
```

If no `adapter` is specified, topology information for all adapters is shown.

Obtaining fabric nameserver data

The *fabric nameserver* is the entity on the fabric that holds all information about devices in the fabric. The FC target sends a variety of defined FC commands to the nameserver to collect the fabric nameserver data.

Step

1. Obtain the fabric nameserver data by entering the following command:

```
fcv nameserver show
```

Example

```
system1> fcv nameserver show
Name Server database connected on adapter 0c:No entries found.

Name Server database connected on adapter 0d:No entries found.

Name Server database connected on adapter 1a:

Port ID           :0xe60c00
Port Type          :N-Port
Port Name          :50:0a:09:81:87:19:66:26
Node Name          :50:0a:09:80:87:19:66:26
Symbolic Port Name :NetApp FC Target Adapter (2532) system1:1a
Symbolic Node Name :NetApp                               FAS6040 (system1)
Fabric Port Name   :20:0c:00:05:1e:0f:7f:a5
Class of Service   :3
FC4 Type           :FCP
```

Checking connectivity of the initiators

You can use the `fc ping` command to check the connectivity of the initiators and to verify the correctness of zoning. This command can also be used to check fabric latency between the initiator and target by using the `-s` option.

Step

1. Check the connectivity and latency by using the following command:

fc ping

Example

```
system1> fc ping
0c 10:00:00:00:c9:46:dc:6d
10:00:00:00:c9:46:dc:6d (0xe71100) is alive

system1> fc ping -s 0c 10:00:00:00:c9:46:dc:6d
76 bytes from 10:00:00:00:c9:46:dc:6d (0xe71100): seq=0 time=0.203
ms
76 bytes from 10:00:00:00:c9:46:dc:6d (0xe71100): seq=1 time=0.438
ms
76 bytes from 10:00:00:00:c9:46:dc:6d (0xe71100): seq=2 time=0.414
ms
76 bytes from 10:00:00:00:c9:46:dc:6d (0xe71100): seq=3 time=0.246
ms
76 bytes from 10:00:00:00:c9:46:dc:6d (0xe71100): seq=4 time=0.196
ms
76 bytes from 10:00:00:00:c9:46:dc:6d (0xe71100): seq=5 time=0.305
ms
--- 10:00:00:00:c9:46:dc:6d ping statistics ---
6 frames transmitted, 6 frames received, 0% frame loss
```

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 *SAN Configuration Guide* (called *Fibre Channel and iSCSI Configuration Guide in Data ONTAP 8.1 and earlier*).

Related information

Configuration and hardware guides on the NetApp Support Site: support.netapp.com

Configuring onboard adapters for target mode

You can configure the onboard adapters for target mode to connect the adapters to the FC fabric or to another storage controller.

Before you begin

The FC protocol service must be licensed on the system.

About this task

Each onboard FC port can be individually configured as an initiator or a target. If you exceed the allowed number of adapter ports, you must set the onboard adapters to initiator or unconfigured before installing the expansion adapters. 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 or initiator FC port, just like the onboard FC ports:

- The NetApp X2056-R6 4-port 8-Gb FC adapter for the 62xx vertical I/O slots (slots 1, 11, and 12).
- The NetApp X1132A-R6 4-port 8-Gb FC adapter for the 62xx and the 32xx

Note: For detailed information about the number of target adapters supported on each hardware platform, see the *Data ONTAP SAN Configuration Guide for 7-Mode*.

Steps

1. Verify that the FC ports are not already configured as target ports by entering the following command:

fcadmin config

Example

```
fcadmin config
Local
Adapter Type                               State                               Status
-----
0a                               initiator CONFIGURED online
0b                               initiator CONFIGURED online
0c                               target                               CONFIGURED online
0d                               target                               CONFIGURED online
The preceding output displays two ports for host access.
-----
```

2. Set the status of the initiator port to offline by entering the following command:

fcadmin config -d *adapter_name*...

Example

```
fcadmin config -d 0a 0b
```

Ports 0a and 0b are offline.

3. Set the onboard ports to operate in target mode by entering the following command:

```
fcadmin config -t target adapter_name...
```

adapter_name is the port number. You can specify more than one port.

Example

```
fcadmin config -t target 0a 0b
```

Ports 0a and 0b are set to target mode.

4. Run the following command to see the change in state for the ports:

```
fcadmin config
```

Example

```
fcadmin config
```

Adapter	Type	Local State	Status
0a	initiator	PENDING (target)	online
0b	initiator	PENDING (target)	online
0c	target	CONFIGURED	online
0d	target	CONFIGURED	online

Note: The available Local State values are CONFIGURED, PENDING, and UNCONFIGURED. Refer to the `fcadmin` man page for detailed descriptions of each value.

Ports 0a and 0b are now in the PENDING state.

5. Reboot each system in the HA pair by entering the following command:

```
reboot
```

6. Verify that the FC ports are online and configured in the correct state for your configuration by entering the following command:

```
fcadmin config
```

Example

```
fcadmin config
```

Adapter	Type	Local State	Status
0a	target	CONFIGURED	online
0b	target	CONFIGURED	online

0c	target	CONFIGURED	online
0d	target	CONFIGURED	online

The preceding output displays for a four-port SAN configuration.

Related information

Configuration and hardware guides on the NetApp Support Site-support.netapp.com

Configuring onboard 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.

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 or initiator FC port, just like the onboard FC ports:

- The NetApp X2056-R6 4-port 8-Gb FC adapter for the 62xx vertical I/O slots (slots 1, 11, and 12).
- The NetApp X1132A-R6 4-port 8-Gb FC adapter for the 62xx and the 32xx.

For detailed information about the number of target adapters supported on each hardware platform, see the *Data ONTAP SAN Configuration Guide for 7-Mode*.

Steps

1. If you have already connected the port to a switch or fabric, take it offline by entering the following command:

```
fcp config -d adapter_name...
```

adapter_name is the port number. You can specify more than one port.

Example

```
fcp config -d 0c 0d
```

Ports 0c and 0d are taken offline.

```
Adapter Type State Status
-----
0a target CONFIGURED online
0b target CONFIGURED online
0c target CONFIGURED offline
0d target CONFIGURED offline
```

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

2. Set the onboard ports to operate in initiator mode by entering the following command:

```
fcadmin config -t initiator adapter_name...
```

adapter_name is the port number. You can specify more than one port.

Example

```
fcadmin config -t initiator 0c 0d
```

Ports 0c and 0d are set to initiator mode.

3. Run the following command to see the change in state for the ports:

```
fcadmin config
```

```
Adapter Type State Status
-----
0a target CONFIGURED online
0b target CONFIGURED online
0c target PENDING (initiator) offline
0d target PENDING (initiator) offline
```

4. Reboot each system in the HA pair by entering the following command:

```
reboot
```

5. Verify that the FC ports are online and configured in the correct state for your configuration by entering the following command:

```
fcadmin config
```

Example

```
fcadmin config
Adapter Type      Local State      Status
-----
0a target          CONFIGURED      online
0b target          CONFIGURED      online
0c initiator       CONFIGURED      online
0d initiator       CONFIGURED      online
```

Note: The available Local State values are CONFIGURED, PENDING, and UNCONFIGURED. Refer to the `fcadmin` man page for detailed descriptions of each value.

The preceding output displays for a four-port SAN configuration.

Commands for displaying adapter information

You can find the list of commands available for displaying information about adapters. The output varies depending on the storage system model.

If you want to display...	Use this command...
Information for all initiator adapters in the system, including firmware level, node name, FC packet size, link data rate, SRAM parity, and various states	storage show adapter
All adapter (HBAs, NICs, and switch ports) configuration and status information	sysconfig [-v] [adapter] <i>adapter</i> is a numerical value only. -v displays additional information about all adapters.
Disks, disk loops, and options configuration information that affects coredumps and takeover	sysconfig -c
FCP traffic information	sysstat -f
How long FCP has been running	uptime
Initiator HBA port address, port name, port name alias, node name, and igroup name connected to target adapters	fcp show initiator [-v] [adapter] -v displays the Fibre Channel host address of the initiator. <i>adapter</i> is the slot number with the port number, a or b; for example, 5a.
Service statistics	availtime
Target adapter configuration information	fcp config

If you want to display...	Use this command...
Target adapter node name, port name, and link state	fcv show adapter [-v][<i>adapter</i>] <i>adapter</i> is the slot number with the port number, a or b; for example, 5a. -v displays additional information about the adapters.
Target adapter statistics	fcv stats [-z] [<i>adapter</i>] -z zeros the statistics. <i>adapter</i> is the slot number with the port number, a or b; for example, 5a.
Information about FCP traffic along with the statistics from partner storage system	sysstat -b
WWNN of the target adapter	fcv nodename

Displaying the status of onboard FC adapters

You can use the `fcadmin config` command to determine the status of the onboard FC adapters.

About this task

This command also displays other important information, including the configuration status of the adapter and whether it is configured as a target or initiator.

Onboard FC adapters are set to initiator mode by default.

Step

1. Enter the following command:

```
fcadmin config
```

Example

```
fcadmin config
Adapter Type      Local
                  State
-----
 0a  initiator    CONFIGURED    online
 0b  initiator    CONFIGURED    online
 0c  target       CONFIGURED    online
 0d  target       CONFIGURED    online
```

Note: The available Local State values are CONFIGURED, PENDING, and UNCONFIGURED. Refer to the fcadmin man page for detailed descriptions of each value.

Displaying information about all adapters

You can use the `sysconfig -v` command to display system configuration and adapter information for all adapters in the system.

Step

1. Enter the following command:

sysconfig -v

Example

```
system1>sysconfig -v
slot 2: Fibre Channel Target Host Adapter 2a
          (Dual-channel, QLogic 2532 (2562) rev. 2, 32-bit,
[ONLINE])
          Firmware rev: 4.6.2
          Host Port Addr: 011200
          Cacheline size: 16
          SRAM parity: Yes
          FC Nodename: 50:0a:09:80:87:29:2a:42
(500a098087292a42)
          FC Portname: 50:0a:09:85:97:29:2a:42
(500a098597292a42)
          Connection: PTP, Fabric
          SFP Vendor Name: AVAGO
          SFP Vendor P/N: AFBR-57D5APZ
          SFP Vendor Rev: B
          SFP Serial No.: AD0820EA06W
          SFP Connector: LC
          SFP Capabilities: 2, 4, 8 Gbit/Sec
                        I/O base 0x0000000000008000, size 0x100
                        memory mapped I/O base 0xfe500000, size 0x4000
slot 2: Fibre Channel Target Host Adapter 2b
          (Dual-channel, QLogic 2532 (2562) rev. 2, 32-bit,
[ONLINE])
          Firmware rev: 4.6.2
          Host Port Addr: 011300
          Cacheline size: 16
          SRAM parity: Yes
          FC Nodename: 50:0a:09:80:87:29:2a:42
(500a098087292a42)
          FC Portname: 50:0a:09:86:97:29:2a:42
(500a098697292a42)
          Connection: PTP, Fabric
          SFP Vendor Name: AVAGO
          SFP Vendor P/N: AFBR-57D5APZ
          SFP Vendor Rev: B
          SFP Serial No.: AD0820EA0ES
          SFP Connector: LC
```

```
SFP Capabilities: 2, 4, 8 Gbit/Sec
                  I/O base 0x00000000000008400, size 0x100
                  memory mapped I/O base 0xfe504000, size 0x4000
```

System configuration information and adapter information for each slot that is used is displayed on the screen. Look for *Fibre Channel Target Host Adapter* to get information about target HBAs.

Note: In the output, in the information about the Dual-channel QLogic HBA, the value 2532 does not specify the model number of the HBA; it refers to the device ID set by QLogic. Also, the output varies according to storage system model.

Displaying brief target adapter information

You can use the `fcv config` command to display information about target adapters in the system, as well as to quickly detect whether the adapters are active and online.

About this task

The output of the `fcv config` command depends on the storage system model.

Step

1. Enter the following command:
fcv config

Example

The `fcv config` command displays the following output:

```
7a:  ONLINE [ADAPTER UP]  PTP  Fabric
      host address 170900
      portname 50:0a:09:83:86:87:a5:09
      nodename 50:0a:09:80:86:87:a5:09
      mediatype ptp  partner adapter 7a

7b:  ONLINE [ADAPTER UP]  PTP  Fabric
      host address 171800
      portname 50:0a:09:8c:86:57:11:22
      nodename 50:0a:09:80:86:57:11:22
      mediatype ptp  partner adapter 7b
```

Example

The following example shows output for the 30xx. The `fcv config` command displays information about the onboard ports connected to the SAN:

```
0c:  ONLINE [ADAPTER UP]  PTP  Fabric
      host address 010900
      portname 50:0a:09:81:86:f7:a8:42
      nodename 50:0a:09:80:86:f7:a8:42
```

```

mediatype ptp partner adapter 0d
0d:  ONLINE [ADAPTER UP]  PTP  Fabric
      host address 010800
      portname 50:0a:09:8a:86:47:a8:32
      nodename 50:0a:09:80:86:47:a8:32
      mediatype ptp partner adapter 0c

```

Displaying detailed target adapter information

You can use the `fcpl show adapter` command to display the node name, port name, and link state of all target adapters in the system.

About this task

Notice that the port name and node name are displayed with and without the separating colons. For Solaris hosts, you use the WWPN without separating colons when you map adapter port names (or these target WWPNs) to the host.

Step

1. Enter the following command:

```
fcpl show adapter -v
```

Example

```

system1> fcpl show adapter -v 4a
Slot: 4a
Description: Fibre Channel Target Adapter 4a (Dual-
channel, QLogic CNA 8112 (8152) rev. 2)
Status: ONLINE
Host Port Address: 0x98d601
Firmware Rev: 5.3.4
MPI Firmware Rev: 1.38.0
PHY Firmware Rev: 1.7.0
FC VLAN ID: 5
FC Nodename: 50:0a:09:80:87:69:68:5a (500a09808769685a)
FC Portname: 50:0a:09:81:87:69:68:5a (500a09818769685a)
Cacheline Size: 16
FC Packet Size: 2048
SRAM Parity: Yes
External GBIC: No
Data Link Rate: 10 GBit
Adapter Type: Local
Fabric Established: Yes
Connection Established: PTP
Mediatype: auto
Partner Adapter: None
Standby: No
Target Port ID: 0x1
Switch Port: brcdcdx_rtp02:214

```

```
Physical Link Rate:      10 GBit
Physical Link Status:    LINK UP
```

The information about the adapter in slot 4 displays.

Note: In the output, in the information about the Dual-channel QLogic HBA, the value 8112 does not specify the model number of the HBA; it refers to the device ID set by QLogic. Also, the output varies according to storage system model.

Following are the definitions of the possible values in the Status field:

Uninitialized	The firmware has not yet been loaded and initialized.
Link not connected	The driver has finished initializing the firmware. However, the link is not physically connected so the adapter is offline.
Online	The adapter is online for FC traffic.
Link disconnected	The adapter is offline due to a Fibre Channel link offline event.
Offline	The adapter is offline for FC traffic.
Offlined by user/system	A user manually took the adapter offline, or the system automatically took the adapter offline.

Displaying the WWNN of a target adapter

You can use the `fcpl nodename` command to display the WWNN of a target adapter in the system.

Step

1. Enter the following command:

```
fcpl nodename
```

Example

```
Fibre Channel nodename: 50:a9:80:00:02:00:8d:b2 (50a9800002008db2)
```

Displaying Initiator information

You can use the `fcpl show initiator` command to display the port names, aliases, and igroup names of HBAs connected to target adapters on the storage system.

Step

1. Enter the following command:

```
fcpl show initiator
```

Example

```

fcpl show initiator
Portname                Alias      Group
10:00:00:00:c9:32:74:28 calculon0 calculon
10:00:00:00:c9:2d:60:dc gaston0    gaston
10:00:00:00:c9:2b:51:1f

Initiators connected on adapter 0b: None connected.

```

Displaying target adapter statistics

You can use the `fcpl stats` command to display important statistics for the target adapters in your system.

Step

1. Enter the following command:

```
fcpl stats -i interval [-c count] [-a | adapter]
```

-i *interval* is the interval, in seconds, at which the statistics are displayed.

-c *count* is the number of intervals. For example, the `fcpl stats -i 10 -c 5` command displays statistics in ten-second intervals, for five intervals.

-a shows statistics for all adapters.

adapter is the slot and port number of a specific target adapter.

Example

```

system1> fcpl stats -i 1
r/s   w/s   o/s   ki/s   ko/s   asvc_t   qlen hba
0      0      0      0      0      0.00    0.00 7a
110    113    0      7104   12120   9.64    1.05 7a
146    68     0      6240   13488   10.28   1.05 7a
106    92     0      5856   10716   12.26   1.06 7a
136    102    0      7696   13964   8.65    1.05 7a

```

Each column displays the following information:

r/s—The number of SCSI read operations per second.

w/s—The number of SCSI write operations per second.

o/s—The number of other SCSI operations per second.

ki/s—Kilobytes of received traffic per second.

ko/s—Kilobytes of send traffic per second.

asvc_t—Average time in milliseconds to process a request

qlen—The average number of outstanding requests pending.

hba—The HBA slot and port number.

To see additional statistics, enter the `fcv stats` command with no variables.

Displaying FC traffic information

You can use the `sysstat -f` command to display FC traffic information, such as operations per second and kilobytes per second.

Step

1. Enter the following command:

```
sysstat -f
```

Example

CPU	NFS	CIFS	FCP	Net in	kB/s out	Disk read	kB/s write	FCP in	kB/s out	Cache age
81%	0	0	6600	0	0	105874	56233	40148	232749	1
78%	0	0	5750	0	0	110831	37875	36519	237349	1
78%	0	0	5755	0	0	111789	37830	36152	236970	1
80%	0	0	7061	0	0	107742	49539	42651	232778	1
78%	0	0	5770	0	0	110739	37901	35933	237980	1
79%	0	0	5693	0	0	108322	47070	36231	234670	1
79%	0	0	5725	0	0	108482	47161	36266	237828	1
79%	0	0	6991	0	0	107032	39465	41792	233754	1
80%	0	0	5945	0	0	110555	48778	36994	235568	1
78%	0	0	5914	0	0	107562	43830	37396	235538	1

The following columns provide information about FCP statistics:

CPU—The percentage of the time that one or more CPUs were busy.

FCP—The number of FCP operations per second.

FCP KB/s—The number of kilobytes per second of incoming and outgoing FCP traffic.

Displaying information about FC protocol traffic from the partner

If you have an HA pair, you might want to obtain information about the amount of traffic coming to the system from its partner.

Step

1. Enter the following command:

```
sysstat -b
```

The following show the columns information about partner traffic:

Partner—The number of partner operations per second.

Partner KB/s—The number of kilobytes per second of incoming and outgoing partner traffic.

Related concepts

[How to manage FC with HA pairs](#) on page 127

Displaying how long the FC service has been running

You can use the `uptime` command to display how long the FC service has been running on the system.

Step

1. Enter the following command:

```
uptime
```

Example

```
12:46am up 2 days, 8:59 102 NFS ops, 2609 CIFS ops, 0 HTTP ops, 0
DAFS ops, 1933084 FCP ops, 0 iSCSI ops
```

Displaying FC protocol service statistics

You can use the `availtime` command to display the FC protocol service statistics.

Step

1. Enter the following command:

```
availtime
```

Example

```
Service statistics as of Mon Jul 1 00:28:37 GMT 2002
System (UP). First recorded (3894833) on Thu May 16 22:34:44 GMT 2002
  P 28, 230257, 170104, Mon Jun 10 08:31:39 GMT 2002
  U 24, 131888, 121180, Fri Jun 7 17:39:36 GMT 2002
NFS (UP). First recorded (3894828) on Thu May 16 22:34:49 GMT 2002
  P 40, 231054, 170169, Mon June 10 08:32:44 GMT 2002
  U 36, 130363, 121261, Fri Jun 7 17:40:57 GMT 2002
FCP  P 19, 1417091, 1222127, Tue Jun 4 14:48:59 GMT 2002
     U 6, 139051, 121246, Fri Jun 7 17:40:42 GMT 2002
```

Unified Ethernet network management

A unified Ethernet network entails running data and storage traffic, including iSCSI, CIFS, NFS, and Fibre Channel, over your existing Ethernet infrastructure.

Unified target adapters (UTAs) are 10-Gb Ethernet adapters that you install on your storage systems, and converged network adapters (CNAs) are 10-Gb Ethernet adapters that you install on your hosts. These adapters are required for running Fibre Channel over Ethernet (FCoE) traffic, IP traffic, or both over your Ethernet network.

Note: UTAs and CNAs are configured and managed just like any other FC or Ethernet port; there are no unique configuration commands. See the *Data ONTAP File Access and Protocols Management Guide for 7-Mode* 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

[iSCSI network management](#) on page 82

[FC SAN management](#) on page 127

Related information

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

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

[Interoperability Matrix: support.netapp.com/NOW/products/interoperability](#)

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.

What data center bridging is

Data center bridging (DCB) is a collection of extensions to the existing Ethernet standard that provides a lossless transport layer for FCoE traffic.

FC provides a reliable, dedicated fabric with no packet loss. Ethernet, however, is inherently lossy, which poses problems for transmitting FC traffic. The DCB standard resolves this problem 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.

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

SAN Configuration Guide For 7-Mode: support.netapp.com

Displaying DCB settings

When you install one or more CNAs or UTAs, you can display the DCB settings associated with the adapters.

About this task

Note that these settings are configured at the switch level, and the storage system simply discovers and displays those pre-configured settings.

Choices

- Enter the following command to include the bandwidth allocation:

dcb show *interface*

- Enter the following command to display whether flow control is enabled for each priority:

dcb priority show *interface*

```
system1> dcb show e2b
Interface PGID Priority Applications Bandwidth
-----
e2b
      0      0 unassigned 10%
      1      1 2 4 5 6 7 unassigned 0%
      2      3 FCoE 90%
```

```
system1>dcb priority show e2b
Interface Priority Applications Flow Control PGID
-----
e2b
      0 unassigned enabled 0
      1 unassigned disabled 1
      2 unassigned disabled 1
      3 FCoE enabled 2
      4 unassigned disabled 1
      5 unassigned disabled 1
      6 unassigned disabled 1
      7 unassigned disabled 1
```

Priority

The relative priorities for frames that have similar traffic handling requirements, such as latency and frame loss. The available priorities, from lowest to highest priority, are 0 to 7. The default priorities are 3 for FCoE traffic and 0 for IP traffic.

Priority group

A collection of priorities bound together for the purpose of bandwidth allocation. A priority group can be associated with multiple priorities.

Priority group ID (PGID)	A numerical ID from 0 to 15 that identifies each priority group.
Bandwidth	The percentage of available bandwidth allocated to each priority group.
Applications	Activities for which bandwidth and priorities are assigned, such as FCoE and IP traffic.
Flow control	The flow control setting (enabled or disabled) for each priority. If priority-based flow control is enabled, then traffic at that priority might be paused to prevent frame loss due to congestion. Enabling priority-based flow control for one priority has no impact on traffic for a different priority.

Disk space management

Data ONTAP provides a number of tools for effectively managing disk space.

You should understand how to perform the following tasks:

- Monitor available disk space
- Configure Data ONTAP to automatically grow a FlexVol volume
- Configure Data ONTAP to automatically delete Snapshot copies when a FlexVol volume begins to run out of free space

Note: For detailed information about disk space management, see the *Data ONTAP Storage Management Guide for 7-Mode*.

Related information

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

Commands to display space information

Seeing information about how space is being used in your aggregates and volumes and their Snapshot copies enables you to manage your storage more effectively.

Use this Data ONTAP command...	To display information about...
<code>aggr status -S</code>	Disk space usage for aggregates
<code>vol status -F</code>	Disk space usage by volumes within an aggregate
<code>vol status -S</code>	Disk space usage for volumes
<code>df</code>	Disk space usage for volumes or aggregates
<code>snap delta</code>	The estimated rate of change of data between Snapshot copies in a volume
<code>snap reclaimable</code>	The estimated amount of space freed if you delete the specified Snapshot copies

For more information about the snap commands, see the *Data ONTAP Data Protection Online Backup and Recovery Guide for 7-Mode*. For more information about the `df` and `aggr status -S` commands, see the appropriate man page.

Examples of disk space monitoring using the df command

You can use the `df` command to monitor disk space on a volume in which you created LUNs.

Note: These examples are written with the assumption that the storage system and host machine are already properly configured.

Monitoring disk space on volumes with LUNs that do not use Snapshot copies

This example illustrates how to monitor disk space on a volume when you create a LUN without using Snapshot copies.

About this task

For this example, assume that you require less than the minimum capacity based on the recommendation of creating a seven-disk volume.

For simplicity, assume the LUN requires only three GB of disk space. For a traditional volume, the volume size must be approximately three GB plus 10 percent. The recommended volume size is approximately 2*3 GB plus the rate of change of data.

Steps

1. From the storage system, create a new traditional volume named `volSPACE` that has approximately 67 GB, and observe the effect on disk space by entering the following commands:

```
vol create volSPACE aggr1 67g
```

```
df -r /vol/volSPACE
```

The following sample output is displayed. There is a snap reserve of 20 percent on the volume, even though the volume is used for LUNs, because snap reserve is set to 20 percent by default.

Filesystem	kbytes	used	avail	reserved	Mounted on
/vol/volSPACE	50119928	1440	50118488	0	/vol/
volSPACE/					
/vol/volSPACE/.snapshot	12529980	0	12529980	0	/vol/
volSPACE/.snapshot					

2. Set the percentage of snap reserve space to 0 and observe the effect on disk space by entering the following commands:

```
snap reserve volSPACE 0
```

```
df -r /vol/volSPACE
```

The following sample output is displayed. The amount of available Snapshot copy space becomes zero, and the 20 percent of Snapshot copy space is added to available space for `/vol/volSPACE`.

Filesystem	kbytes	used	avail	reserved	Mounted on
/vol/volspace/	62649908	1440	62648468	0	/vol/volspace/
/vol/volspace/.snapshot	0	0	0	0	/vol/
volspace/.snapshot					

3. Create a LUN named /vol/volspace/lun0 and observe the effect on disk space by entering the following commands:

```
lun create -s 3g -t aix /vol/volspace/lun0
df -r /vol/volspace
```

The following sample output is displayed. Three GB of space is used because this is the amount of space specified for the LUN, and LUN space reservation is enabled by default.

Filesystem	kbytes	used	avail	reserved	Mounted on
/vol/volspace/	62649908	3150268	59499640	0	/vol/
volspace/					
/vol/volspace/.snapshot	0	0	0	0	/vol/
volspace/.snapshot					

4. Create an igroup named aix_host and map the LUN to it by entering the following commands (assuming that the host node name is iqn.1996-04.aixhost.host1). Depending on your host, you might need to create WWNN persistent bindings. These commands have no effect on disk space.

```
igroup create -i -t aix aix_host iqn.1996-04.aixhost.host1
lun map /vol/volspace/lun0 aix_host 0
```

5. 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.
6. From the storage system, ensure that creating the file system on the LUN and writing data to it has no effect on space on the storage system by entering the following command:

```
df -r /vol/volspace
```

The following sample output is displayed. From the storage system, the amount of space used by the LUN remains 3 GB.

Filesystem	kbytes	used	avail	reserved	Mounted on
/vol/volspace/	62649908	3150268	59499640	0	/vol/
volspace/					
/vol/volspace/.snapshot	0	0	0	0	/vol/
volspace/.snapshot					

7. Turn off space reservations and see the effect on space by entering the following commands:

```
lun set reservation /vol/volspace/lun0 disable
df -r /vol/volspace
```

The following sample output is displayed. The 3 GB of space for the LUN is no longer reserved, so it is not counted as used space; it is now available space. Any other requests to write data to the volume can occupy all of the available space, including the 3 GB that the LUN expects to

have. If the available space is used before the LUN is written to, write operations to the LUN fail. To restore the reserved space for the LUN, turn space reservations on.

Filesystem	kbytes	used	avail	reserved	Mounted on
/vol/volspace/	62649908	144	62649584	0	/vol/
volspace/					
/vol/volspace/.snapshot	0	0	0	0	/vol/
volspace/.snapshot					

Monitoring disk space on volumes with LUNs that use Snapshot copies

This example illustrates how to monitor disk space on a volume when taking Snapshot copies.

About this task

In this example, you start with a new volume, the LUN requires 3 GB of disk space, and fractional overwrite reserve is set to 100 percent.

Steps

1. From the storage system, create a new FlexVol volume named volspace that has approximately 67 GB, and observe the effect on disk space by entering the following commands:

```
vol create volspace aggr1 67g
```

```
df -r /vol/volspace
```

The following sample output is displayed. There is a snap reserve of 20 percent on the volume, even though the volume will be used for LUNs, because snap reserve is set to 20 percent by default.

Filesystem	kbytes	used	avail	reserved	Mounted on
/vol/volspace	50119928	1440	50118488	0	/vol/volspace/
/vol/volspace/.snapshot	12529980	0	12529980	0	/vol/
volspace/.snapshot					

2. Set the percentage of snap reserve space to zero by entering the following command:

```
snap reserve volspace 0
```

3. Create a LUN (/vol/volspace/lun0) by entering the following commands:

```
lun create -s 6g -t aix /vol/volspace/lun0
```

```
df -r /vol/volspace
```

The following sample output is displayed. Approximately 6 GB of space is taken from available space and is displayed as used space for the LUN:

Filesystem	kbytes	used	avail	reserved	Mounted on
/vol/volspace/	62649908	6300536	56169372	0	/vol/
volspace/					
/vol/volspace/.snapshot	0	0	0	0	/vol/
volspace/.snapshot					

4. Create an igroup named `aix_host` and map the LUN to it by entering the following commands (assuming that the host node name is `iqn.1996-04.aixhost.host1`). Depending on your host, you might need to create WWNN persistent bindings. These commands have no effect on disk space.

```
igroup create -i -t aix aix_host iqn.1996-04.aixhost.host1  
lun map /vol/volSPACE/lun0 aix_host 0
```

5. 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. These commands have no effect on disk space.
6. From the host, write data to the file system (the LUN on the storage system). This has no effect on disk space.
7. Ensure that the active file system is in a quiesced or synchronized state.
8. Take a Snapshot copy of the active file system named `snap1`, write 1 GB of data to it, and observe the effect on disk space by entering the following commands:

```
snap create volSPACE snap1  
df -r /vol/volSPACE
```

The following sample output is displayed. The first Snapshot copy reserves enough space to overwrite every block of data in the active file system, so you see 12 GB of used space, the 6-GB LUN (which has 1 GB of data written to it), and one Snapshot copy. Notice that 6 GB appears in the reserved column to ensure write operations to the LUN do not fail. If you disable space reservation, this space is returned to available space.

Filesystem	kbytes	used	avail	reserved	
Mounted on					
/vol/volSPACE/	62649908	12601072	49808836	6300536	/vol/
volSPACE/					
/vol/volSPACE/.snapshot	0	180	0	0	/vol/
volSPACE/.snapshot					

9. From the host, write another 1 GB of data to the LUN. Then, from the storage system, observe the effect on disk space by entering the following commands:

```
df -r /vol/volSPACE
```

The following sample output is displayed. The amount of data stored in the active file system does not change. You just overwrote 1 GB of old data with 1 GB of new data. However, the Snapshot copy requires the old data to be retained. Before the write operation, there was only 1 GB of data, and after the write operation, there was 1 GB of new data and 1 GB of data in a Snapshot copy. Notice that the used space increases for the Snapshot copy by 1 GB, and the available space for the volume decreases by 1 GB.

Filesystem	kbytes	used	avail	reserved	Mounted
on					
/vol/volSPACE/	62649908	12601072	47758748	0	/vol/
volSPACE/					
/vol/volSPACE/.snapshot	0	1050088	0	0	/vol/
volSPACE/.snapshot					

10. Ensure that the active file system is in a quiesced or synchronized state.
11. Take a Snapshot copy of the active file system named snap2 and observe the effect on disk space by entering the following command:

dr -r /vol/volspace

The following sample output is displayed. Because the first Snapshot copy reserved enough space to overwrite every block, only 44 blocks are used to account for the second Snapshot copy.

Filesystem Mounted on	kbytes	used	avail	reserved	
/vol/volspace/	62649908	12601072	47758748	6300536	/vol/
volspace/					
/vol/volspace/.snapshot	0	1050136	0	0	/vol/
volspace/.snapshot					

12. From the host, write 2 GB of data to the LUN and observe the effect on disk space by entering the following command:

df -r /vol/volspace

The following sample output is displayed. The second write operation requires the amount of space actually used if it overwrites data in a Snapshot copy.

Filesystem Mounted on	kbytes	used	avail	reserved	
/vol/volspace/	62649908	12601072	4608427	6300536	/vol/
volspace/					
/vol/volspace/.snapshot	0	3150371	0	0	/vol/
volspace/					
.snapshot					

Moving your volumes nondisruptively

NetApp DataMotion for Volumes enables you to nondisruptively move a volume from one aggregate to another within the same controller for capacity utilization, improved performance, and to satisfy service-level agreements. In a SAN environment, FlexVol volumes and the LUNs in the volumes are moved nondisruptively from one aggregate to another.

In a volume move, SCSI applications accessing different LUNs in the volume can continue to run during the move. Applications that use FC and iSCSI to access a LUN in the volume that is being moved do not see any I/O disruptions during the volume move. You can continue to access data in the volume during and after the volume move.

The volume move occurs in three phases: setup phase, data copy phase, and cutover phase.

Ways to use volume move

You can perform a nondisruptive volume move in different scenarios, such as moving it from a busy aggregate to a less busy aggregate or from a high-speed disk to a lower-speed disk.

You can move the volume in the following scenarios:

- From a high-speed disk to a lower-speed disk or from a lower-speed disk to a high-speed disk, to satisfy SLA requirements.
- From a full aggregate to an aggregate that has space for growth.
- From an aggregate laid out on third-party disks to an aggregate laid out on NetApp disks by using V-Series systems.
- Between different RAID types, such as RAID-DP and RAID4.
- Between different types of disk drives, such as array LUNs, SSDs, FC, SATA, or SAS.

Requirements for performing a volume move

Before you move a volume nondisruptively, you must be aware of the types of volumes you can move and the operations that might conflict with the volume move. The volume move does not start if the volume has unsupported settings or if there are conflicting operations.

- Your FAS system or V-Series system must be running Data ONTAP 8.0.1 7-Mode or later.
- You can move only one 7-Mode FlexVol volume at a time.
- The volume must be online.
- The source volume must be consistent.
- You cannot move the following types of volumes:
 - A root volume
 - A FlexClone volume
 - A FlexCache volume
 - A volume that is the destination of any replication relationship, such as volume SnapMirror or qtree SnapMirror
 - A volume that is a SnapVault destination

Note: During a volume move, you must not initiate qtree SnapMirror or SnapVault relationships from the destination volume.

- A read-only volume
- A volume in a nondefault vFiler unit
- A volume from a 64-bit aggregate to a 32-bit aggregate
- The source volume should not be exported to NFS or CIFS clients when the volume move operation is in progress.

There is a small window of time when you can export the source volume over NFS or CIFS before the volume move enters the cutover phase. However, if you do so, the cutover phase might not be successfully completed. If the cutover phase is not completed, there is no disruption to SCSI clients because the volume move rolls back to continue with the data copy phase.

- The volume guarantee option must not be set to `file`.
- Deduplication operations must not be running on the source volume.

If deduplication is active, the volume move is paused and the cutover phase is not initiated.

For more information about deduplication operations, see the *Data ONTAP Storage Management Guide for 7-Mode*.

- The following conflicting operations must not be running:

- SnapRestore of the source volume or the containing aggregate
- WAFLIron operation on the source or the destination aggregate
- Active LUN clone split operations on the source volume
- Revert operation on the storage system

Note: FlexClone volumes in the source volume are not moved along with the source volume. Fingerprint databases and change logs in the source volume are moved along with the source volume.

Related concepts

[How the setup phase of volume move works](#) on page 172

[How the data copy phase of volume move works](#) on page 173

[How the cutover phase of volume move works](#) on page 173

Related information

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

[NetApp Interoperability Matrix-support.netapp.com/NOW/products/interoperability](#)

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

How the setup phase of volume move works

The setup phase creates a temporary destination volume in the destination aggregate and initiates data transfer from the source volume to the destination volume.

During the setup phase, the system checks whether the volume you plan to move meets the specified requirements. If any of these checks fail, then the volume move is terminated and an error message is displayed. You should follow the guidance of the error message before you manually resume the volume move.

Related concepts

[Requirements for performing a volume move](#) on page 171

[How the data copy phase of volume move works](#) on page 173

[How the cutover phase of volume move works](#) on page 173

Related tasks

[Resuming the volume move operation](#) on page 176

How the data copy phase of volume move works

The data copy phase follows the setup phase of a volume move operation. In the data copy phase, incremental data is transferred automatically from the source volume to the destination volume, after which the cutover phase can begin.

After each block of data is transferred, the volume move determines whether the cutover phase can be initiated.

If a SnapRestore or a WAFLIron operation is started on the source volume, the destination volume, or the containing aggregate, the volume move is canceled and an appropriate error message is recorded in the log file.

Note: During the data copy phase, if you attempt SnapMirror migrate on the source volume, then the volume move pauses, and you cannot resume or abort the volume move operation.

If the volume move finds any unsupported settings or conflicting operations before entering the cutover phase, the volume move operation is paused and the reason for the pause is displayed. You must resolve the issue before you can manually resume the volume move.

Related concepts

[Requirements for performing a volume move](#) on page 171

[How the setup phase of volume move works](#) on page 172

[How the cutover phase of volume move works](#) on page 173

Related tasks

[Resuming the volume move operation](#) on page 176

How the cutover phase of volume move works

The cutover phase is the final phase of the volume move. During the cutover phase, the data in the source volume and the destination volume is synchronized. I/O operations are redirected to the destination volume and the volume move is complete.

Note: The host application might encounter I/O disruptions if storage system reboot, nondisruptive upgrade (NDU), shutdown, takeover, or giveback occurs during the volume move.

If the volume move is not completed within the specified cutover period (default 60 seconds), then the cutover phase is timed out, logging the appropriate error messages, and the volume move reverts to the data copy phase.

If the cutover phase is successful, it results in the following:

- The contents of the destination volume are identical to the source volume.
- The destination volume takes the identity of the source volume.
- After the volume is moved, the LUN at the destination starts processing I/O operations.
- The source volume is destroyed, unless you choose to retain it.

Depending on the number of cutover attempts, the volume move tries to enter the cutover phase again. If cutover is not completed within the specified number of cutover attempts, then the volume move is paused and an appropriate error message is recorded in the log file. You can then manually resume the volume move.

Related concepts

[Requirements for performing a volume move](#) on page 171

[How the setup phase of volume move works](#) on page 172

[How the data copy phase of volume move works](#) on page 173

[How the setup phase of volume move works](#) on page 172

[How the data copy phase of volume move works](#) on page 173

Related tasks

[Performing the volume move operation](#) on page 174

[Resuming the volume move operation](#) on page 176

Performing the volume move operation

You can nondisruptively move a volume from one aggregate to another within a storage system. You can continue to access data in the LUNs during the volume move.

Before you begin

Before the volume move enters the cutover phase, you must ensure that any existing synchronous SnapMirror relationships established on the source volume are destroyed. You can resynchronize the SnapMirror relationships after the volume move is completed.

About this task

- A temporary volume is created at the beginning of the volume move. You should not change the contents, state, or attributes of the destination volume, or create any replication, disaster recovery, SnapVault, or qtree SnapMirror relationship with other volumes for the duration of the move.
- MetroCluster relationships are not affected by the volume move.
- If your volume guarantee is set to none, the fractional reserve of the volume is automatically set to 0 after the move is completed.

Step

1. Start the volume move by entering the following command:
vol move start srcvol dstaggr [-k] [-m | -r num_cutover_attempts] [-w cutover_window] [-o] [-d]

srcvol specifies the source volume.

dstaggr specifies the destination aggregate.

-k retains the source volume after a successful move. The source volume remains offline.

-m specifies that the volume move does not initiate automatic cutover. The system continuously runs updates and you can initiate manual cutover at any point during the volume move.

num_cutover_attempts specifies the number of cutover attempts. The minimum number of cutover attempts is one and the default number of attempts is three. If cutover cannot be completed in the specified number of attempts, then the volume move is paused.

cutover_window specifies the duration of the cutover window. The default and minimum value is 60 seconds.

-o displays warning messages on the console and the operation continues.

-d runs all the data copy phase checks. If any of the checks fail, error messages are displayed on the console and the operation is terminated.

Result

If the volume move is successful, the destination volume retains the following:

- Snapshot copies of the source volume
- Attributes of the LUNs from the source volume in the corresponding LUNs in the destination volume

Related concepts

[How the setup phase of volume move works](#) on page 172

[How the data copy phase of volume move works](#) on page 173

Pausing the volume move operation

You can manually pause the volume move during the setup phase or the data copy phase to complete any high priority I/O operations.

Step

1. Pause the volume move by entering the following command:

vol move pause srcvol

Example

```
system1> vol move pause vol1
Wed Aug 29 08:11:40 GMT [system1: replication.src.err:error]:
SnapMirror: source transfer from vol1 to system1:
ndm_dstvol_1188375081 : transfer failed.
Wed Aug 29 08:11:41 GMT [system1: replication.dst.err:error]:
SnapMirror: destination transfer from 127.0.0.1:vol1 to
ndm_dstvol_1188375081 : replication transfer failed to complete.
```

```
Wed Aug 29 08:11:41 GMT [system1: vol.move.paused:info]:
Move of volume vol1 to aggregate aggr1 paused : User initiated
```

Resuming the volume move operation

When the volume move is manually or automatically paused, you can resume it by running the `vol move resume` command. On resuming, the volume move runs the same set of checks that were run during the data copy phase. You can add to or change the options you specified when you started the volume move.

Step

1. Resume the volume move operation by entering the following command:

```
vol move resume srcvol [-k] [-m | -r num_cutover_attempts] [-w cutover_window] [-o]
```

Example

```
system1> vol move resume vol1 -k -r 8 -w 120
Wed Aug 29 08:15:14 GMT [system1: vol.move.resume:info]:
Move of volume vol1 to aggregate aggr1 was resumed.
system1> Wed Aug 29 08:15:14 GMT [system1:
vol.move.transferStart:info]: Baseline transfer from volume vol1
to ndm_dstvol_1188375081 started.
```

Monitoring the volume move status

You can use the `vol move status` command to display information about the volume that is moved.

About this task

Note: If you are running the `vol move status` command in a continuous loop during cutover phase, you might see a message indicating that `vol move` is complete even before actual completion. This may not be indicating actual cutover completion. To confirm, wait a few sections and run `vol move status` again.

Step

1. Obtain the status of the volume move operation by entering the following command:

```
vol move status srcvol [-v]
```

-v provides additional information about the destination volume name, amount of data transferred, the time taken for the data transfer, and the amount of data that is currently being transferred.

Example

```
system1> vol move status vol1 -v
Source           : vol1
Destination      : aggr1:ndm_dstvol_1188375081
State            : move
Cutover Attempts : 3
Cutover Time     : 60
Last Completed Transfer:
      Data Transferred = 324 KB      Time Taken = 1 s
Current Transfer Size = 0 KB
```

Performing manual cutover of the volume move operation

If the volume move is unable to complete automatic cutover in the specified number of cutover attempts, you can initiate manual cutover. You can specify the -m option when starting or resuming the volume move to initiate cutover and increase the probability of completing the volume move within the cutover period.

Before you begin

Before starting manual cutover, you should perform any prerequisites based on the failure observed through EMS in the automatic cutover.

Step

1. Manually cut over the volume move operation by entering the following command:

```
vol move cutover srcvol [-w cutover_window]
```

Canceling the volume move operation

You can cancel the volume move if you want to complete any high priority operations.

Step

1. Cancel the volume move operation by entering the following command:

```
vol move abort srcvol
```

Working with VMware VAAI features for ESX hosts

Data ONTAP 8.0.1 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. You can determine the extent to which your system is using the VAAI features by checking the statistics contained in the VAAI counters.

The VAAI feature set consists of the following:

- **Extended copy**
This feature enables the host to initiate the transfer of data between the source and destination without involving the host in the data transfer. This results in saving ESX CPU cycles and increasing the network throughput. The extended copy feature is used in scenarios such as cloning a virtual machine. When invoked by the ESX host, the extended copy 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
 If this feature cannot be invoked, the ESX host automatically uses the standard ESX 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 scenarios such as zero-filling a file.
- **VERIFY AND WRITE**
This feature bypasses certain file access concurrency limitations, which speeds up operations such as booting up a virtual machine.

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 8.0.1 or later.
- (Extended copy only) Both the LUNs and the igroups must specify VMware as the OS type.
- (Extended copy only) The source and the destination of the VMware copy operation must be hosted on the same storage system.

It does not matter whether the VMware datastores are on different LUNs or volumes within that storage system.

Note: The extended copy feature currently does not support copying data between VMware datastores that are hosted on different storage systems.

Methods for determining whether VAAI features are supported

To confirm whether the ESX operating system supports the VAAI features, you can check either the Virtual Storage Console (VSC) or the statistics produced by the VAAI counters.

- When you are at the VSC, you can look at the VAAI Capable option. If it is displayed as Enabled, then the storage system is capable of using the VAAI features.
- To view the statistics on the VAAI features, you can use the `stats show vstorage` command. When you enter this command without an option, it displays all the counters associated with the VAAI features. When you enter it with the name of a counter as an option (`stats show vstorage:counter_name`), it displays information for only that counter.

By checking the requests counter for a feature, you can determine whether the ESX host is using that feature. This counter specifies how many requests for that feature have been sent to the storage system. The counter value increases as the ESX host invokes the feature.

The following table lists the requests counters for each feature:

Feature	Counter
Extended copy	<code>xcopy_copy_reqs</code>
WRITE SAME	<code>writesame_reqs</code>
VERIFY AND WRITE	<code>vaw_reqs</code>

Statistics collected for VAAI counters

The VAAI counters supply numerous statistics that provide information such as which features the ESX host is using, how they are performing, and how much data is being operated on by the features.

Each of the following counters supplies information for a single vFiler unit.

<code>xcopy_copy_reqs</code>	The number of requests for the extended copy feature.
<code>xcopy_abort_reqs</code>	The number of requests to abort the extended copy feature commands.
<code>xcopy_status_reqs</code>	The number of requests for status information about the extended copy feature commands.
<code>xcopy_total_data</code>	The sum of the kilobytes of data that was successfully copied using extended copy. This is a measurement of data copied at the NetApp storage system rather than through the network.
<code>xcopy_invalid_parms</code>	The number of extended copy requests that had invalid parameters.
<code>xcopy_authorization_failures</code>	The number of unauthorized requests for the extended copy feature.

xcopy_authentication_failures	The number of requests for the extended copy feature that could not be authenticated.
xcopy_copy_failures	The total number of extended copy requests that failed during copy operations.
xcopy_copyErr_isDir	The number of extended copy requests that were sent to a directory instead of a file.
xcopy_copyErr_data_unrecov	The number of extended copy requests received that failed due to an unrecoverable RAID error.
xcopy_copyErr_offline	The number of extended copy requests that failed because the volume was offline.
xcopy_copyErr_staleFH	The number of extended copy requests that failed because the request referenced an invalid file handle.
xcopy_copyErr_IO	The number of extended copy requests that failed because there was no I/O available on the storage system.
xcopy_copyErr_noSpace	The number of extended copy requests that failed because of an internal I/O error.
xcopy_copyErr_diskQuota	The number of extended copy requests that failed because the disk quota on the storage system was exceeded.
xcopy_copyErr_readOnly	The number of extended copy requests that failed because the copy destination was read-only.
xcopy_copyErr_other	The number of extended copy requests that failed due to a generic copy operation failure.
xcopy_intravol_moves	The number of extended copy requests for copy operations where the copy source and the copy destination were within the same volume.
xcopy_intervol_moves	The number of extended copy requests for copy operations where the copy source and the copy destination were on different volumes.
xcopy_one2one_moves	The number of extended copy requests for copy operations where the copy source and the copy destination were within the same LUN.
xcopy_one2many_moves	The number of extended copy requests for copy operations between one copy source and multiple copy destinations.
writesame_reqs	The sum of the WRITE SAME requests.
writesame_holepunch_reqs	The number of requests for WRITE SAME operations that were used to perform hole punching (freeing of blocks).

writesame_total_data	The sum of the kilobytes of data that was successfully written using the WRITE SAME requests.
vaw_reqs	The sum of VAW requests.
vaw_miscompares	The sum of VAW requests that resulted in a miscompare (contention for resource).

Viewing statistics for the VAAI features

You can use the `stats show` command with the option `vstorage` to display the statistics that the counters collected about the VAAI features extended copy, WRITE SAME, and VERIFY AND WRITE.

Step

1. To view the statistics for the VAAI features, complete the appropriate action:

To view...	Enter...
All the statistics	The command: stats show vstorage
A specific statistic	The <code>stats show vstorage</code> command with the name of the counter that contains the statistics you want to see: stats show vstorage:counter_name

Example

The following example uses the `stats show vstorage` command to display information from all the counters for the VAAI features:

```
TESTER1*> stats show vstorage
vstorage:vfiler0:xcopy_copy_reqs:1139
vstorage:vfiler0:xcopy_abort_reqs:0
vstorage:vfiler0:xcopy_status_reqs:0
vstorage:vfiler0:xcopy_total_data:4046848
vstorage:vfiler0:xcopy_invalid_parms:0
vstorage:vfiler0:xcopy_authorization_failures:0
vstorage:vfiler0:xcopy_authentication_failures:0
vstorage:vfiler0:xcopy_copy_failures:73
vstorage:vfiler0:xcopy_copyErr_isDir:0
vstorage:vfiler0:xcopy_copyErr_data_unrecov:0
vstorage:vfiler0:xcopy_copyErr_offline:0
vstorage:vfiler0:xcopy_copyErr_staleFH:0
vstorage:vfiler0:xcopy_copyErr_IO:0
vstorage:vfiler0:xcopy_copyErr_noSpace:0
vstorage:vfiler0:xcopy_copyErr_diskQuota:0
vstorage:vfiler0:xcopy_copyErr_readOnly:0
vstorage:vfiler0:xcopy_copyErr_other:0
vstorage:vfiler0:xcopy_intravol_moves:530
vstorage:vfiler0:xcopy_intervol_moves:536
```

```
vstorage:vfiler0:xcopy_one2one_moves:0  
vstorage:vfiler0:xcopy_one2many_moves:0  
vstorage:vfiler0:writesame_reqs:0  
vstorage:vfiler0:writesame_holepunch_reqs:0  
vstorage:vfiler0:writesame_total_data:0  
vstorage:vfiler0:vaw_reqs:0  
vstorage:vfiler0:vaw_miscompares:0  
TESTER1*>
```

In the following example, the command displays only the information collected by the `xcopy_abort_reqs` counter:

```
TESTER1*> stats show vstorage:vfiler0:xcopy_abort_reqs  
vstorage:vfiler0:xcopy_abort_reqs:0  
TESTER1*>
```

Data protection with Data ONTAP

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

Snapshot copies are a standard feature of Data ONTAP. A Snapshot copy is a frozen, read-only image of the entire Data ONTAP file system, or WAFL (Write Anywhere File Layout) volume, that reflects the state of the LUN or the file system at the time the Snapshot copy is created. The other data protection methods rely on Snapshot copies or create, use, and destroy Snapshot copies, as required.

Data protection methods

This section describes the various methods for protecting your data with Data ONTAP.

Snapshot copy

Make point-in-time copies of a volume.

"volume copy" command

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

FlexClone LUNs (FlexClone license required)

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.

Backup vaults (SnapVault license required)

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

Store and retrieve data on tape.

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 *Data ONTAP Data Protection Tape Backup and Recovery Guide for 7-Mode*.

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 *Data ONTAP Data Protection Online Backup and Recovery Guide for 7-Mode*.

Related information

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

LUN clones

A LUN clone is a point-in-time, writable copy of a LUN in a Snapshot copy. Changes made to the parent LUN after the clone is created are not reflected in the Snapshot copy.

A LUN clone shares space with the LUN in the backing Snapshot copy. When you clone a LUN, and new data is written to the LUN, the LUN clone still depends on data in the backing Snapshot copy. The clone does not require additional disk space until changes are made to it.

You cannot delete the backing Snapshot copy until you split the clone from it. When you split the clone from the backing Snapshot copy, the data is copied from the Snapshot copy to the clone, thereby removing any dependence on the Snapshot copy. After the splitting operation, both the backing Snapshot copy and the clone occupy their own space.

Note: Cloning is not NVLOG protected, so if the storage system panics during a clone operation, the operation is restarted from the beginning on a reboot or takeover.

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.

Differences between FlexClone LUNs and LUN clones

Data ONTAP provides two LUN cloning capabilities—LUN clone with the support of a Snapshot copy and FlexClone LUN. However, there are a few differences between these two LUN cloning techniques.

The following table lists the key differences between the two LUN cloning techniques:

FlexClone LUN	LUN clone
To create a FlexClone LUN, you should use the <code>clone start</code> command.	To create a LUN clone, you should use the <code>lun clone create</code> command.

FlexClone LUN	LUN clone
You do not need to create a Snapshot copy manually.	You must create a Snapshot copy manually before creating a LUN clone, because a LUN clone uses a backing Snapshot copy
A temporary Snapshot copy is created during the cloning operation. The Snapshot copy is deleted immediately after the cloning operation. However, you can prevent the Snapshot copy creation by using the <code>-n</code> option of the <code>clone start</code> command.	A LUN clone is coupled with a Snapshot copy.
A FlexClone LUN is independent of Snapshot copies. Therefore, no splitting is required.	When a LUN clone is split from the backing Snapshot copy, it uses extra storage space. The amount of extra space used depends on the type of clone split.
You can clone a complete LUN or a sub-LUN. To clone a sub-LUN, you should know the block range of the parent entity and clone entity.	You can only clone a complete LUN.
FlexClone LUNs are best for situations where you need to keep the clone for a long time.	LUN clones are best when you need a clone only for a short time.
No Snapshot copy management is required.	You need to manage Snapshot copies if you keep the LUN clones for a long time.

For more information about FlexClone LUNs, see the *Data ONTAP Storage Management Guide for 7-Mode*.

Cloning LUNs

You can use LUN clones to create multiple readable and writable copies of a LUN.

Before you begin

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

About this task

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

Steps

1. Create a LUN by entering the following command:

```
lun create -s size -t lun type lun_path
```

Example

```
lun create -s 100g -t solaris /vol/vol1/lun0
```

2. Create a Snapshot copy of the volume containing the LUN to be cloned by entering the following command:

```
snap create volume_name snapshot_name
```

Example

```
snap create vol1 mysnap
```

3. Create the LUN clone by entering the following command:

```
lun clone create clone_lun_path -b parent_lun_path parent_snap
```

clone_lun_path is the path to the clone you are creating, for example, /vol/vol1/lun0clone.

parent_lun_path is the path to the original LUN.

parent_snap is the name of the Snapshot copy of the original LUN.

Example

```
lun clone create /vol/vol1/lun0_clone -b /vol/vol1/lun0 mysnap
```

Result

The LUN clone is created.

LUN clone splits

After you clone a LUN, you can split the clone from the backing Snapshot copy.

The LUN clone split technology was significantly improved to create greater space efficiency. However, note that you must wait until the LUN clone split is complete before you can take additional Snapshot copies.

Splitting the clone from the backing Snapshot copy

If you want to delete the backing Snapshot copy, you can split the LUN clone from the backing Snapshot copy without taking the LUN offline. Any data from the Snapshot copy that the LUN clone depended on is copied to the LUN clone.

About this task

You cannot delete the backing Snapshot copy or create a new Snapshot copy until the LUN clone split is complete.

Step

1. Begin the clone split operation by entering the following command:

```
lun clone split start lun_path
```

lun_path is the path to the cloned LUN.

Result

The Snapshot copy can be deleted.

Displaying the progress of a clone-splitting operation

Because clone splitting is a copy operation and might take considerable time to complete, you can check the status of a clone splitting operation that is in progress.

Step

1. Enter the following command:

```
lun clone split status lun_path
```

lun_path is the path to the cloned LUN.

Stopping the clone-splitting process

You can use the `lun clone split` command to stop a clone split that is in progress.

Step

1. Enter the following command:

```
lun clone split stop lun_path
```

lun_path is the path to the cloned LUN.

Deleting Snapshot copies

After you split the LUN clone from the backing Snapshot copy, you have removed any dependence on that Snapshot copy so it can be safely deleted.

Step

1. Delete the Snapshot copy by entering the following command:

```
snap delete vol-name snapshot-name
```

Example

```
snap delete vol2 snap2
```

Result

The Snapshot copy is deleted.

Deleting backing Snapshot copies of deleted LUN clones

Prior to Data ONTAP 7.3, the system automatically locked all backing Snapshot copies when Snapshot copies of LUN clones were taken. Starting with Data ONTAP 7.3, you can enable the system to only lock backing Snapshot copies for the active LUN clone. If you do this, when you delete the active LUN clone, you can delete the base Snapshot copy without having to first delete all of the more recent backing Snapshot copies.

About this task

This behavior is not enabled by default; you can use the `snapshot_clone_dependency` volume option to enable it. If this option is set to `off`, you might still be required to delete all subsequent Snapshot copies before deleting the base Snapshot copy.

If you enable this option, you are not required to rediscover the LUNs. If you perform a subsequent volume `snap restore` operation, the system restores whichever value was present at the time the Snapshot copy was taken.

Step

1. Enable this behavior by entering the following command:

```
vol options volume_name snapshot_clone_dependency on
```

Examples of deleting backing Snapshot copies of deleted LUN clones

You can use the `snapshot_clone_dependency` option to determine whether you can delete the base Snapshot copy without deleting the more recent Snapshot copies after deleting a LUN clone. This option is set to `off` by default.

Example with `snapshot_clone_dependency` set to `off`

The following example illustrates how all newer backing Snapshot copies must be deleted before deleting the base Snapshot copy when a LUN clone is deleted.

You can set the `snapshot_clone_dependency` option to `off` by entering the following command:

```
vol options volume_name snapshot_clone_dependency off
```

You can create a new LUN clone, `lun_s1` from the LUN in Snapshot copy `snap1`. Also, you should run the `lun show -v` command to show that `lun_s1` is backed by `snap1`.

```
system1> lun clone create /vol/vol1/lun_s1 -b /vol/vol1/lun snap1

system1>lun show -v
/vol/vol1/lun_s1 32m (33554432) (r/w, online)
Serial#: BYjB3?-iq3hU
```

```

Backed by: /vol/vol1/.snapshot/snap1/lun
Share: none
Space Reservation: enabled
Multiprotocol Type: linux
Occupied Size:      0 (0)
Creation Time: Tue Oct 19 10:49:13 GMT 2010
Cluster Shared Volume Information: 0x0

```

You should run the `snap list` command to show that `snap1` is busy, as expected.

```

system1> snap list vol1
Volume vol1
working...

```

%/used	%/total	date	name	
24% (24%)	0% (0%)	Dec 20 02:40	snap1	(busy, LUNs)

When you create a new Snapshot copy, `snap2`, it contains a copy of `lun_s1`, which is still backed by the LUN in `snap1`.

```

system1> snap create vol1 snap2
system1> snap list vol1
Volume vol1
working...

```

%/used	%/total	date	name	
24% (24%)	0% (0%)	Dec 20 02:41	snap2	
43% (31%)	0% (0%)	Dec 20 02:40	snap1	(busy, LUNs)

You should run the `lun snap usage` command to show this dependency.

```

system1> lun snap usage vol1 snap1
Active:
    LUN: /vol/vol1/lun_s1
    Backed By: /vol/vol1/.snapshot/snap1/lun
Snapshot - snap2:
    LUN: /vol/vol1/.snapshot/snap2/lun_s1
    Backed By: /vol/vol1/.snapshot/snap1/lun

```

Then you should delete the LUN clone `lun_s1`.

```

system1> lun destroy /vol/vol1/lun_s1
Wed Dec 20 02:42:23 GMT [wafl.inode.fill.disable:info]: fill
reservation disabled for inode 3087 (vol vol1).
Wed Dec 20 02:42:23 GMT [wafl.inode.overwrite.disable:info]:
overwrite reservation disabled for inode 3087 (vol vol1).
Wed Dec 20 02:42:23 GMT [lun.destroy:info]: LUN /vol/vol1/lun_s1
destroyed

```

```
system1> lun show
/vol/vol1/lun          30m (31457280)      (r/w,
online)
```

You should run the `lun snap usage` command to show that snap2 still has a dependency on snap1.

```
system1> lun snap usage vol1 snap1
Snapshot - snap2:
  LUN: /vol/vol1/.snapshot/snap2/lun_s1
  Backed By: /vol/vol1/.snapshot/snap1/lun
```

You should run the `snap list` command to show that snap1 is still busy.

```
system1> snap list vol1
Volume vol1
working...

  %/used      %/total   date           name
  -----
  39% (39%)   0% ( 0%)   Dec 20 02:41   snap2
  53% (33%)   0% ( 0%)   Dec 20 02:40   snap1           (busy, LUNs)
```

Since snap1 is still busy, you cannot delete it until you delete the more recent Snapshot copy, snap2.

Example with `snapshot_clone_dependency` set to on

The following example illustrates how you can delete a base Snapshot copy without deleting all newer backing Snapshot copies when a LUN clone is deleted.

You can set the `snapshot_clone_dependency` option to on by entering the following command:

vol options *volume_name* snapshot_clone_dependency on

You can create a new LUN clone, `lun_s1`, from the LUN in Snapshot copy snap1. You should run the `lun show -v` command to show that `lun_s1` is backed by snap1.

```
system1> lun clone create /vol/vol1/lun_s1 -b /vol/vol1/lun snap1

system1> lun show -v
/vol/vol1/lun_s1 32m (33554432) (r/w, online)
  Serial#: BYjB3?-iq3hU
  Backed by: /vol/vol1/.snapshot/snap1/lun
  Share: none
  Space Reservation: enabled
  Multiprotocol Type: linux
  Occupied Size:      0 (0)
```

```
Creation Time: Tue Oct 19 10:49:13 GMT 2010
Cluster Shared Volume Information: 0x0
```

You should run the `snap list` command to show that `snap1` is busy, as expected.

```
system1> snap list vol1
Volume vol1
working...
```

%/used	%/total	date	name
24% (24%)	0% (0%)	Dec 20 02:40	snap1 (busy, LUNs)

When you create a new Snapshot copy, `snap2`, it contains a copy of `lun_s1`, which is still backed by the LUN in `snap1`.

```
system1> snap create vol1 snap2
system1> snap list vol1
Volume vol1
working...
```

%/used	%/total	date	name
24% (24%)	0% (0%)	Dec 20 02:41	snap2
43% (31%)	0% (0%)	Dec 20 02:40	snap1 (busy, LUNs)

You should run the `lun snap usage` command to show this dependency.

```
system1> lun snap usage vol1 snap1
Active:
  LUN: /vol/vol1/lun_s1
  Backed By: /vol/vol1/.snapshot/snap1/lun
Snapshot - snap2:
  LUN: /vol/vol1/.snapshot/snap2/lun_s1
  Backed By: /vol/vol1/.snapshot/snap1/lun
```

Then you can delete the LUN clone `lun_s1`.

```
system1> lun destroy /vol/vol1/lun_s1
Wed Dec 20 02:42:23 GMT [wafl.inode.fill.disable:info]: fill
reservation disabled for inode 3087 (vol vol1).
Wed Dec 20 02:42:23 GMT [wafl.inode.overwrite.disable:info]:
overwrite reservation disabled for inode 3087 (vol vol1).
Wed Dec 20 02:42:23 GMT [lun.destroy:info]: LUN /vol/vol1/lun_s1
destroyed
```

```
system1> lun show
/vol/vol1/lun 30m (31457280) (r/w,
online)
```

You should run the `lun snap usage` command to show that snap2 still has a dependency on snap1.

```
system1> lun snap usage vol1 snap1
Snapshot - snap2:
  LUN: /vol/vol1/.snapshot/snap2/lun_s1
  Backed By: /vol/vol1/.snapshot/snap1/lun
```

You should run the `snap list` command to show that snap1 is no longer busy.

```
system1> snap list vol1
Volume vol1
working...

  %/used      %/total    date           name
  -----
  39% (39%)   0% ( 0%)   Dec 20 02:41   snap2
  53% (33%)   0% ( 0%)   Dec 20 02:40   snap1
```

Since snap1 is no longer busy, you can delete it without first deleting snap2.

```
system1> snap delete vol1 snap1
Wed Dec 20 02:42:55 GMT [wafl.snap.delete:info]: Snapshot copy
snap1 on volume vol1 was deleted by the Data ONTAP function
snapcmd_delete.
The unique ID for this Snapshot copy is (1, 6).
```

```
system1> snap list vol1
Volume vol1
working...

  %/used      %/total    date           name
  -----
  38% (38%)   0% ( 0%)   Dec 20 02:41   snap2
```

Deleting busy Snapshot copies

A Snapshot copy is in a busy state if there are any LUN clones backed by data in that Snapshot copy because the Snapshot copy contains data that is used by the LUN clone. These LUN clones can exist either in the active file system or in some other Snapshot copy.

About this task

You can use the `lun snap usage` command to list all the LUNs backed by data in the specified Snapshot copy. That command also lists the corresponding Snapshot copies in which such LUNs exist.

The `lun snap usage` command displays the following information:

- LUN clones that are holding a lock on the Snapshot copy given as input to this command
- Snapshots in which these LUN clones exist

Steps

1. Identify all Snapshot copies that are in a busy state, locked by LUNs, by entering the following command:

```
snap list vol-name
```

Example

```
snap list vol2
```

The following message is displayed:

```
Volume vol2
working...

  %/used      %/total    date           name
  -----
  0% ( 0%)    0% ( 0%)    Jan 14 04:35   snap3
  0% ( 0%)    0% ( 0%)    Jan 14 03:35   snap2
  42% (42%)   22% (22%)   Dec 12 18:38   snap1
  42% ( 0%)   22% ( 0%)   Dec 12 03:13   snap0 (busy, LUNs)
```

2. Identify the LUNs and the Snapshot copies that contain them by entering the following command:
lun snap usage [-s] vol_name snap_name

Use the `-s` option to only display the relevant backing LUNs and Snapshot copies that must be deleted.

Note: The `-s` option is particularly useful in making SnapDrive output more readable. For example:

```
lun snap usage -s vol2 snap0
You need to delete the following snapshots before deleting snapshot
"snap0":
/vol/vol1/.snapshot/snap1
/vol/vol2/.snapshot/snap2
```

Example

```
lun snap usage vol2 snap0
```

The following message is displayed:

```
active:
  LUN:      /vol/vol2/lunC
```

```

    Backed By: /vol/vol12/.snapshot/snap0/lunA
snap2:
  LUN:      /vol/vol12/.snapshot/snap2/lunB
  Backed By: /vol/vol12/.snapshot/snap0/lunA
snap1:
  LUN:      /vol/vol11/.snapshot/snap1/lunB
  Backed By: /vol/vol12/.snapshot/snap0/lunA

```

Note: The LUNs are backed by lunA in the snap0 Snapshot copy.

In some cases, the path for LUN clones backed by a Snapshot copy cannot be determined. In those instances, a message is displayed so that those Snapshot copies can be identified. You must still delete these Snapshot copies in order to free the busy backing Snapshot copy. For example:

lun snap usage vol12 snap0

```

Snapshot - snap2:
  LUN: Unable to determine the path of the LUN
  Backed By: Unable to determine the path of the LUN
  LUN:      /vol/vol12/.snapshot/snap2/lunB
  Backed By: /vol/vol12/.snapshot/snap0/lunA

```

3. Delete all the LUNs in the active file system that are displayed by the `lun snap usage` command by entering the following command:

lun destroy [-f] lun_path [lun_path ...]

Example

lun destroy /vol/vol12/lunC

4. Delete all the Snapshot copies that are displayed by the `lun snap usage` command in the order they appear, by entering the following command:

snap delete vol-name snapshot-name

Example

snap delete vol12 snap2

snap delete vol12 snap1

All the Snapshot copies containing lunB are now deleted and snap0 is no longer busy.

5. Delete the Snapshot copy by entering the following command:

snap delete vol-name snapshot-name

Example

snap delete vol12 snap0

Restoring a Snapshot copy of a LUN in a volume

You can use SnapRestore to restore a Snapshot copy of a LUN and the volume that contains it to its state when the Snapshot copy was taken. You can use SnapRestore to restore an entire volume or a single LUN.

Before you begin

Before using SnapRestore, you must perform the following tasks:

- Always unmount the LUN before you run the `snap restore` command on a volume containing the LUN or before you run a single file SnapRestore of the LUN. For a single file SnapRestore, you must also take the LUN offline.
- Check available space; SnapRestore does not revert the Snapshot copy if sufficient space is unavailable.

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 a single LUN is restored, it must be taken offline or be unmapped prior to recovery. Using SnapRestore on a LUN, or on a volume that contains LUNs, without stopping all host access to those LUNs, can cause data corruption and system errors.

Steps

1. From the host, stop all host access to the LUN.
2. From the host, if the LUN contains a host file system mounted on a host, unmount the LUN on that host.
3. From the storage system, unmap the LUN by entering the following command:

```
lun unmap lun_path initiator-group
```

4. Enter the following command:

```
snap restore [-f] [-t vol] volume_name [-s snapshot_name]
```

-f suppresses the warning message and the prompt for confirmation. This option is useful for scripts.

-t *vol volume_name* specifies the volume name to restore.

volume_name is the name of the volume to be restored. Enter the name only, not the complete path. You can enter only one volume name.

`-s snapshot_name` specifies the name of the Snapshot copy from which to restore the data. You can enter only one Snapshot copy name.

Example

snap restore -s payroll_lun_backup.2 -t vol /vol/payroll_lun

```
storage_system> WARNING! This will restore a volume from a snapshot
into the active filesystem. If the volume already exists in the
active filesystem, it will be overwritten with the contents from the
snapshot.
Are you sure you want to do this? y
You have selected file /vol/payroll_lun, snapshot payroll_lun_backup.2
Proceed with restore? y
```

If you did not use the `-f` option, Data ONTAP displays a warning message and prompts you to confirm your decision to restore the volume.

5. Press **y** to confirm that you want to restore the volume.

Data ONTAP displays the name of the volume and the name of the Snapshot copy for the reversion. If you did not use the `-f` option, Data ONTAP prompts you to decide whether to proceed with the reversion.

6. Decide if you want to continue with the reversion.
 - If you want to continue the reversion, press **y**. The storage system reverts the volume from the selected Snapshot copy.
 - If you do not want to continue the reversion, press **n** or **Ctrl-C**. The volume is not reverted and you are returned to a storage system prompt.
7. Enter the following command to unmap the existing old maps that you do not want to keep.

lun unmap lun_path initiator-group

8. Remap the LUN by entering the following command:

lun map lun_path initiator-group

9. From the host, remount the LUN if it was mounted on a host.

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

11. From the storage system, bring the restored LUN online by entering the following command:

lun online lun_path

After you finish

After you use SnapRestore to update a LUN from a Snapshot copy, you also need to restart any applications you closed down and remount the volume from the host side.

Restoring a single LUN

You can use SnapRestore to restore a single LUN without restoring the volume that contains it.

Steps

1. Notify users that you are going to restore a LUN so that they know that the current data in the LUN will be replaced by that of the selected Snapshot copy.
2. Enter the following command:

```
snap restore [-f] [-t file] [-s snapshot_name] [-r restore_as_path]  
path_and_LUN_name
```

-f suppresses the warning message and the prompt for confirmation.

-t file specifies that you are entering the name of a file to revert.

-s *snapshot_name* specifies the name of the Snapshot copy from which to restore the data.

-r *restore_as_path* restores the file to a location in the volume different from the location in the Snapshot copy. For example, if you specify /vol/vol0/vol3/mylun as the argument to -r, SnapRestore restores the file called mylun to the location /vol/vol0/vol3 instead of to the path structure indicated by the path in *path_and_lun_name*.

path_and_LUN_name is the complete path to the name of the LUN to be restored. You can enter only one path name.

A LUN can be restored only to the volume where it was originally. The directory structure to which a LUN is to be restored must be the same as specified in the path. If this directory structure no longer exists, you must re-create it before restoring the file.

Unless you enter -r and a path name, only the LUN at the end of the *path_and_lun_name* is reverted.

If you did not use the -f option, Data ONTAP displays a warning message and prompts you to confirm your decision to restore the LUN.

3. Type the following character to confirm that you want to restore the file:

y

Data ONTAP displays the name of the LUN and the name of the Snapshot copy for the restore operation. If you did not use the -f option, Data ONTAP prompts you to decide whether to proceed with the restore operation.

4. Type the following character to continue with the restore operation:

y

Data ONTAP restores the LUN from the selected Snapshot copy.

Example of a single LUN restore

```
snap restore -t file -s payroll_backup_friday /vol/vol1/payroll_luns
```

```
storage_system> WARNING! This will restore a file from a snapshot
into the active filesystem.
If the file already exists in the active filesystem, it will be
overwritten with the contents from the snapshot.
Are you sure you want to do this? y
You have selected file /vol/vol1/payroll_luns, snapshot
payroll_backup_friday
Proceed with restore? y
```

Data ONTAP restores the LUN called payroll_backup_friday to the existing volume and directory structure /vol/vol1/payroll_luns.

After a LUN is restored with SnapRestore, all data and all relevant user-visible attributes for that LUN in the active file system are identical to that contained in the Snapshot copy.

Backing up SAN systems to tape

In most cases, backup of SAN systems to tape takes place 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.


Before you begin

You must have completed the following tasks:

- Created the production LUN
- Created the igroup to which the LUN will belong
The igroup must include the WWPN of the application server.
- Mapped the LUN to the igroup
- Formatted the LUN and made it accessible to the host

About this task

Configure volumes as SAN-only or NAS-only, and 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 WAFL volume or qtree or spread across multiple WAFL 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.

To map a LUN within a Snapshot copy for backup, complete the following steps.

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 the command provided by your host operating system, or by using SnapDrive for Windows or SnapDrive for UNIX.
2. Create a Snapshot copy by entering the following command:
snap create volume_name snapshot_name

Example

snap create vol1 payroll_backup

3. To create a clone of the production LUN, enter the following command:

```
lun clone create clone_lunpath -b parent_lunpath parent_snap
```

Example

```
lun clone create /vol/vol1/qtrees_1/payroll_lun_clone -b /vol/vol1/  
qtrees_1/payroll_lun payroll_backup
```

4. Create an igroup that includes the WWPN of the backup server by entering the following command:

```
igroup create -f -t ostype group [node ...]
```

Example

```
igroup create -f -t windows_2008 backup_server 10:00:00:00:d3:6d:0f:e1
```

Data ONTAP creates an igroup that includes the WWPN (10:00:00:00:d3:6d:0f:e1) of the Windows backup server.

5. To map the LUN clone you created in Step 3 to the backup host, enter the following command:

```
lun map lun_path initiator-group LUN_ID
```

Example

```
lun map /vol/vol1/qtrees_1/payroll_lun_clone backup_server 1
```

Data ONTAP maps the LUN clone (/vol/vol1/qtrees_1/payroll_lun_clone) to the igroup called backup_server with a SCSI ID of 1.

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. Take the LUN clone offline by entering the following command:

```
lun offline /vol/vol_name/qtrees_name/lun_name
```

Example

```
lun offline /vol/vol1/qtrees_1/payroll_lun_clone
```

9. Remove the LUN clone by entering the following command:

```
lun destroy lun_path
```

Example

```
lun destroy /vol/vol1/qtrees_1/payroll_lun_clone
```

10. Remove the Snapshot copy by entering the following command:

```
snap delete volume_name lun_name
```

Example

```
snap delete vol1 payroll_backup
```

Using volume copy to copy LUNs

You can use the `vol copy` command to copy LUNs; however, this requires that applications accessing the LUNs are quiesced and offline prior to the copy operation.

Before you begin

The contents of the host file system buffers must be saved to disk before running `vol copy` commands on the storage system.

Note: The term *LUNs* in this context refers to the LUNs that Data ONTAP serves to clients, not to the array LUNs used for storage on a storage array.

About this task

The `vol copy` command enables you to copy data from one WAFL volume to another, either within the same storage system or to a different storage system. The result of the `vol copy` command is a restricted volume containing the same data that was on the source storage system at the time you initiate the copy operation.

Step

1. To copy a volume containing a LUN to the same or different storage system, enter the following command:

```
vol copy start -S source:source_volume dest:dest_volume
```

-S copies all Snapshot copies in the source volume to the destination volume. If the source volume has Snapshot copy-backed LUNs, you must use the -S option to ensure that the Snapshot copies are copied to the destination volume.

If the copying takes place between two storage systems, you can enter the `vol 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.

Example

```
vol copy start -S systemA:vol0 systemB:vol1
```

Basic block access concepts

In iSCSI networks and FC fabrics, storage systems are targets that have storage target devices, which are referred to as LUNs, or 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

What ALUA is

Data ONTAP 7.2 added support for the Asymmetric Logical Unit Access (ALUA) features of SCSI, also known as SCSI Target Port Groups or Target Port Group Support.

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.

Note: You cannot enable ALUA on iSCSI igroups.

Attention: You must ensure that your host supports ALUA before enabling it. Enabling ALUA for a host that does not support it can cause host failures during cluster failover.

Related tasks

[Enabling ALUA](#) on page 68

Related information

[Interoperability Matrix: support.netapp.com/NOW/products/interoperability](http://support.netapp.com/NOW/products/interoperability)

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](http://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.

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 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 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

Storage system node name

Each storage system has a default node name based on a reverse domain name and the serial number of the storage system's non-volatile RAM (NVRAM) card.

The node name is displayed in the following format:

`iqn.1992-08.com.netapp:sn.serial-number`

The following example shows the default node name for a storage system with the serial number 12345678:

`iqn.1992-08.com.netapp:sn.12345678`

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`

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.

What target portal groups are

A target portal group is a set of network portals within an iSCSI node over which an iSCSI session is conducted.

In a target, a network portal is identified by its IP address and listening TCP port. For storage systems, each network interface can have one or more IP addresses and therefore one or more network portals. A network interface can be an Ethernet port, virtual local area network (VLAN), or interface group.

The assignment of target portals to portal groups is important for two reasons:

- The iSCSI protocol allows only one session between a specific iSCSI initiator port and a single portal group on the target.
- All connections within an iSCSI session must use target portals that belong to the same portal group.

By default, Data ONTAP maps each Ethernet interface on the storage system to its own default portal group. You can create new portal groups that contain multiple interfaces.

You can have only one session between an initiator and target using a given portal group. To support some multipath I/O (MPIO) solutions, you need to have separate portal groups for each path. Other initiators, including the Microsoft iSCSI initiator version 2.0, support MPIO to a single target portal group by using different initiator session IDs (ISIDs) with a single initiator node name.

Note: Although this configuration is supported, it is not recommended for NetApp storage systems. For more information, see the Technical Report TR-3441 on *iSCSI Multipathing*.

Related information

[iSCSI Multipathing Possibilities on Windows with Data ONTAP-media.netapp.com/documents/tr-3441.pdf](http://media.netapp.com/documents/tr-3441.pdf)

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, and it is configured and enabled for use by both the initiator and the storage system, the storage system automatically registers its IP address, node name, and portal groups with the iSNS server when the iSNS service is started. The iSCSI initiator can query the iSNS server to discover the storage system as a target device.

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

Currently available iSNS servers support different versions of the iSNS specification. Depending on which iSNS server you are using, you may have to set a configuration parameter in the storage system.

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.

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 system appears as a single iSCSI target node with one iSCSI node name. For storage systems with a MultiStore license enabled, each vFiler unit is a target with a different iSCSI node name.

On the storage system, the interface can be an Ethernet port, interface group, 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 storage system supports up to 1,024 simultaneous sessions, depending on its memory capacity. 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.

You can change the assignment of target portals to portal groups as needed to support multi-connection sessions, multiple sessions, and multipath I/O.

Each session has an Initiator Session ID (ISID), a number that is determined by the initiator.

How iSCSI works with HA pairs

HA pairs provide high availability because one system in the HA pair can take over if its partner fails. During failover, the working system assumes the IP addresses of the failed partner and can continue to support iSCSI LUNs.

The two systems in the HA pair should have identical networking hardware with equivalent network configurations. The target portal group tags associated with each networking interface must be the same on both systems in the configuration. This ensures that the hosts see the same IP addresses and target portal group tags whether connected to the original storage system or connected to the partner during failover.

Setting up the iSCSI protocol on a host and storage system

The procedure for setting up the iSCSI protocol on a host and storage system follows the same basic sequence for all host types.

About this task

You must alternate between setting up the host and the storage system in the order shown below.

Steps

1. Install the initiator HBA and driver or software initiator on the host and record or change the host's iSCSI node name.

It is recommended that you use the host name as part of the initiator node name to make it easier to associate the node name with the host.

2. Configure the storage system, including the following:

- Licensing and starting the iSCSI service
- Optionally configuring CHAP
- Creating LUNs, creating an igroup that contains the host's iSCSI node name, and mapping the LUNs to that igroup

Note: If you are using SnapDrive, do not manually configure LUNs. You must configure them using SnapDrive after it is installed.

3. Configure the initiator on the host, including the following:

- Setting initiator parameters, including the IP address of the target on the storage system
- Optionally configuring CHAP
- Starting the iSCSI service

4. Access the LUNs from the host, including the following:

- Creating file systems on the LUNs and mounting them, or configuring the LUNs as raw devices
- Creating persistent mappings of LUNs to file systems

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

[FC SAN management](#) on page 127

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

[FC SAN management](#) on page 127

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.

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 optical 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) and a worldwide port name (WWPN).

How WWPNs are used

WWPNs identify each port on an adapter. They are used for creating an initiator group and for uniquely identifying a storage system's HBA target ports.

- **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 a storage system's HBA target ports**
The storage system's WWPNs uniquely identify each target port on the system. The host operating system uses the combination of the WWNN and WWPN to identify storage system adapters and host target IDs. Some operating systems require persistent binding to ensure that the LUN appears at the same target ID on the host.

Related concepts

[*Required information for mapping a LUN to an igroup*](#) on page 75

[*How to make LUNs available on specific FC target ports*](#) on page 77

How storage systems are identified

When the FC protocol service is first initialized, it assigns a WWNN to a storage system based on the serial number of its NVRAM adapter. The WWNN is stored on disk.

Each target port on the HBAs installed in the storage system has a unique WWPN. Both the WWNN and the WWPN are a 64-bit address represented in the following format:
nn:nn:nn:nn:nn:nn:nn:nn, where n represents a hexadecimal value.

You can use commands such as `fcpl show adapter`, `fcpl config`, `sysconfig -v`, or `fcpl nodename` to see the system's WWNN as `FC Nodename` or `nodename`, or the system's WWPN as `FC portname` or `portname`.

How hosts are identified


You can use the `fcpl show initiator` command to see all of the WWPNs, and any associated aliases, of the FC initiators that have logged on to the storage system. Data ONTAP displays the WWPN as `Portname`.

To know which WWPNs are associated with a specific host, see the FC Host Utilities documentation for your host. These documents describe commands supplied by the Host Utilities or the vendor of the initiator, or methods that show the mapping between the host and its WWPN. For example, for Windows hosts, you should use the `lputilnt`, `HBAnywhere`, or `SANsurfer` applications, and for UNIX hosts, you should use the `sanlun` command.

How FC switches are identified

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

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



- Port 0, WWPN 20:00:00:60:69:51:06:b4
- Port 1, WWPN 20:01:00:60:69:51:06:b4
- Port 14, WWPN 20:0e:00:60:69:51:06:b4
- Port 15, WWPN 20:0f:00:60:69:51:06:b4

Copyright information

Copyright © 1994–2013 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, Cryptainer, CryptoShred, CyberSnap, Data Center Fitness, Data ONTAP, DataFabric, DataFort, Decru, Decru DataFort, DenseStak, Engenio, Engenio logo, E-Stack, ExpressPod, FAServer, FastStak, FilerView, 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, Mars, Manage ONTAP, 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

10-Gb [161](#)

10-Gb Ethernet adapters [161](#)

A

access lists

about [90](#)

creating [90](#)

displaying [91](#)

removing interfaces from [91](#)

adapters

changing the speed for [138](#)

changing the WWPN for [141](#)

configuring for initiator mode [150](#)

configuring for target mode [148](#)

displaying brief target adapter information [155](#)

displaying detailed target adapter information [156](#)

displaying information about all [154](#)

displaying information for FCP [152](#)

displaying statistics for target adapters [158](#)

aggregates

creating [28](#)

defined [21](#)

aliases

for WWPNs [143](#)

alignment

LUN [61](#)

ALUA

defined [204](#)

enabling [68](#)

igroup [68](#)

authentication

defining default for CHAP [98](#)

iSCSI [96](#)

autodelete

setting options for [39](#)

setting volume options for [42](#)

use conditions [33](#)

volume size [24](#)

autogrow

how Data ONTAP can add space for FlexVol

volumes automatically [33](#)

autosizing

how Data ONTAP can add space for FlexVol

volumes automatically [33](#)

B

backing up SAN systems [199](#)

best practices

storage provisioning [22](#)

Block access [203](#)

C

capacity [11](#), [27](#)

CHAP

and RADIUS [102](#)

authenticate

iSCSI initiator [99](#)

defined [208](#)

defining default authentication [98](#)

guidelines [97](#)

iSCSI authentication [96](#)

using with vFile units [96](#)

cluster failover

avoiding igroup mapping conflicts with [128](#)

multipathing requirements for [130](#)

overriding mapping conflicts [129](#)

understanding [127](#)

configuration options

volumes [9](#), [30](#)

configure volumes

autodelete [42](#)

configuring

thin provisioning [13](#)

configuring LUN

autodelete [33](#)

configuring volumes

autodelete [33](#)

create_ucode option

changing with the command line [41](#)

creating

aggregates [28](#)

cutover phase

cutover attempts [173](#)

volume move [173](#)

D

data center bridging

defined [162](#)

- data copy phase
 - volume move [173](#)
 - Data ONTAP
 - caution when upgrading [106](#)
 - Data ONTAP options
 - iscsi.isns.rev [93](#)
 - iscsi.max_connections_per_session [82](#)
 - iscsi.max_error_recovery_level [83](#)
 - data protection [183](#)
 - DataMotion for Volumes
 - about [170](#)
 - DCB
 - defined [162](#)
 - DCB (data center bridging) switch
 - for FCoE [161](#)
 - DCB settings [163](#)
 - df command
 - monitoring disk space using [166](#)
 - disk space
 - monitoring with Snapshot copies [168](#)
 - monitoring without Snapshot copies [166](#)
 - disk space management [165](#)
 - displaying
 - space information, commands for [165](#)
- ## E
- enabling
 - ALUA [68](#)
 - report_scsi_name [69](#)
 - error recovery level
 - enabling levels 1 and 2 [83](#)
 - Ethernet [88](#), [161](#), [203](#), [205](#)
 - eui type designator [207](#)
 - example
 - thin provisioning [11](#)
 - extended copy feature
 - environment [178](#)
 - invoked automatically [178](#)
 - statistics collected [179](#)
 - VAAI feature [177](#)
 - viewing statistics [181](#)
 - when the standard copy operation is used [177](#)
- ## F
- FC
 - changing the adapter speed [138](#)
 - checking interfaces [56](#)
 - displaying adapters [152](#)
 - managing in HA pairs [127](#)
 - managing systems with onboard adapters [147](#)
 - storage system nodes [211](#)
 - FC license
 - disabling [136](#)
 - enabling [136](#)
 - FC service
 - displaying statistics for [160](#)
 - starting and stopping [137](#)
 - FCoE
 - data center bridging [162](#)
 - target adapters [161](#)
 - FCP
 - changing the WWNN [143](#)
 - defined [210](#)
 - host nodes [212](#)
 - node connection [211](#)
 - node identification [211](#)
 - nodes defined [210](#)
 - switch nodes [212](#)
 - taking adapters offline and online [138](#)
 - FCP commands
 - fc config [138](#), [152](#)
 - fc nodename [152](#)
 - fc portname set [141](#)
 - fc show [152](#)
 - fc start [137](#)
 - fc stats [152](#)
 - fc status [135](#)
 - fc stop [137](#)
 - license [135](#)
 - storage show adapter [152](#)
 - fc ping
 - connectivity [147](#)
 - fabric latency [147](#)
 - FCP service
 - displaying how long running [160](#)
 - displaying partner's traffic information [159](#)
 - displaying traffic information about [159](#)
 - verifying the service is licensed [135](#)
 - verifying the service is running [135](#)
 - FCP target service
 - enabling [136](#)
 - Fibre Channel over Ethernet (FCoE)
 - overview [161](#)
 - FlexClone files and FlexClone LUNs
 - differences between FlexClone LUNs and LUN clones [185](#)
 - FlexClone LUNs
 - reasons for using [185](#)

- flexible volumes
 - described [21](#)
- FlexVol volumes
 - fractional reserve
 - considerations for setting [31](#)
 - how Data ONTAP can automatically add space for [33](#)
- fractional reserve
 - considerations for setting [31](#)
- free space
 - how Data ONTAP can increase automatically for FlexVol volumes [33](#)

G

- going offline
 - LUNs [12](#)
- guidelines
 - CHAP authentication [97](#)
 - LUN layout [51](#)
 - LUN mapping [76](#)
 - LUN type [49](#)
 - provisioning [22](#)
 - space allocation [51](#)
 - thin provisioning [13](#)

H

- HA pairs
 - and controller failover [127](#)
 - and iSCSI [209](#)
 - using with iSCSI [122](#)
- HBA
 - displaying information about [157](#)
- head swap
 - changing WWPNs [141](#)
- host
 - initiator
 - node name [66](#)
 - iSCSI implementation [205](#)
 - storage system connection [203](#)
- host bus adapters
 - displaying information about [157](#)
- Host Utilities
 - defined [203](#)

I

- I/O
 - alignment issues [61](#)

- igroup
 - WWPN [211](#)
- igroup commands
 - for vFiler units [67](#)
 - igroup add [79](#)
 - igroup create [45](#)
 - igroup destroy [78](#)
 - igroup remove [79](#)
 - igroup rename [80](#)
 - igroup set [81](#)
 - igroup set alua [68](#)
 - igroup show [80](#)
- igroup commands for iSCSI
 - igroup create [64](#)
- igroup mapping conflicts
 - avoiding during cluster failover [128](#)
- igroup show
 - vtic output [68, 69, 80](#)
- igroup throttles
 - borrowing queue resources [72](#)
 - creating [71](#)
 - defined [70](#)
 - destroying [71](#)
 - displaying information about [72](#)
 - displaying LUN statistics for [73](#)
 - displaying usage information [73](#)
 - how Data ONTAP uses [70](#)
 - how port sets affect [130](#)
 - how to use [70](#)
- igroups
 - borrowing queue resources for [72](#)
 - configuration [68](#)
 - mapping to LUNs [75](#)
 - requirements for creating [65](#)
- initiator
 - node
 - name [66](#)
 - node name
 - login [207](#)
- initiator groups
 - adding [79](#)
 - binding to port sets [132](#)
 - creating for FC using sanlun [66](#)
 - creating for iSCSI [64](#)
 - defined [63](#)
 - destroying [78](#)
 - displaying [80](#)
 - name rules [65](#)
 - naming [65](#)
 - ostype of [66](#)

- removing initiators from [79](#)
 - renaming [80](#)
 - requirements for creating [65](#)
 - setting the ostype for [81](#)
 - showing port set bindings [134](#)
 - type of [66](#)
 - unmapping LUNs from [77](#)
 - initiators
 - configuring adapters as [150](#)
 - displaying for iSCSI [95](#)
 - interface
 - disabling for iSCSI [89](#)
 - enabling for iSCSI [89](#)
 - IP addresses, displaying for iSCSI [89](#)
 - iqn type designator [206](#)
 - iSCSI
 - access lists [90](#)
 - connection, displaying [121](#)
 - creating access lists [90](#)
 - creating target portal groups [107](#)
 - default TCP port [207](#)
 - destroying target portal groups [108](#)
 - displaying access lists [91](#)
 - displaying initiators [95](#)
 - displaying statistics [116](#)
 - enabling error recovery levels 1 and 2 [83](#)
 - enabling on interface [89](#)
 - error messages [125](#)
 - explained [205](#)
 - host implementation [205](#)
 - how communication sessions work [209](#)
 - how nodes are identified [206](#)
 - implementation on the storage system [205](#)
 - iSNS [92](#)
 - license [84](#)
 - multi-connection sessions, enabling [82](#)
 - node
 - name [66](#)
 - node name rules [86](#)
 - nodes defined [205](#)
 - RADIUS [99](#)
 - removing interfaces from access lists [91](#)
 - security [96](#)
 - service, verifying [83](#)
 - session, displaying [120](#)
 - setup procedure [209](#)
 - target alias [87](#)
 - target IP addresses [89](#)
 - target node name [86](#)
 - target portal groups defined [105](#), [207](#)
 - troubleshooting [123](#)
 - using with HA pairs [209](#)
 - with HA pairs [122](#)
 - iscsi commands
 - iscsi alias [87](#)
 - iscsi connection [121](#)
 - iscsi initiator [95](#)
 - iscsi interface [88](#)
 - iscsi isns [93](#)
 - iscsi nodename [86](#)
 - iscsi portal [89](#)
 - iscsi security [97](#)
 - iscsi session [120](#)
 - iscsi start [85](#)
 - iscsi stats [116](#)
 - iscsi status [83](#)
 - iscsi stop [86](#)
 - iscsi tpgroup [106](#)
 - iSCSI license
 - deleting [85](#)
 - disabling [85](#)
 - enabling [84](#)
 - iSCSI service
 - disabling [85](#)
 - iSCSI target service
 - enabling [84](#)
 - iscsi.isns.rev option [93](#)
 - iscsi.max_connections_per_session option [82](#)
 - iscsi.max_error_recovery_level option [83](#)
 - iSNS
 - defined [208](#)
 - disabling [94](#)
 - server versions [92](#)
 - service for iSCSI [92](#)
 - updating immediately [94](#)
 - with vFiler units [94](#)
 - ISNS
 - and IPv6 [93](#)
 - registering [93](#)
- ## L
- license
 - FC [210](#)
 - iSCSI [84](#)
 - Logical
 - Block
 - Provisioning [17](#)
 - login
 - initiator

- checks [207](#)
- LUN
 - alignment issues [61](#)
 - autosize [33](#)
 - LUN [16](#)
 - on line [11](#)
 - provisioning [34](#)
 - read-only [76](#)
 - roadmap [8](#)
 - Snapshot copies [33](#)
 - space reserved [33](#)
 - space-reserved [18, 20](#)
 - staying online [16](#)
 - thin [34](#)
 - thinly provisioned [16](#)
 - thinly-provisioned
 - considerations [11](#)
 - strategies [11](#)
 - workflow [8](#)
- LUN clones
 - creating [186](#)
 - defined [185](#)
 - deleting Snapshot copies [188, 189](#)
 - displaying progress of split [188](#)
 - splitting from Snapshot copy [187](#)
 - stopping split [188](#)
- lun commands
 - lun clone create [186](#)
 - lun clone split [187, 188](#)
 - lun config_check [56](#)
 - lun destroy [61](#)
 - lun help [52](#)
 - lun map [45](#)
 - lun move [54](#)
 - lun offline [53](#)
 - lun online [53](#)
 - lun set reservation [55](#)
 - lun setup [44](#)
 - lun share [56](#)
 - lun show [59](#)
 - lun snap usage [193](#)
 - lun stats [58](#)
 - lun unmap [77, 78](#)
- LUN creation
 - description attribute [51](#)
 - host operating system type [49](#)
 - information required for [48](#)
 - LUN ID requirement [75](#)
 - ostype [49](#)
 - path name [48](#)
 - size specifiers [50](#)
 - space reservation default [51](#)
- LUN ID
 - range [75](#)
- LUN not visible [123](#)
- LUN reservations
 - how they work [54](#)
- LUN serial numbers
 - displaying
 - changing [57](#)
- LUN type
 - deciding [9](#)
- LUNs
 - autosize [42](#)
 - bringing online [53](#)
 - calculating rate of change [10](#)
 - checking settings for [56](#)
 - configuring [36, 37](#)
 - controlling availability [53](#)
 - creating [48](#)
 - displaying information [59](#)
 - displaying mapping [59](#)
 - displaying reads, writes, and operations for [58](#)
 - displaying serial numbers for [57](#)
 - enabling space reservations [55](#)
 - host operating system type [49](#)
 - layout [51](#)
 - management [52](#)
 - mapping guidelines [76](#)
 - mapping to igroups [75](#)
 - modifying description [54](#)
 - multiprotocol type [49](#)
 - offline [12](#)
 - ostype [49](#)
 - pre-allocation [18](#)
 - provisioning [36, 37](#)
 - removing [61](#)
 - renaming [54](#)
 - reserve
 - Snapshot [19](#)
 - restoring [198](#)
 - Snapshot
 - reserve [19](#)
 - snapshot copies [42](#)
 - snapshot copy [42](#)
 - space reserved [42](#)
 - space-reserved [18, 19, 36, 37](#)
 - statistics for igroup throttles [73](#)
 - taking offline [53](#)
 - thinly provisioned [18](#)

troubleshooting [123](#)
 unmapping from initiator group [77](#)

M

managing
 volumes [13](#)
 mapping conflicts
 overriding [129](#)
 misaligned
 LUNs and I/O [61](#)
 moving volumes
 DataMotion for Volumes [170](#)
 MPIO [63](#)
 multi-connection sessions
 enabling [82](#)
 multipathing
 requirements for cluster failover [130](#)
 MultiStore
 creating LUNs for vFiler units [46](#)

N

name rules
 igroups [65](#)
 iSCSI node name [86](#)
 node
 name
 iSCSI [66](#)
 node name
 rules for iSCSI [86](#)
 storage system [207](#)
 node type designator
 eui [207](#)
 iqn [206](#)
 nodes
 FCP [210](#)
 iSCSI [205](#)

O

offline
 LUN
 going offline [12](#)
 prevention [12](#)
 onboard adapters
 configuring for target mode [148](#)
 options
 iscsi.isns.rev [93](#)
 iscsi.max_connections_per_session [82](#)

 iscsi.max_error_recovery_level [83](#)
 ostype
 determining [49](#)
 displaying [59](#)
 required for LUN creation [48](#)
 setting [81](#)
 out of space
 LUN [16](#)
 over-provisioning example [11](#)
 over-subscribed storage [9, 11, 18, 30](#)

P

paths [204](#)
 plex
 defined [21](#)
 port set commands
 port set add [133](#)
 port set create [131](#)
 port set destroy [134](#)
 port set remove [133](#)
 port set show [134](#)
 port sets
 adding ports [133](#)
 binding to igroups [132](#)
 creating [131](#)
 defined [130](#)
 destroying [134](#)
 how they affect igroup throttles [130](#)
 how upgrades affect [130](#)
 removing [133](#)
 showing igroup bindings [134](#)
 unbinding igroups [132](#)
 viewing ports in [134](#)
 Protocols
 supported types [82](#)
 provisioning
 best practices [13](#)
 guidelines [22](#)
 methods of [44, 45](#)
 options [9, 30](#)
 thin [11](#)

Q

qtrees
 defined [21](#)
 quotas [51](#)

R**RADIUS**

- adding a RADIUS server [101](#)
- clearing statistics for [104](#)
- defining as the authentication method [100](#)
- displaying statistics for [104](#)
- displaying the status of [103](#)
- enabling for CHAP authentication [102](#)
- overview [99](#)
- removing a RADIUS server [104](#)
- server
 - client service [99](#)
 - starting the client service [101](#)
 - stopping the service [103](#)

RAID-level mirroring

- described [21](#)

rate of change

- LUN
 - type
 - determining [9](#)
 - rate of change [9](#)

Rate of change

- calculating [10](#)

reclaim

- space [16](#)

report_scsi_name

- automatic enablement [69](#)
- igroup [69](#)
- manually enabling [69](#)

reservations

- how they work [54](#)

reserves

- considerations for setting fractional [31](#)

restoring

- LUNs [198](#)

resuming volume move

- data copy phase [176](#)

roadmap

- LUN [8](#)

S**san****SAN systems**

- backing up [199](#)

sanlun

- creating igroups for FC [66](#)

SCSI

- SBC-3 standard [17](#)

- thin provisioning [17](#)

SCSI command [68](#)**serial numbers**

- for LUNs [57](#)

session

- checks [207](#)

setup phase

- volume move [172](#)

snap commands

- snap restore [189](#), [196](#)

snap reserve

- setting the percentage [41](#)

SnapDrive

- about [204](#)

SnapMirror destinations

- mapping read-only LUNs [76](#)

Snapshot

- copy
 - duration [23](#)
 - room [23](#)

Snapshot copies

- autodelete [30](#)
- deleting busy [193](#)
- no pre-allocation [19](#)
- not space-reserved [20](#)
- pre-allocated [18](#), [20](#)
- schedule, turning off [40](#)

Snapshot reserve

- without pre-allocated [34](#), [36](#), [37](#)

space

- how Data ONTAP can automatically add FlexVol volume [33](#)
- reclamation [16](#), [17](#)

space allocation

- guidelines [51](#)
- host support [16](#)
- LUN [9](#), [30](#)
- Snapshot copies [9](#), [30](#)
- space-reserved LUN [18–20](#)
- thin [18](#)
- volume [9](#), [30](#)

space information

- commands to display [165](#)

Space Reclamation [14](#)**space reservations**

- See* reservations

space-reserved

- LUN [18–20](#)

space

- allocation [18](#)
- statistics
 - collected for VAAI features [179](#)
 - displaying for iSCSI [116](#)
- stats command
 - viewing statistics for VAAI features [181](#)
- storage administrator [11](#)
- storage efficiency [9](#), [11](#), [18](#), [19](#), [30](#)
- storage system node name
 - defined [207](#)
- storage units
 - configuring [28](#)
 - types [21](#)
- SyncMirror
 - plexes [21](#)

T

- target adapters
 - displaying statistics [158](#)
 - displaying WWNN [157](#)
 - FCoE [161](#)
- target alias for iSCSI [87](#)
- target node name, iSCSI [86](#)
- target portal groups
 - about [105](#)
 - adding interfaces [108](#)
 - adding IP addresses to IP-based groups [115](#)
 - caution when upgrading Data ONTAP [106](#)
 - creating [107](#)
 - creating IP-based [114](#)
 - defined [207](#)
 - deleting IP-based groups [115](#)
 - destroying [108](#)
 - displaying information about IP-based groups [113](#)
 - enabling IP-based [111](#)
 - removing interfaces [109](#)
 - removing IP addresses [116](#)
 - upgrade and revert implications for [110](#)
- targets
 - configuring adapters as [148](#)
- TCP port
 - default for iSCSI [207](#)
- thin provisioned
 - LUN [34](#)
- thin provisioning
 - about [11](#)
 - best practices [13](#)
- thinly-provisioned
 - LUN [11](#)

- traditional volumes
 - described [21](#)
- troubleshooting
 - iSCSI error messages [125](#)
 - LUN [123](#)
- troubleshooting iSCSI [123](#)

U

- unified Ethernet
 - overview [161](#)
- unified target adapters
 - data center bridging [162](#)
 - managing [161](#)
- UTA [161](#)

V

- VAAI features
 - copy offload [177](#)
 - extended copy feature [177](#)
 - methods for determining support of [179](#)
 - statistics collected [179](#)
 - VERIFY AND WRITE feature [177](#)
 - viewing statistics [181](#)
 - WRITE SAME feature [177](#)
- VERIFY AND WRITE feature
 - environment [178](#)
 - invoked automatically [178](#)
 - statistics collected [179](#)
 - VAAI feature [177](#)
 - viewing statistics [181](#)
- vFiler units
 - authentication using CHAP [96](#)
 - creating LUNs for [46](#)
 - using iSCSI igroups with [67](#)
 - with iSNS [94](#)
- volume
 - configuring [34](#), [36](#), [37](#)
 - space-reserved LUN [18–20](#)
 - thinly provisioned [20](#)
- volume autosize [13](#)
- volume move
 - abort [177](#)
 - about [170](#)
 - automatic cutover [177](#)
 - cancel [177](#)
 - conflicting operations [171](#)
 - cutover phase
 - temporary destination volume [173](#)

- data copy phase [173](#)
- data transfer [176](#)
- DataMotion for Volumes [170](#)
- destination volume [170](#), [174](#)
- high priority
 - I/O operations [175](#)
- manual cutover [177](#)
- operations supported [171](#)
- operations unsupported [171](#)
- pausing [175](#), [176](#)
- requirements [171–173](#)
- resuming volume move [176](#)
- scenarios [170](#)
- setup phase [172](#)
- SLA requirements [170](#)
- source volume [174](#)
- temporary volume [174](#)
- volume status [176](#)
- volume size
 - autodelete [24](#)
 - no Snapshot copies [27](#)
 - Snapshot copies [25](#)
- volumes
 - autosizing [30](#), [31](#)
 - configuration options [9](#), [30](#)
 - configuring [9](#), [30](#), [34](#)
 - creating [29](#)
 - default settings [39](#)
 - estimating [25](#)
 - estimating required size of [30](#)
 - fractional reserve
 - considerations for setting [31](#)
 - moving nondisruptively [170](#)
 - required size [25](#)
 - snap_delete [42](#)
 - Snapshot

- reserve [18](#)
 - space
 - allocation [34](#)
 - space reservation [42](#)
 - thinly provisioned LUN [18](#)
- vtic in igroup show output [68](#), [69](#), [80](#)

W

- workflow
 - LUN [8](#)
- WRITE SAME feature
 - environment [178](#)
 - invoked automatically [178](#)
 - statistics collected [179](#)
 - VAAI feature [177](#)
 - viewing statistics [181](#)
- WWNN
 - changing [143](#)
 - displaying for a target adapter [157](#)
- WWPN
 - assignment [212](#)
 - changing for a target adapter [141](#)
 - usage [211](#)
- WWPN aliases
 - about [143](#)
 - creating [143](#)
 - displaying [144](#)
 - removing [144](#)

Z

- zero fat provisioning [11](#)