

SAN storage management

ONTAP 9

NetApp August 19, 2025

This PDF was generated from https://docs.netapp.com/us-en/ontap/san-admin/san-host-provisioning-concept.html on August 19, 2025. Always check docs.netapp.com for the latest.

Table of Contents

SAN storage management	1
SAN concepts	1
SAN provisioning with iSCSI	1
iSCSI service management	2
SAN provisioning with FC	8
SAN provisioning with NVMe	9
SAN volumes	10
SAN host-side space management	15
About igroups	16
Specify initiator WWPNs and iSCSI node names for an igroup	17
Storage virtualization with VMware and Microsoft copy offload	17
SAN administration	24
SAN provisioning	24
NVMe provisioning	34
Manage LUNs	45
Manage igroups and portsets	59
Manage iSCSI protocol	65
Manage FC protocol	72
Manage NVMe protocol	74
Manage systems with FC adapters	84
Manage LIFs for all SAN protocols	92
Enable ONTAP space allocation for SAN protocols	98
Recommended volume and file or LUN configuration combinations	100
SAN data protection	105
Learn about ONTAP data protection methods for SAN environments	105
Restore a single LUN from an ONTAP snapshot	106
Restore all LUNs in a volume from an ONTAP snapshot.	108
Protect your data with ONTAP FlexClone LUNs	109
Configure and use SnapVault backups in a SAN environment	110
Recommended configuration to connect a host backup system to ONTAP.	118
Use a host backup system to protect a LUN on your ONTAP storage system	119
SAN configuration reference	121
Learn about ONTAP SAN configuration	121
iSCSI configurations	121
FC configurations	124
FCoE configurations	131
FC and FCoE zoning	135
Requirements for SAN hosts connected to ONTAP and non-NetApp systems	138
SAN configurations in a MetroCluster environment	139
ONTAP support for SAN host multipathing	141
Configuration limits	142

SAN storage management

SAN concepts

SAN provisioning with iSCSI

In SAN environments, storage systems are targets that have storage target devices. For iSCSI and FC, the storage target devices are referred to as LUNs (logical units). For Non-Volatile Memory Express (NVMe) over Fibre Channel, the storage target devices are referred to as namespaces.

You configure storage by creating LUNs for iSCSI and FC or by creating namespaces for NVMe. The LUNs or namespaces are then accessed by hosts using Internet Small Computer Systems Interface (iSCSI) or Fibre Channel (FC) protocol networks.

To connect to iSCSI networks, hosts can use standard Ethernet network adapters (NICs), TCP offload engine (TOE) cards with software initiators, converged network adapters (CNAs), or dedicated iSCSI host bus adapters (HBAs).

To connect to FC networks, hosts require FC HBAs or CNAs.

Supported FC protocols include:

- FC
- FCoE
- NVMe

iSCSI target node network connections and names

iSCSI target nodes can connect to the network in several ways:

- Over Ethernet interfaces using software that is integrated into ONTAP.
- Over multiple system interfaces, with an interface used for iSCSI that can also transmit traffic for other protocols, such as SMB and NFS.
- Using a unified target adapter (UTA) or a converged network adapter (CNA).

Every iSCSI node must have a node name.

The two formats, or type designators, for iSCSI node names are *iqn* and *eui*. The SVM iSCSI target always uses the iqn-type designator. The initiator can use either the iqn-type or eui-type designator.

Storage system node name

Each SVM running iSCSI has a default node name based on a reverse domain name and a unique encoding number.

The node name is displayed in the following format:

iqn.1992-08.com.netapp:sn.unique-encoding-number

The following example shows the default node name for a storage system with a unique encoding number:

TCP port for iSCSI

The iSCSI protocol is configured in ONTAP to use TCP port number 3260.

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.

Related information

NetApp Documentation: ONTAP SAN Host Configuration

iSCSI service management

iSCSI service management

You can manage the availability of the iSCSI service on the iSCSI logical interfaces of the storage virtual machine (SVM) by using the vserver iscsi interface enable or vserver iscsi interface disable commands.

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

How iSCSI is implemented on the host

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

You can implement iSCSI in one of the following ways:

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

The iSCSI protocol processing is still performed by host software.

How iSCSI authentication works

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

iSCSI authentication methods are:

• Challenge Handshake Authentication Protocol (CHAP)--The initiator logs in using a CHAP user name and password.

You can specify a CHAP password or generate a hexadecimal secret password. There are two types of CHAP user names and passwords:

• Inbound—The storage system authenticates the initiator.

Inbound settings are required if you are using CHAP authentication.

• Outbound—This is an optional setting to enable the initiator to authenticate the storage system.

You can use outbound settings only if you define 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 the 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.

Related information

Windows Multipathing Options with Data ONTAP: Fibre Channel and iSCSI

iSCSI initiator security management

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.

iSCSI endpoint isolation

Existing iSCSI security commands can accept an IP address range, or multiple IP addresses.

All iSCSI initiators must provide origination IP addresses when establishing a session or connection with a target. This new functionality prevents an initiator from logging into the cluster if the origination IP address is unsupported or unknown, providing a unique identification scheme. Any initiator originating from an unsupported or unknown IP address will have their login rejected at the iSCSI session layer, preventing the initiator from accessing any LUN or volume within the cluster.

Implement this new functionality with two new commands to help manage pre-existing entries.

Add initiator address range

Improve iSCSI initiator security management by adding an IP address range, or multiple IP addresses with the vserver iscsi security add-initiator-address-range command.

cluster1::> vserver iscsi security add-initiator-address-range

Remove initiator address range

Remove an IP address range, or multiple IP addresses, with the vserver iscsi security removeinitiator-address-range command.

cluster1::> vserver iscsi security remove-initiator-address-range

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.

Authentication	Outbound	Inbound	Match?
Unidirectional	Host initiator user name and password	Storage user name and password	Must match
Bidirectional	Host initiator user name and password	Storage user name and password	Must match
Bidirectional	Storage user name and password	Host initiator user name and password	Must match

The outbound user name and password for the host initiator must be different than the outbound user name and password for the storage system.

Guidelines for using CHAP authentication

You should follow certain guidelines when using CHAP authentication.

- If you define an inbound user name and password on the storage system, you must use the same user name and password for outbound CHAP settings on the initiator. If you also define an outbound user name and password on the storage system to enable bidirectional authentication, you must use the same user name and password for inbound CHAP settings on the initiator.
- You cannot use the same user name and password for inbound and outbound settings on the storage system.
- CHAP user names can be 1 to 128 bytes.

A null user name is not allowed.

• CHAP passwords (secrets) can be 1 to 512 bytes.

Passwords can be hexadecimal values or strings. For hexadecimal values, you should enter the value with a prefix of "0x" or "0X". A null password is not allowed.

ONTAP allows the use of special characters, non-English letters, numbers and spaces for CHAP passwords (secrets). However, this is subject to host restrictions. If any of these are not allowed by your specific host, they cannot be used.



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.

For additional restrictions, you should see the initiator's documentation.

How using iSCSI interface access lists to limit initiator interfaces can increase performance and security

ISCSI interface access lists can be used to limit the number of LIFs in an SVM that an initiator can access, thereby increasing performance and security.

When an initiator begins a discovery session using an iSCSI SendTargets command, it receives the IP addresses associated with the LIF (network interface) that is in the access list. By default, all initiators have access to all iSCSI LIFs in the SVM. You can use the access list to restrict the number of LIFs in an SVM that an initiator has access to.

Internet Storage Name Service (iSNS) in ONTAP

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 IQN's, and portal groups.

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

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

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

What an iSNS server does

An iSNS server uses the Internet Storage Name Service (iSNS) protocol to maintain information about active iSCSI devices on the network, including their IP addresses, iSCSI node names (IQNs), 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 can obtain these servers from a vendor supported by NetApp.

How SVMs interact with an iSNS server

The iSNS server communicates with each storage virtual machine (SVM) through the SVM management LIF.

The management LIF registers all iSCSI target node name, alias, and portal information with the iSNS service for a specific SVM.

In the following example, SVM "VS1" uses SVM management LIF "VS1_mgmt_lif" to register with the iSNS server. During iSNS registration, an SVM sends all the iSCSI LIFs through the SVM management LIF to the iSNS Server. After the iSNS registration is complete, the iSNS server has a list of all the LIFs serving iSCSI in "VS1". If a cluster contains multiple SVMs, each SVM must register individually with the iSNS server to use the iSNS service.



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



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

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

A delay of a few minutes might occur between the time you make the configuration changes and when ONTAP sends the update to the iSNS server. Force an immediate update of the iSNS information on the iSNS server: vserver iscsi isns update. Learn more about vserver iscsi isns update in the ONTAP command reference.

Commands for managing iSNS

ONTAP provides commands to manage your iSNS service.

If you want to	Use this command
Configure an iSNS service	vserver iscsi isns create
Start an iSNS service	vserver iscsi isns start
Modify an iSNS service	vserver iscsi isns modify
Display iSNS service configuration	vserver iscsi isns show
Force an update of registered iSNS information	vserver iscsi isns update

Stop an iSNS service	vserver iscsi isns stop
Remove an iSNS service	vserver iscsi isns delete
View the man page for a command	man command name

Learn more about vserver iscsi isns in the ONTAP command reference.

SAN provisioning with FC

You should be aware of the important concepts that are required to understand how ONTAP implements an FC SAN.

How FC target nodes connect to the network

Storage systems and hosts have adapters so that they can be connected to FC switches with cables.

When a node is connected to the FC SAN, each SVM registers the World Wide Port Name (WWPN) of its LIF with the switch Fabric Name Service. The WWNN of the SVM and the WWPN of each LIF is automatically assigned by ONTAP..



Direct-connection to nodes from hosts with FC is not supported, NPIV is required and this requires a switch to be used.With iSCSI sessions, communication works with connections that are either network routed or direct-connect. However, both of these methods are supported with ONTAP.

How FC nodes are identified

Each SVM configured with FC is identified by a worldwide node name (WWNN).

How WWPNs are used

WWPNs identify each LIF in an SVM configured to support FC. These LIFs use the physical FC ports in each node in the cluster, which can be FC target cards, UTA or UTA2 configured as FC or FCoE in the nodes.

• Creating an initiator group

The WWPNs of the host's HBAs are used to create an initiator group (igroup). An igroup is used to control host access to specific LUNs. You can create an igroup by specifying a collection of WWPNs of initiators in an FC network. When you map a LUN on a storage system to an igroup, you can grant all the initiators in that group access to that LUN. If a host's WWPN is not in an igroup that is mapped to a LUN, that host does not have access to the LUN. This means that the LUNs do not appear as disks on that host.

You can also create port sets to make a LUN visible only on specific target ports. A port set consists of a group of FC target ports. You can bind an igroup to a port set. Any host in the igroup can access the LUNs only by connecting to the target ports in the port set.

• Uniquely identifying FC LIFs

WWPNs uniquely identify each FC logical interface. The host operating system uses the combination of the WWNN and WWPN to identify SVMs and FC LIFs. Some operating systems require persistent binding to

ensure that the LUN appears at the same target ID on the host.

How worldwide name assignments work

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

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

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

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

How FC switches are identified

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

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

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

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

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

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

SAN provisioning with NVMe

Beginning with ONTAP 9.4, NVMe/FC is supported in SAN environment. NVMe/FC enables storage administrators to provision namespaces and subsystems and then map the namespaces to subsystems, similar to the way LUNs are provisioned and mapped to igroups for FC and iSCSI.

An NVMe namespace is a quantity of non-volatile memory that can be formatted into logical blocks. Namespaces are the equivalent of LUNs for FC and iSCSI protocols, and an NVMe subsystem is analogous to an igroup. An NVMe subsystem can be associated with initiators so that namespaces within the subsystem can be accessed by the associated initiators.



Although analogous in function, NVMe namespaces do not support all features supported by LUNs.

Beginning with ONTAP 9.5 a license is required to support host-facing data access with NVMe. If NVMe is enabled in ONTAP 9.4, a 90 day grace period is given to acquire the license after upgrading to ONTAP 9.5. If you have ONTAP One, the NVMe licenses is included. You can enable the license using the following command:

system license add -license-code NVMe_license_key

Related information

NetApp Technical Report 4684: Implementing and Configuring Modern SANs with NVMe/FC

SAN volumes

About SAN volumes overview

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



Putting SAN LUNs and NAS shares in the same FlexVol volume is not recommended. You should provision separate FlexVol volumes specifically for your SAN LUNs and you should provision separate FlexVol volumes specifically to your NAS shares. This simplifies management and replication deployments and parallels the way FlexVol volumes are supported in Active IQ Unified Manager (formerly OnCommand Unified Manager).

Thin provisioning for volumes

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

You create a thin-provisioned FlexVol volume by setting its -space-guarantee option to none.

Thick provisioning for volumes

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

You create a thick-provisioned FlexVol volume by setting its -space-slo (service level objective) option to thick.

Semi-thick provisioning for volumes

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

Note: The following functionality is not supported on volumes that use semi-thick provisioning:

- · Storage efficiency technologies such as deduplication, compression, and compaction
- Microsoft Offloaded Data Transfer (ODX)

You create a semi-thick-provisioned FlexVol volume by setting its -space-slo (service level objective) option to semi-thick.

Use with space-reserved files and LUNs

A space-reserved file or LUN is one for which storage is allocated when it is created. Historically, NetApp has used the term "thin-provisioned LUN" to mean a LUN for which space reservation is disabled (a non-space-reserved LUN).

Note: Non-space-reserved files are not generally referred to as "thin-provisioned files".

The following table summarizes the major differences in how the three volume provisioning options can be used with space-reserved files and LUNs:

Volume provisioning	LUN/file space reservation	Overwrites	Protection data ²	Storage efficiency ³
Thick	Supported	Guaranteed ¹	Guaranteed	Supported
Thin	No effect	None	Guaranteed	Supported
Semi-thick	Supported	Best effort ¹	Best effort	Not supported

Notes

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

Support for SCSI thin-provisioned LUNs

ONTAP supports T10 SCSI thin-provisioned LUNs as well as NetApp thin-provisioned LUNs. T10 SCSI thin provisioning enables host applications to support SCSI features including LUN space reclamation and LUN space monitoring capabilities for blocks environments. T10 SCSI thin provisioning must be supported by your SCSI host software.

You use the ONTAP space-allocation setting to enable/disable support for the T10 thin provisioning on a LUN. You use the ONTAP space-allocation enable setting to enable T10 SCSI thin provisioning on a LUN.

The [-space-allocation {enabled|disabled}] command in the ONTAP command reference has more information to enable/disable support for the T10 thin provisioning and to enable T10 SCSI thin provisioning on a LUN.

Configure volume provisioning options

You can configure a volume for thin provisioning, thick provisioning, or semi-thick provisioning.

About this task

Setting the -space-slo option to thick ensures the following:

- The entire volume is preallocated in the aggregate. You cannot use the volume create or volume modify command to configure the volume's -space-guarantee option.
- 100% of the space required for overwrites is reserved. You cannot use the volume modify command to configure the volume's -fractional-reserve option

Setting the -space-slo option to semi-thick ensures the following:

- The entire volume is preallocated in the aggregate. You cannot use the volume create or volume modify command to configure the volume's -space-guarantee option.
- No space is reserved for overwrites. You can use the volume modify command to configure the volume's -fractional-reserve option.
- Automatic deletion of snapshots is enabled.

Step

1. Configure volume provisioning options:

```
volume create -vserver vserver_name -volume volume_name -aggregate
aggregate name -space-slo none|thick|semi-thick -space-guarantee none|volume
```

The -space-guarantee option defaults to none for AFF systems and for non-AFF DP volumes. Otherwise, it defaults to volume. For existing FlexVol volumes, use the volume modify command to configure provisioning options.

The following command configures vol1 on SVM vs1 for thin provisioning:

```
cluster1::> volume create -vserver vs1 -volume vol1 -space-guarantee
none
```

The following command configures vol1 on SVM vs1 for thick provisioning:

cluster1::> volume create -vserver vs1 -volume vol1 -space-slo thick

The following command configures vol1 on SVM vs1 for semi-thick provisioning:

```
cluster1::> volume create -vserver vs1 -volume vol1 -space-slo semi-
thick
```

SAN volume configuration options

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

Autogrow

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

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

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

Autoshrink

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

Snapshot autodelete

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

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

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

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

Snapshot reserve

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

Requirement for moving volumes in SAN environments

Before you move a volume that contains LUNs or namespaces, you must meet certain requirements.

• For volumes containing one or more LUNs, you should have a minimum of two paths per LUN (LIFs) connecting to each node in the cluster.

This eliminates single points of failure and enables the system to survive component failures.

• For volumes containing namespaces, the cluster must be running ONTAP 9.6 or later.

Volume move is not supported for NVMe configurations running ONTAP 9.5.

Considerations for setting fractional reserve

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

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

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

- Deduplication is not in use
- Compression is not in use
- FlexClone sub-files are not in use
- All FlexClone files and FlexClone LUNs are enabled for automatic deletion

This is not the default setting. You must explicitly enable automatic deletion, either at creation time or by modifying the FlexClone file or FlexClone LUN after it is created.

- ODX and FlexClone copy offload are not in use
- Volume guarantee is set to volume
- File or LUN space reservation is enabled
- Volume Snapshot reserve is set to 0
- Volume snapshot automatic deletion is enabled with a commitment level of destroy, a destroy list of lun_clone, vol_clone, cifs_share, file_clone, sfsr, and a trigger of volume

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

Note that if your rate of change is high, in rare cases the snapshot automatic deletion could fall behind,

resulting in the volume running out of space, even with all of the above required configuration settings in use.

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

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

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

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

SAN host-side space management

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

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

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

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

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

Simplified host management with SnapCenter

You can use SnapCenter software to simplify some of the management and data protection tasks associated with iSCSI and FC storage. SnapCenter is an optional management package for Windows and UNIX hosts.

You can use SnapCenter Software to easily create virtual disks from pools of storage that can be distributed among several storage systems and to automate storage provisioning tasks and simplify the process of creating snapshots and clones from snapshots consistent with host data.

See NetApp product documentation for more information on SnapCenter.

Related links

Enable ONTAP space allocation for SAN protocols

About igroups

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

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

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

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

Example of how igroups give LUN access

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

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

Hosts with HBA WWPNs, IQNs, or EUIs	igroups	WWPNs, IQNs, EUIs added to igroups	LUNs mapped to igroups
Host1, single-path (iSCSI software initiator) iqn.1991- 05.com.microsoft:host1	group1	iqn.1991- 05.com.microsoft:host1	/vol/vol2/lun1
Host2, multipath (two HBAs) 10:00:00:00:c9:2b:6b:3c 10:00:00:00:c9:2b:02:3c	group2	10:00:00:00:c9:2b:6b:3c 10:00:00:00:c9:2b:02:3c	/vol/vol2/lun2

Hosts with HBA WWPNs, IQNs, or EUIs	igroups	WWPNs, IQNs, EUIs added to igroups	LUNs mapped to igroups
Host3, multipath, clustered with host 4 10:00:00:00:c9:2b:32:1b 10:00:00:00:c9:2b:41:02	group3	10:00:00:00:c9:2b:32:1b 10:00:00:00:c9:2b:41:02 10:00:00:00:c9:2b:51:2c 10:00:00:00:c9:2b:47:a2	/vol/vol2/qtree1/lu n3
Host4, multipath, clustered (not visible 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/qtree2/lu n4 /vol/vol2/qtree1/lu n5

Specify initiator WWPNs and iSCSI node names for an igroup

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

Follow the instructions in your Host Utilities documentation to obtain WWPNs and to find the iSCSI node names associated with a specific host. For hosts running ESX software, use Virtual Storage Console.

Storage virtualization with VMware and Microsoft copy offload

Storage virtualization with VMware and Microsoft copy offload overview

VMware and Microsoft support copy offload operations to increase performance and network throughput. You must configure your system to meet the requirements of the VMware and Windows operating system environments to use their respective copy offload functions.

When using VMware and Microsoft copy offload in virtualized environments, your LUNs must be aligned. Unaligned LUNs can degrade performance.

Advantages of using a virtualized SAN environment

Creating a virtualized environment by using storage virtual machines (SVMs) and LIFs enables you to expand your SAN environment to all of the nodes in your cluster.

• Distributed management

You can log in to any node in the SVM to administer all of the nodes in a cluster.

Increased data access

With MPIO and ALUA, you have access to your data through any active iSCSI or FC LIFs for the SVM.

Controlled LUN access

If you use SLM and portsets, you can limit which LIFs an initiator can use to access LUNs.

How LUN access works in a virtualized environment

In a virtualized environment, LIFs enable hosts (clients) to access LUNs through optimized and unoptimized paths.

A LIF is a logical interface that connects the SVM to a physical port. Although multiple SVMs can have multiple LIFs on the same port, a LIF belongs to one SVM. You can access LUNs through the SVMs LIFs.

Example of LUN access with a single SVM in a cluster

In the following example, Host 1 connects to LIF1.1 and LIF1.2 in SVM-1 to access LUN1. LIF1.1 uses the physical port node1:0c and LIF1.2 uses the node2:0c. LIF1.1 and LIF1.2 belongs only to SVM-1. If a new LUN is created on node 1 or node 2, for SVM-1, then it can use these same LIFs. If a new SVM is created, then new LIFs can be created using physical ports 0c or 0d on both the nodes.



Example of LUN access with multiple SVMs in a cluster

A physical port can support multiple LIFs serving different SVMs. Because LIFs are associated with a particular SVM, the cluster nodes can send the incoming data traffic to the correct SVM. In the following example, each node from 1 through 4 has a LIF for SVM-2 using the physical port 0c on each node. Host 1 connects to LIF1.1 and LIF1.2 in SVM-1 to access LUN1. Host 2 connects to LIF2-1 and LIF2-2 in SVM-2 to access LUN2. Both SVMs are sharing the physical port 0c on the nodes 1 and 2. SVM-2 has additional LIFs that Host 2 is using to access LUNs 3 and 4. These LIFs are using physical port 0c on nodes 3 and 4. Multiple SVMs can share the physical ports on the nodes.



Example of an active or optimized path to a LUN from a host system

In an active or optimized path, the data traffic does not travel over the cluster network; it travels the most direct route to the LUN. The active or optimized path to LUN1 is through LIF1.1 in node1, using physical port 0c. Host 2 has two active or optimized paths, one path to node1, LIF2.1, which is sharing physical port 0c and the other path to node4, LIF2.4, which is using physical port 0c.



Example of an active or unoptimized path (indirect) path to a LUN from a host system

In an active or unoptimized path (indirect) path, the data traffic travels over the cluster network. This issue occurs only if all the active or optimized paths from a host are unavailable to handle traffic. If the path from Host 2 to SVM-2 LIF2.4 is lost, then access to LUN3 and LUN4 traverses the cluster network. Access from Host 2 uses LIF2.3 on node3. Then the traffic enters the cluster network switch and backs up to node4 for access to the LUN3 and LUN4. It will then traverse back over the cluster network switch and then back out through LIF2.3 to Host 2. This active or unoptimized path is used until the path to LIF2.4 is restored or a new LIF is established for SVM-2 on another physical port on node 4.



Improve VMware VAAI performance for ESX hosts

ONTAP supports certain VMware vStorage APIs for Array Integration (VAAI) features when the ESX host is running ESX 4.1 or later. These features help offload operations from the ESX host to the storage system and increase the network throughput. The ESX host enables the features automatically in the correct environment.

The VAAI feature supports the following SCSI commands:

• EXTENDED_COPY

This feature enables the host to initiate the transfer of data between the LUNs or within a LUN without involving the host in the data transfer. This results in saving ESX CPU cycles and increasing the network throughput. The extended copy feature, also known as "copy offload," is used in scenarios such as cloning a virtual machine. When invoked by the ESX host, the copy offload feature copies the data within the storage system rather than going through the host network. Copy offload transfers data in the following ways:

- Within a LUN
- Between LUNs within a volume
- Between LUNs on different volumes within a storage virtual machine (SVM)
- Between LUNs on different SVMs within a cluster
 If this feature cannot be invoked, the ESX host automatically uses the standard READ and WRITE commands for the copy operation.
- WRITE_SAME

This feature offloads the work of writing a repeated pattern, such as all zeros, to a storage array. The ESX host uses this feature in operations such as zero-filling a file.

• COMPARE_AND_WRITE

This feature bypasses certain file access concurrency limits, which speeds up operations such as booting up virtual machines.

Requirements for using the VAAI environment

The VAAI features are part of the ESX operating system and are automatically invoked by the ESX host when you have set up the correct environment.

The environment requirements are as follows:

- The ESX host must be running ESX 4.1 or later.
- The NetApp storage system that is hosting the VMware datastore must be running ONTAP.
- (Copy offload only) The source and the destination of the VMware copy operation must be hosted on the same storage system within the same cluster.



The copy offload feature currently does not support copying data between VMware datastores that are hosted on different storage systems.

Determine if VAAI features are supported by ESX

To confirm whether the ESX operating system supports the VAAI features, you can check the vSphere Client or use any other means of accessing the host. ONTAP supports the SCSI commands by default.

You can check your ESX host advanced settings to determine whether VAAI features are enabled. The table indicates which SCSI commands correspond to ESX control names.

SCSI command	ESX control name (VAAI feature)
EXTENDED_COPY	HardwareAcceleratedMove
WRITE_SAME	HardwareAcceleratedInit
COMPARE_AND_WRITE	HardwareAcceleratedLocking

Microsoft Offloaded Data Transfer (ODX)

Microsoft Offloaded Data Transfer (ODX), also known as *copy offload*, enables direct data transfers within a storage device or between compatible storage devices without transferring the data through the host computer.

ONTAP supports ODX for both the SMB and SAN protocols.

In non-ODX file transfers, the data is read from the source and is transferred across the network to the host. The host transfers the data back over the network to the destination. In ODX file transfer, the data is copied directly from the source to the destination without passing through the host.

Because ODX offloaded copies are performed directly between the source and destination, significant performance benefits are realized if copies are performed within the same volume, including faster copy time for same volume copies, reduced utilization of CPU and memory on the client, and reduced network I/O bandwidth utilization. If copies are across volumes, there might not be significant performance gains compared to host-based copies.

For SAN environments, ODX is only available when it is supported by both the host and the storage system. Client computers that support ODX and have ODX enabled automatically and transparently use offloaded file transfer when moving or copying files. ODX is used regardless of whether you drag-and-drop files through Windows Explorer or use command-line file copy commands, or whether a client application initiates file copy requests.

Requirements for using ODX

If you plan to use ODX for copy offloads, you need to be familiar with volume support considerations, system requirements, and software capability requirements.

To use ODX, your system must have the following:

• ONTAP

ODX is automatically enabled in supported versions of ONTAP.

• Minimum source volume of 2 GB

For optimal performance, the source volume should be greater than 260 GB.

· ODX support on the Windows client

ODX is supported in Windows Server 2012 or later and in Windows 8 or later. The Interoperability Matrix contains the latest information about supported Windows clients.

NetApp Interoperability Matrix Tool

Copy application support for ODX

The application that performs the data transfer must support ODX. Application operations that support ODX include the following:

- Hyper-V management operations, such as creating and converting virtual hard disks (VHDs), managing snapshots, and copying files between virtual machines
- Windows Explorer operations
- · Windows PowerShell copy commands
- Windows command prompt copy commands The Microsoft TechNet Library contains more information about supported ODX applications on Windows servers and clients.
- If you use compressed volumes, the compression group size must be 8K.

32K compression group size is not supported.

ODX does not work with the following volume types:

- · Source volumes with capacities of less than 2 GB
- Read-only volumes
- FlexCache volumes



ODX is supported on FlexCache origin volumes.

· Semi-thick provisioned volumes

Special system file requirements

You can delete ODX files found in qtrees. Do not remove or modify any other ODX system files unless you are told by technical support to do so.

When using the ODX feature, there are ODX system files that exist in every volume of the system. These files enable point-in-time representation of data used during the ODX transfer. The following system files are in the root level of each volume that contains LUNs or files to which data was offloaded:

- .copy-offload (a hidden directory)
- .tokens (file under the hidden .copy-offload directory)

You can use the copy-offload delete-tokens -path dir_path -node *node_name* command to delete a qtree containing an ODX file.

Use cases for ODX

You should be aware of the use cases for using ODX on SVMs so that you can determine under what circumstances ODX provides you with performance benefits.

Windows servers and clients that support ODX use copy offload as the default way of copying data across remote servers. If the Windows server or client does not support ODX or the ODX copy offload fails at any point, the copy or move operation falls back to traditional reads and writes for the copy or move operation.

The following use cases support using ODX copies and moves:

• Intra-volume

The source and destination files or LUNs are within the same volume.

· Inter-volume, same node, same SVM

The source and destination files or LUNs are on different volumes that are located on the same node. The data is owned by the same SVM.

· Inter-volume, different nodes, same SVM

The source and destination files or LUNs are on different volumes that are located on different nodes. The data is owned by the same SVM.

Inter-SVM, same node

The source and destination file or LUNs are on different volumes that are located on the same node. The data is owned by different SVMs.

Inter-SVM, different nodes

The source and destination file or LUNs are on different volumes that are located on different nodes. The data is owned by different SVMs.

Inter-cluster

The source and destination LUNs are on different volumes that are located on different nodes across clusters. This is only supported for SAN and does not work for SMB.

There are some additional special use cases:

• With the ONTAP ODX implementation, you can use ODX to copy files between SMB shares and FC or iSCSI attached virtual drives.

You can use Windows Explorer, the Windows CLI or PowerShell, Hyper-V, or other applications that support ODX to copy or move files seamlessly using ODX copy offload between SMB shares and connected LUNs, provided that the SMB shares and LUNs are on the same cluster.

- Hyper-V provides some additional use cases for ODX copy offload:
 - You can use ODX copy offload pass-through with Hyper-V to copy data within or across virtual hard disk (VHD) files or to copy data between mapped SMB shares and connected iSCSI LUNs within the same cluster.

This allows copies from guest operating systems to pass through to the underlying storage.

- When creating fixed-sized VHDs, ODX is used for initializing the disk with zeros, using a well-known zeroed token.
- ODX copy offload is used for virtual machine storage migration if the source and destination storage is on the same cluster.



To take advantage of the use cases for ODX copy offload pass-through with Hyper-V, the guest operating system must support ODX and the guest operating system's disks must be SCSI disks backed by storage (either SMB or SAN) that supports ODX. IDE disks on the guest operating system do not support ODX pass-through.

SAN administration

SAN provisioning

SAN management overview

The content in this section shows you how to configure and manage SAN environments with the ONTAP command line interface (CLI) and System Manager in ONTAP 9.7 and later releases.

If you are using the classic System Manager (available only in ONTAP 9.7 and earlier), see these topics:

- iSCSI protocol
- FC/FCoE protocol

You can use the iSCSI and FC protocols to provide storage in a SAN environment.



With iSCSI and FC, storage targets are called LUNs (logical units) and are presented to hosts as standard block devices. You create LUNs and then map them to initiator groups (igroups). Initiator groups are tables of FC host WWPs and iSCSI host node names and control which initiators have access to which LUNs.

FC targets connect to the network through FC switches and host-side adapters and are identified by worldwide port names (WWPNs). iSCSI targets connect to the network through standard Ethernet network adapters (NICs), TCP offload engine (TOE) cards with software initiators, converged network adapters (CNAs) or dedicated host bust adapters (HBAs) and are identified by iSCSI qualified names (IQNs).

For more information

If you have an ASA r2 storage system (ASA A1K, ASA A90, ASA A70, ASA A50, ASA A30, or ASA A20), see the ASA r2 storage system documentation.

Learn about All-Flash SAN Array configurations

The NetApp All-Flash SAN Arrays (ASAs) are available beginning with ONTAP 9.7. ASAs are all-flash SAN-only solutions built on proven AFF NetApp platforms.

ASA platforms include the following:

- ASA A150
- ASA A250

- ASA A400
- ASA A800
- ASA A900
- ASA C250
- ASA C400
- ASA C800



Beginning with ONTAP 9.16.0, a simplified ONTAP experience specific to SAN-only customers is available on ASA r2 systems (ASA A1K, ASA A90, ASA A70, ASA A50, ASA A30, or ASA A20). If you have an ASA r2 system, see the ASA r2 system documentation.

ASA platforms use symmetric active-active for multipathing. All paths are active/optimized so in the event of a storage failover, the host does not need to wait for the ALUA transition of the failover paths to resume I/O. This reduces time to failover.

Set up an ASA

All-Flash SAN Arrays (ASAs) follow the same setup procedure as non-ASA systems.

System Manager guides you through the procedures necessary to initialize your cluster, create a local tier, configure protocols, and provision storage for your ASA.

Get started with ONTAP cluster set up.

ASA host settings and utilities

Host settings for setting up All-Flash SAN Arrays (ASAs) are the same as those for all other SAN hosts.

You can download the NetApp Host Utilities software for your specific hosts from the support site.

Ways to identify an ASA system

You can identify an ASA system using System Manager or using the ONTAP command line interface (CLI).

• From the System Manager dashboard: Click Cluster > Overview and then select the system node.

The **PERSONALITY** is displayed as **All-Flash SAN Array**.

• From the CLI: Enter the san config show command.

The "All-Flash SAN Array" value returns as true for ASA systems.

Learn more about san config show in the ONTAP command reference.

Related information

- Technical Report 4968: NetApp All-SAN Array Data Availability and Integrity
- NetApp Technical Report 4080: Best Practices for Modern SAN

Configure switches for FCoE

You must configure your switches for FCoE before your FC service can run over the

existing Ethernet infrastructure.

Before you begin

• Your SAN configuration must be supported.

For more information about supported configurations, see the NetApp Interoperability Matrix Tool.

• A Unified Target Adapter (UTA) must be installed on your storage system.

If you are using a UTA2, it must be set to cna mode.

• A converged network adapter (CNA) must be installed on your host.

Steps

- 1. Use your switch documentation to configure your switches for FCoE.
- 2. Verify that the DCB settings for each node in the cluster have been correctly configured.

run -node nodel -command dcb show

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

3. Verify that the FCoE login is working when the FC target port online status is true.

```
fcp adapter show -fields node,adapter,status,state,speed,fabric-
established,physical-protocol
```

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

Related information

- NetApp Interoperability Matrix Tool
- NetApp Technical Report 3800: Fibre Channel over Ethernet (FCoE) End-to-End Deployment Guide
- Cisco MDS 9000 NX-OS and SAN-OS Software Configuration Guides
- Brocade products

System Requirements

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

- The Interoperability Matrix must list your SAN configuration as supported.
- Your SAN environment must meet the SAN host and controller configuration limits specified in NetApp Hardware Universe for your version of the ONTAP software.
- A supported version of Host Utilities must be installed.

The Host Utilities documentation provides more information.

• You must have SAN LIFs on the LUN owning node and the owning node's HA partner.

Related information

- NetApp Interoperability Matrix Tool
- ONTAP SAN Host Configuration
- NetApp Technical Report 4017: Fibre Channel SAN Best Practices

What to know before you create a LUN

Before you begin setting up your LUNs on your cluster, you need to review these LUN guidelines.

Why actual LUN sizes slightly vary

You should be aware of the following regarding the size of your LUNs.

- When you create a LUN , the actual size of the LUN might vary slightly based on the OS type of the LUN. The LUN OS type cannot be modified after the LUN is created.
- If you create a LUN at the max LUN size, be aware that the actual size of the LUN might be slightly less. ONTAP rounds down the limit to be slightly less.
- The metadata for each LUN requires approximately 64 KB of space in the containing aggregate. When you create a LUN, you must ensure that the containing aggregate has enough space for the LUN's metadata. If the aggregate does not contain enough space for the LUN's metadata, some hosts might not be able to access the LUN.

Guidelines for assigning LUN IDs

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

Guidelines for mapping LUNs to igroups

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

Verify and add your protocol FC or iSCSI license

Before you can enable block access for a storage virtual machine (SVM) with FC or iSCSI, you must have a license. The FC and iSCSI licenses are included with ONTAP One.

System Manager

If you don't have ONTAP One, verify and add your FC or iSCSI license with ONTAP System Manager (9.7 and later).

- 1. In System Manager, select Cluster > Settings > Licenses
- ^{2.} If the license is not listed, select + Add and enter the license key.
- 3. Select Add.

CLI

If you don't have ONTAP One, verify and add your FC or iSCSI license with the ONTAP CLI.

1. Verify that you have a active license for FC or iSCSI.

system license show

Package	Туре	Description	Expiration
			_
Base	site	Cluster Base License	-
NFS	site	NFS License	-
CIFS	site	CIFS License	-
iscsi	site	iSCSI License	-
FCP	site	FCP License	-

2. If you do not have a active license for FC or iSCSI, add your license code.

license add -license-code <your license code>

Provision SAN storage

This procedure creates new LUNs on an existing storage VM which already has the FC or iSCSI protocol configured.

About this task

This procedure applies to FAS, AFF, and ASA systems. If you have an ASA r2 system (ASA A1K, ASA A90, ASA A70, ASA A50, ASA A30, ASA A20, or ASA C30), follow these steps to provision your storage. ASA r2 systems provide a simplified ONTAP experience specific to SAN-only customers.

If you need to create a new storage VM and configure the FC or iSCSI protocol, see Configure an SVM for FC or Configure an SVM for iSCSI.

If the FC license is not enabled, the LIFs and SVMs appear to be online but the operational status is down.

LUNs appear to your host as disk devices.



Asymmetric logical unit access (ALUA) is always enabled during LUN creation. You cannot change the ALUA setting.

You must use single initiator zoning for all of the FC LIFs in the SVM to host the initiators.

Beginning with ONTAP 9.8, when you provision storage, QoS is enabled by default. You can disable QoS, or choose a custom QoS policy during the provisioning process or at a later time.

System Manager

Create LUNs to provide storage for a SAN host using the FC or iSCSI protocol with ONTAP System Manager (9.7 and later).

To complete this task using System Manager Classic (available with 9.7 and earlier) refer to iSCSI configuration for Red Hat Enterprise Linux

Steps

- 1. Install the appropriate SAN host utilities on your host.
- 2. In System Manager, click **Storage > LUNs** and then click **Add**.
- 3. Enter the required information to create the LUN.
- 4. You can click **More Options** to do any of the following, depending upon your version of ONTAP.

Option	Available beginning with
 Assign QoS policy to LUNs instead of parent volume 	ONTAP 9.10.1
 More Options > Storage and Optimization 	
 Select Performance Service Level. 	
 To apply the QoS policy to individual LUNs instead of the entire volume, select Apply these performance limits enforcements to each LUN. 	
By default, performance limits are applied at the volume level.	
Create a new initiator group using existing initiator groups	ONTAP 9.9.1
 More Options > HOST INFORMATION 	
 Select New initiator group using existing initiator groups. The OS type for an igroup containing other igroups cannot be changed after it has been created. 	
 Add a description to your igroup or host initiator The description serves as an alias for the igroup or host initiator. 	ONTAP 9.9.1
 More Options > HOST INFORMATION 	
Create your LUN on an existing volume	ONTAP 9.9.1
By default, a new LUN is created in a new volume.	
 More Options > Add LUNs 	
 Select Group related LUNs. 	

Disable QoS	or choose a custom QoS policy	ONTAP 9.8
• More Op	tions > Storage and Optimization	
∘ Select Pe	rformance Service Level.	
i	In ONTAP 9.9.1 and later, if you select a custom QoS policy, you can also select manual placement on a specified local tier.	

- 5. For FC, zone your FC switches by WWPN. Use one zone per initiator and include all target ports in each zone.
- 6. Discover LUNs on your host.

For VMware vSphere, use Virtual Storage Console (VSC) to discover and initialize your LUNs.

- 7. Initialize the LUNs and optionally, create file systems.
- 8. Verify that the host can write and read data on the LUN.

CLI

Create LUNs to provide storage for a SAN host using the FC or iSCSI protocol with the ONTAP CLI.

1. Verify that you have a license for FC or iSCSI.

system license show

Package	Туре	Description	Expiration
			_
Base	site	Cluster Base License	-
NFS	site	NFS License	-
CIFS	site	CIFS License	-
iscsi	site	iSCSI License	-
FCP	site	FCP License	-

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

license add -license-code <your_license_code>

3. Enable your protocol service on the SVM:

For iSCSI:

vserver iscsi create -vserver <svm name> -target-alias <svm name>

For FC:

vserver fcp create -vserver <svm_name> -status-admin up

4. Create two LIFs for the SVMs on each node:

```
network interface create -vserver <svm_name> -lif <lif_name> -role
data -data-protocol <iscsi|fc> -home-node <node_name> -home-port
<port_name> -address <ip_address> -netmask <netmask>
```

NetApp supports a minimum of one iSCSI or FC LIF per node for each SVM serving data. However, two LIFS per node are required for redundancy. For iSCSI, it is recommended that you configure a minimum of two LIFs per node in separate Ethernet networks.

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

network interface show -vserver <svm name> <lif name>

6. Create your LUNs:

```
lun create -vserver <svm_name> -volume <volume_name> -lun <lun_name>
-size <lun_size> -ostype linux -space-reserve <enabled|disabled>
```

Your LUN name cannot exceed 255 characters and cannot contain spaces.



The NVFAIL option is automatically enabled when a LUN is created in a volume.

7. Create your igroups:

igroup create -vserver <svm_name> -igroup <igroup_name> -protocol <fcp|iscsi|mixed> -ostype linux -initiator <initiator name>

8. Map your LUNs to igroups:

```
lun mapping create -vserver <svm_name> -volume <volume_name> -lun
<lun name> -igroup <igroup name>
```

9. Verify that your LUNs are configured correctly:

lun show -vserver <svm_name>

- 10. Optionally, Create a port set and bind to an igroup.
- 11. Follow steps in your host documentation for enabling block access on your specific hosts.
- 12. Use the Host Utilities to complete the FC or iSCSI mapping and to discover your LUNs on the host.

Related information

- SAN Administration overview
- ONTAP SAN Host Configuration
- · View and manage SAN initiator groups in System Manager
- NetApp Technical Report 4017: Fibre Channel SAN Best Practices

NVMe provisioning

NVMe Overview

You can use the non-volatile memory express (NVMe) protocol to provide storage in a SAN environment. The NVMe protocol is optimized for performance with solid state storage.

For NVMe, storage targets are called namespaces. An NVMe namespace is a quantity of non-volatile storage that can be formatted into logical blocks and presented to a host as a standard block device. You create namespaces and subsystems, and then map the namespaces to the subsystems, similar to the way LUNs are provisioned and mapped to igroups for FC and iSCSI.

NVMe targets are connected to the network through a standard FC infrastructure using FC switches or a standard TCP infrastructure using Ethernet switches and host-side adapters.

Support for NVMe varies based on your version of ONTAP. See NVMe support and limitations for details.

What NVMe is

The nonvolatile memory express (NVMe) protocol is a transport protocol used for accessing nonvolatile storage media.

NVMe over Fabrics (NVMeoF) is a specification-defined extension to NVMe that enables NVMe-based communication over connections other than PCIe. This interface allows for external storage enclosures to be connected to a server.

NVMe is designed to provide efficient access to storage devices built with non-volatile memory, from flash technology to higher performing, persistent memory technologies. As such, it does not have the same limitations as storage protocols designed for hard disk drives. Flash and solid state devices (SSDs) are a type of non-volatile memory (NVM). NVM is a type of memory that keeps its content during a power outage. NVMe is a way that you can access that memory.

The benefits of NVMe include increased speeds, productivity, throughput, and capacity for data transfer. Specific characteristics include the following:
• NVMe is designed to have up to 64 thousand queues.

Each queue in turn can have up to 64 thousand concurrent commands.

- NVMe is supported by multiple hardware and software vendors
- NMVe is more productive with Flash technologies enabling faster response times
- NVMe allows for multiple data requests for each "request" sent to the SSD.

NVMe takes less time to decode a "request" and does not require thread locking in a multithreaded program.

• NVMe supports functionality that prevents bottlenecking at the CPU level and enables massive scalability as systems expand.

About NVMe namespaces

An NVMe namespace is a quantity of non-volatile memory (NVM) that can be formatted into logical blocks. Namespaces are used when a storage virtual machine is configured with the NVMe protocol and are the equivalent of LUNs for FC and iSCSI protocols.

One or more namespaces are provisioned and connected to an NVMe host. Each namespace can support various block sizes.

The NVMe protocol provides access to namespaces through multiple controllers. Using NVMe drivers, which are supported on most operating systems, solid state drive (SSD) namespaces appear as standard-block devices on which file systems and applications can be deployed without any modification.

A namespace ID (NSID) is an identifier used by a controller to provide access to a namespace. When setting the NSID for a host or host group, you also configure the accessibility to a volume by a host. A logical block can only be mapped to a single host group at a time, and a given host group does not have any duplicate NSIDs.

About NVMe subsystems

An NVMe subsystem includes one or more NVMe controllers, namespaces, NVM subsystem ports, an NVM storage medium, and an interface between the controller and the NVM storage medium. When you create an NVMe namespace, by default it is not mapped to a subsystem. You can also choose to map it a new or existing subsystem.

Related information

- · Learn to provision NVMe storage on ASA, AFF, and FAS systems
- Learn to map an NVMe namespace to a subsystem on ASA AFF and FAS systems.
- Configure SAN hosts and cloud clients
- Learn to provision SAN storage on ASA r2 (ASA A1K, ASA A90, ASA A70, ASA A50, ASA A30, or ASA A20) storage systems.

NVMe license requirements

Beginning with ONTAP 9.5 a license is required to support NVMe. If NVMe is enabled in ONTAP 9.4, a 90 day grace period is given to acquire the license after upgrading to ONTAP 9.5.

You can enable the license using the following command:

system license add -license-code NVMe_license_key

NVMe configuration, support, and limitations

Beginning with ONTAP 9.4, the non-volatile memory express (NVMe) protocol is available for SAN environments. FC-NVMe uses the same physical setup and zoning practice as traditional FC networks but allows for greater bandwidth, increased IOPs and reduced latency than FC-SCSI.

NVMe support and limitations vary based on your version of ONTAP, your platform and your configuration. For details on your specific configuration, see the NetApp Interoperability Matrix Tool. For supported limits, see Hardware Universe.



The maximum nodes per cluster is available in Hardware Universe under **Supported Platform Mixing**.

Configuration

- You can set up your NVMe configuration using a single fabric or multifabric.
- You should configure one management LIF for every SVM supporting SAN.
- The use of heterogeneous FC switch fabrics is not supported, except in the case of embedded blade switches.

Specific exceptions are listed on the NetApp Interoperability Matrix Tool.

• Cascade, partial mesh, full mesh, core-edge, and director fabrics are all industry-standard methods of connecting FC switches to a fabric, and all are supported.

A fabric can consist of one or multiple switches, and the storage controllers can be connected to multiple switches.

Features

The following NVMe features are supported based on your version of ONTAP.

Beginning with ONTAP	NVMe suppo	orts
9.17.1	 SnapMirror active sync NVMe/FC and NVMe/TCF host access for VMware workloads. 	
	i	NVMe/TCP with VMware depends on the resolution of VMware Bug ID: TR1049746.
9.15.1	 Four-nod NVMe/TC 	e MetroCluster IP configurations on CP

9.14.1	 Setting the host priority at the subsystem (host- level QoS)
9.12.1	 Four-node MetroCluster IP configurations on NVMe/FC MetroCluster configurations are not supported for front-end NVMe networks before ONTAP 9.12.1. MetroCluster configurations are not supported on NVMe/TCP.
9.10.1	Resizing a namespace
9.9.1	 Namespaces and LUNs coexistence on the same volume
9.8	 Protocol co-existence SCSI, NAS and NVMe protocols can exist on the same storage virtual machine (SVM). Prior to ONTAP 9.8, NVMe can be the only protocol on the SVM.
9.6	 512 byte blocks and 4096 byte blocks for namespaces 4096 is the default value. 512 should only be used if the host operating system does not support 4096 byte blocks. Volume move with mapped namespaces
9.5	Multipath HA pair failover/giveback

Protocols

The following NVMe protocols are supported.

Protocol	Beginning with ONTAP	Allowed by
TCP	9.10.1	Default
FC	9.4	Default

Beginning with ONTAP 9.8, you can configure SCSI, NAS and NVMe protocols on the same storage virtual machine (SVM).

In ONTAP 9.7 and earlier, NVMe can be the only protocol on the SVM.

Namespaces

When working with NVMe namespaces, you should be aware of the following:

- For ONTAP 9.15.1 and earlier, ONTAP does not support the NVMe DataSet Management (deallocate) command with NVMe for space reclamation.
- You cannot use SnapRestore to restore a namespace from a LUN or vice-versa.
- The space guarantee for namespaces is the same as the space guarantee of the containing volume.
- You cannot create a namespace on a volume transition from Data ONTAP operating in 7-Mode.
- Namespaces do not support the following:
 - Renaming
 - Inter-volume move
 - Inter-volume copy
 - · Copy on Demand

Additional limitations

The following ONTAP features are not supported by NVMe configurations:

- Virtual Storage Console
- Persistent reservations

The following applies only to nodes running ONTAP 9.4:

- NVMe LIFs and namespaces must be hosted on the same node.
- The NVMe service must be created before the NVMe LIF is created.

Related information

Best practices for modern SAN

Configure a storage VM for NVMe

If you want to use the NVMe protocol on a node, you must configure your SVM specifically for NVMe.

Before you begin

Your FC or Ethernet adapters must support NVMe. Supported adapters are listed in the NetApp Hardware Universe.

Example 3. Steps

System Manager

Configure an storage VM for NVMe with ONTAP System Manager (9.7 and later).

To configure NVMe on a new storage VM	To configure NVMe on an existing storage VM
 In System Manager, click Storage > Storage VMs and then click Add. 	 In System Manager, click Storage > Storage VMs.
2. Enter a name for the storage VM.	2. Click on the storage VM you want to configure.
 Select NVMe for the Access Protocol. Select Enable NVMe/FC or Enable NVMe/TCP 	 Click on the Settings tab, and then click next to the NVMe protocol.
and Save .	 Select Enable NVMe/FC or Enable NVMe/TCP and Save.

CLI

Configure an storage VM for NVMe with the ONTAP CLI.

1. If you do not want to use an existing SVM, create one:

vserver create -vserver <SVM name>

a. Verify that the SVM is created:

vserver show

2. Verify that you have NVMe or TCP capable adapters installed in your cluster:

For NVMe:

network fcp adapter show -data-protocols-supported fc-nvme

For TCP:

network port show

Learn more about network port show in the ONTAP command reference.

3. If you are running ONTAP 9.7 or earlier, remove all protocols from the SVM:

```
vserver remove-protocols -vserver <SVM_name> -protocols
iscsi,fcp,nfs,cifs,ndmp
```

Beginning with ONTAP 9.8, it is not necessary to remove other protocols when adding NVMe.

4. Add the NVMe protocol to the SVM:

vserver add-protocols -vserver <SVM name> -protocols nvme

5. If you are running ONTAP 9.7 or earlier, verify that NVMe is the only protocol allowed on the SVM:

vserver show -vserver <SVM name> -fields allowed-protocols

NVMe should be the only protocol displayed under the allowed protocols column.

6. Create the NVMe service:

vserver nvme create -vserver <SVM name>

7. Verify that the NVMe service was created:

vserver nvme show -vserver <SVM_name>

The Administrative Status of the SVM should be listed as up. Learn more about up in the ONTAP command reference.

- 8. Create an NVMe/FC LIF:
 - For ONTAP 9.9.1 or earlier, FC:

```
network interface create -vserver <SVM_name> -lif <lif_name>
-role data -data-protocol fc-nvme -home-node <home_node> -home
-port <home_port>
```

• For ONTAP 9.10.1 or later, FC:

```
network interface create -vserver <SVM_name> -lif <lif_name>
-service-policy <default-data-nvme-tcp | default-data-nvme-fc>
-data-protocol <fc-nvme> -home-node <home_node> -home-port
<home_port> -status-admin up -failover-policy disabled -firewall
-policy data -auto-revert false -failover-group <failover_group>
-is-dns-update-enabled false
```

• For ONTAP 9.10.1 or later, TCP:

network interface create -vserver <SVM_name> -lif <lif_name> -address <ip address> -netmask <netmask_value> -service-policy <default-data-nvme-tcp> -data-protocol <nvme-tcp> -home-node <home_node> -home-port <home_port> -status-admin up -failover -policy disabled -firewall-policy data -auto-revert false -failover-group <failover group> -is-dns-update-enabled false

- 9. Create an NVMe/FC LIF on the HA partner node:
 - For ONTAP 9.9.1 or earlier, FC:

network interface create -vserver <SVM_name> -lif <lif_name> -role data -data-protocol fc-nvme -home-node <home_node> -home -port <home_port>

• For ONTAP 9.10.1 or later, FC:

network interface create -vserver <SVM_name> -lif <lif_name> -service-policy <default-data-nvme-fc> -data-protocol <fc-nvme> -home-node <home_node> -home-port <home_port> -status-admin up -failover-policy disabled -firewall-policy data -auto-revert false -failover-group <failover_group> -is-dns-update-enabled false

• For ONTAP 9.10.1 or later, TCP:

network interface create -vserver <SVM_name> -lif <lif_name> -address <ip address> -netmask <netmask_value> -service-policy <default-data-nvme-tcp> -data-protocol <nvme-tcp> -home-node <home_node> -home-port <home_port> -status-admin up -failover -policy disabled -firewall-policy data -auto-revert false -failover-group <failover_group> -is-dns-update-enabled false

10. Verify the NVMe/FC LIFs were created:

network interface show -vserver <SVM_name>

11. Create volume on the same node as the LIF:

vol create -vserver <SVM_name> -volume <vol_name> -aggregate
<aggregate name> -size <volume size>

Provision NVMe storage

Use these steps to create namespaces and provision storage for any NVMe supported host on an existing storage VM.

About this task

This procedure applies to FAS, AFF, and ASA systems. If you have an ASA r2 system (ASA A1K, ASA A90, ASA A70, ASA A50, ASA A30, ASA A20, or ASA C30), follow these steps to provision your storage. ASA r2 systems provide a simplified ONTAP experience specific to SAN-only customers.

Beginning with ONTAP 9.8, when you provision storage, QoS is enabled by default. You can disable QoS or choose a custom QoS policy during the provisioning process or at a later time.

Before you begin

Your storage VM must be configured for NVME, and your FC or TCP transport should already be set up.

System Manager

Using ONTAP System Manager (9.7 and later), create namespaces to provide storage using the NVMe protocol.

Steps

1. In System Manager, click Storage > NVMe Namespaces and then click Add.

If you need to create a new subsystem, click More Options.

- 2. If you are running ONTAP 9.8 or later and you want to disable QoS or choose a custom QoS policy, click **More Options** and then, under **Storage and Optimization** select **Performance Service Level**.
- 3. Zone your FC switches by WWPN. Use one zone per initiator and include all target ports in each zone.
- 4. On your host, discover the new namespaces.
- 5. Initialize the namespace and format it with a file system.
- 6. Verify that your host can write and read data on the namespace.

CLI

Using the ONTAP CLI, create namespaces to provide storage using the NVMe protocol.

This procedure creates an NVMe namespace and subsystem on an existing storage VM which has already been configured for the NVMe protocol, then maps the namespace to the subsystem to allow data access from your host system.

If you need to configure the storage VM for NVMe, see Configure an SVM for NVMe.

Steps

1. Verify that the SVM is configured for NVMe:

vserver show -vserver <svm name> -fields allowed-protocols

NVMe should be displayed under the allowed-protocols column.

2. Create the NVMe namespace:



The volume you reference with the -path parameter must already exist or you will need to create one before running this command.

```
vserver nvme namespace create -vserver <svm_name> -path <path> -size
<size of namespace> -ostype <OS type>
```

3. Create the NVMe subsystem:

vserver nvme subsystem create -vserver <svm_name> -subsystem
<name of subsystem> -ostype <OS type>

The NVMe subsystem name is case sensitive. It must contain 1 to 96 characters. Special characters are allowed.

4. Verify that the subsystem was created:

vserver nvme subsystem show -vserver <svm name>

The nyme subsystem should be displayed under the Subsystem column.

- 5. Obtain the NQN from the host.
- 6. Add the host NQN to the subsystem:

vserver nvme subsystem host add -vserver <svm_name> -subsystem
<subsystem name> -host-nqn <Host NQN>

7. Map the namespace to the subsystem:

vserver nvme subsystem map add -vserver <svm_name> -subsystem
<subsystem name> -path <path>

A namespace can only be mapped to a single subsystem.

8. Verify that the namespace is mapped to the subsystem:

vserver nvme namespace show -vserver <svm name> -instance

The subsystem should be listed as the Attached subsystem.

Map an NVMe namespace to a subsystem

Mapping an NVMe namespace to a subsystem allows data access from your host. You can map an NVMe namespace to a subsystem when you provision storage or you can do it after your storage has been provisioned.

Beginning with ONTAP 9.17.1, if you are using a SnapMirror active sync configuration, you can add an SVM to a host as a proximal vserver while adding the host to an NVMe subsystem. Active-optimized paths for a namespace in an NVMe subsystem are published to a host only from the SVM that is configured as proximal vserver.

Beginning with ONTAP 9.14.1, you can prioritize resource allocation for specific hosts. By default, when a host is added to the NVMe subsystem, it is given regular priority. You can use the ONTAP command line interface (CLI) to manually change the default priority from regular to high. Hosts assigned a high priority are allocated larger I/O queue counts and queue-depths.



If you want to give a high priority to a host that was added to a subsystem in ONTAP 9.13.1 or earlier, you can change the host priority.

Before you begin

Your namespace and subsystem should already be created. If you need to create a namespace and subsystem, see Provision NVMe storage.

Map an NVMe namespace

Steps

- 1. Obtain the NQN from the host.
- 2. Add the host NQN to the subsystem:

vserver nvme subsystem host add -vserver <SVM_name> -subsystem
<subsystem name> -host-nqn <Host NQN :subsystem. subsystem name>

If you want to change the default priority of the host from regular to high, use the -priority high option. This option is available beginning with ONTAP 9.14.1. Learn more about vserver nvme subsystem host add in the ONTAP command reference.

If you want to add an SVM as a proximal-vserver to a host while adding the host to an NVMe subsystem in a SnapMirror active sync configuration, you can use the -proximal-vservers option. This option is available beginning with ONTAP 9.17.1. You can add the source or destination SVM, or both. The SVM in which you are running this command is the default.

3. Map the namespace to the subsystem:

```
vserver nvme subsystem map add -vserver <SVM_name> -subsystem
<subsystem_name> -path <path>
```

A namespace can only be mapped to a single subsystem. Learn more about vserver nvme subsystem map add in the ONTAP command reference.

4. Verify that the namespace is mapped to the subsystem:

vserver nvme namespace show -vserver <SVM_name> -instance

The subsystem should be listed as the Attached subsystem. Learn more about vserver nvme namespace show in the ONTAP command reference.

Manage LUNs

Edit LUN QoS policy group

Beginning with ONTAP 9.10.1, you can use System Manager to assign or remove Quality of Service (QoS) policies on multiple LUNs at the same time.



If the QoS policy is assigned at the volume level, it must be changed at the volume level. You can only edit the QoS policy at the LUN level if it was originally assigned at the LUN level.

Steps

- 1. In System Manager, click Storage > LUNs.
- 2. Select the LUN or LUNs you want to edit.

If you are editing more than one LUN at a time, the LUNs must belong to the same Storage Virtual Machine (SVM). If you select LUNs that do not belong to the same SVM, the option to edit the QoS Policy Group is not displayed.

3. Click More and select Edit QoS Policy Group.

Convert a LUN into a namespace

Beginning with ONTAP 9.11.1, you can use the ONTAP CLI to in-place convert an existing LUN to an NVMe namespace.

Before you begin

- Specified LUN should not have any existing maps to an igroup.
- LUN should not be in a MetroCluster configured SVM or in a SnapMirror active sync relationship.
- · LUN should not be a protocol endpoint or bound to a protocol endpoint.
- LUN should not have non-zero prefix and/or suffix stream.
- LUN should not be part of a snapshot or on the destination side of SnapMirror relationship as a read-only LUN.

Step

1. Convert a LUN to an NVMe namespace:

```
vserver nvme namespace convert-from-lun -vserver -lun-path
```

Take a LUN offline

Beginning with ONTAP 9.10.1 you can use System Manager to take LUNs offline. Prior to ONTAP 9.10.1, you must use the ONTAP CLI to take LUNs offline.

System Manager

Steps

- 1. In System Manager, click **Storage>LUNs**.
- 2. Take a single LUN or multiple LUNs offline

If you want to	Do this
Take a single LUN offline	Next to the LUN name, click and select Take Offline.
Take multiple LUNs offline	a. Select the LUNs you want to take offline.b. Click More and select Take Offline.

CLI

You can only take one LUN offline at a time when using the CLI.

Step

1. Take the LUN offline:

lun offline <lun_name> -vserver <SVM_name>

Resize a LUN in ONTAP

You can increase or decrease the size of a LUN.

About this task

This procedure applies to FAS, AFF, and ASA systems. If you have an ASA r2 system (ASA A1K, ASA A90, ASA A70, ASA A50, ASA A30, ASA A20, or ASA C30), follow these steps to increase the size of a storage unit. ASA r2 systems provide a simplified ONTAP experience specific to SAN-only customers.



Solaris LUNs cannot be resized.

Increase the size of a LUN

The size to which you can increase your LUN varies depending upon your version of ONTAP.

ONTAP version	Maximum LUN size
ONTAP 9.12.1P2 and later	128 TB for AFF, FAS, and ASA platforms
ONTAP 9.8 and later	128 TB for All-Flash SAN Array (ASA) platforms16 TB for non-ASA platforms
ONTAP 9.5, 9.6, 9.7	16TB

ONTAP 9.4 or earlier	10 times the original LUN size, but not greater than 16TB, which is the maximum LUN size.
	For example, if you create a 100 GB LUN, you can only grow it to 1,000 GB.
	The actual maximum size of the LUN might not be exactly 16TB. ONTAP rounds down the limit to be slightly less.

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

Example 4. Steps

System Manager

Increase the size of a LUN with ONTAP System Manager (9.7 and later).

- 1. In System Manager, click **Storage > LUNs**.
- 2. Click i and select Edit.
- 3. Under Storage and Optimization increase the size of the LUN and Save.

CLI

Increase the size of a LUN with the ONTAP CLI.

1. Increase the size of the LUN:

```
lun resize -vserver <SVM_name> -volume <volume_name> -lun <lun_name>
-size <lun size>
```

Learn more about lun resize in the ONTAP command reference.

2. Verify the increased LUN size:

lun show -vserver <SVM_name>



```
(\mathbf{i})
```

set -unit B

lun show -fields max-resize-size -volume volume_name -lun
lun_name

Learn more about lun show in the ONTAP command reference.

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

Decrease the size of a LUN

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

After you decrease the size of your LUN, ONTAP automatically notifies the initiator that the LUN size has decreased. However, additional steps might be required on your host for the host to recognize the new LUN

size. Check your host documentation for specific information about decreasing the size of the host file structure.

Move a LUN

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

Before you begin

If your LUN is using Selective LUN Map (SLM), you should modify the SLM reporting-nodes list to include the destination node and its HA partner before you move your LUN.

About this task

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

Data protection through snapshots occurs at the volume level. Therefore, when you move a LUN, it falls under the data protection scheme of the destination volume. If you do not have snapshots established for the destination volume, snapshots of the LUN are not created. Also, all of the snapshots of the LUN stay in the original volume until those snapshots are deleted.

You cannot move a LUN to the following volumes:

- A SnapMirror destination volume
- The SVM root volume

You cannot move the following types of LUNs:

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

When the nodes in a cluster are on different ONTAP versions, you can move a LUN between volumes on different nodes only if the source is on a later version than the destination. For example, if the source volume's node is on ONTAP 9.15.1 and the destination volume's node is on ONTAP 9.16.1, you cannot move the LUN. You can move LUNs between volumes on nodes that are on the same ONTAP version.



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

System Manager

Move a LUN with ONTAP System Manager (9.7 and later).

Beginning with ONTAP 9.10.1, you can use System Manager to create a new volume when you move a single LUN. In ONTAP 9.8 and 9.9.1, the volume to which you are moving your LUN must exist before you begin the LUN move.

Steps

- 1. In System Manager, click Storage>LUNs.
- 2. Right click the LUN you want to move, then click and select **Move LUN**.

In ONTAP 9.10.1, select to move the LUN to An existing volume or to a New volume.

If you select to create a new volume, provide the volume specifications.

3. Click Move.

CLI

Move a LUN with the ONTAP CLI.

1. Move the LUN:

lun move start

During a very brief period, the LUN is visible on both the origin and destination volume. This is expected and is resolved upon completion of the move.

2. Track the status of the move and verify successful completion:

lun move show

Related information

Selective LUN Map

Delete LUNs

You can delete a LUN from a storage virtual machine (SVM) if you no longer need the LUN.

Before you begin

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

Steps

1. Verify that the application or host is not using the LUN.

2. Unmap the LUN from the igroup:

lun mapping delete -vserver <SVM_name> -volume <volume_name> -lun
<LUN_name> -igroup <igroup_name>

3. Delete the LUN:

```
lun delete -vserver <SVM name> -volume <volume name> -lun <LUN name>
```

4. Verify that you deleted the LUN:

lun show -vserver <SVM name>

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

What to know before copying LUNs

You should be aware of certain things before copying a LUN.

Cluster administrators can copy a LUN across storage virtual machines (SVMs) within the cluster by using the lun copy command. Cluster administrators must establish the storage virtual machine (SVM) peering relationship using the vserver peer create command before an inter-SVM LUN copy operation is performed. There must be enough space in the source volume for a SIS clone.

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

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

Copying LUNs is a nondisruptive operation.

You cannot copy the following types of LUNs:

- · A LUN that has been created from a file
- · A LUN that is in NVFAIL state
- · A LUN that is in a load-sharing relationship

• A protocol-endpoint class LUN

Learn more about lun copy in the ONTAP command reference.

Examine configured and used space of a LUN

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

Step

1. View the configured space versus the actual space used for a LUN:

lun show

The following example show the configured space versus the actual space used by the LUNs in the vs3 storage virtual machine (SVM):

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

```
vserver path
                             size space-reserve size-used
                              ----- ------ ------
_____ _ ____
vs3 /vol/vol0/lun1 50.01GB disabled
vs3 /vol/vol0/lun1_backup 50.01GB disabled
                             50.01GB disabled
                                                  25.00GB
                                                 32.15GB
      /vol/vol0/lun2
                             75.00GB disabled
vs3
                                                 0B
vs3 /vol/volspace/lun0
                            5.00GB enabled
                                                4.50GB
4 entries were displayed.
```

Learn more about lun show in the ONTAP command reference.

Control and monitor I/O performance to LUNs by using Storage QoS

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

About this task

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

You can also assign storage virtual machines (SVMs) with FlexVol volumes and LUNs to policy groups.

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

• The LUN must be contained by the SVM to which the policy group belongs.

You specify the SVM when you create the policy group.

• If you assign a LUN to a policy group, then you cannot assign the LUN's containing volume or SVM to a policy group.

For more information about how to use Storage QoS, see the System administration reference.

Steps

1. Use the qos policy-group create command to create a policy group.

Learn more about qos policy-group create in the ONTAP command reference.

2. Use the lun create command or the lun modify command with the -qos-policy-group parameter to assign a LUN to a policy group.

Learn more about lun in the ONTAP command reference.

- 3. Use the qos statistics commands to view performance data.
- 4. If necessary, use the gos policy-group modify command to adjust the policy group's maximum throughput limit.

Learn more about qos policy-group modify in the ONTAP command reference.

Tools available to effectively monitor your LUNs

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

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

Capabilities and restrictions of transitioned LUNs

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

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

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

You can do the following with transitioned LUNs:

- View the LUN using the lun show command
- View the inventory of LUNs transitioned from the 7-Mode volume using the transition 7-mode show command
- Restore a volume from a 7-Mode snapshot

Restoring the volume transitions all of the LUNs captured in the snapshot

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

You cannot do the following with transitioned LUNs:

· Access snapshot-backed LUN clones captured in the volume

Related information

- Copy-based transition
- lun show

I/O misalignments on properly aligned LUNs overview

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

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



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

If you require further investigation, see the Knowledge Base article How to identify unaligned IO on LUNs?

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

- Windows Unified Host Utilities 7.1
- Provision SAN storage documentation

Achieve I/O alignment using LUN OS types

For ONTAP 9.7 or earlier, you should use the recommended ONTAP LUN ostype value that most closely matches your operating system to achieve I/O alignment with your OS partitioning scheme.

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



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

For more information on LUNs provisioned in ONTAP 9.7 or earlier, see the KB article How to identify unaligned IO on LUNs.



By default, new LUNs that are provisioned in ONTAP 9.8 or later have a prefix and suffix size of zero for all LUN OS types. The I/O should be aligned with the supported host OS by default.

Special I/O alignment considerations for Linux

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

For RHEL 5 and earlier and SLES 10 and earlier, if the LUN will be used without a volume manager, you should partition the LUN to have one partition that begins at an aligned offset, which is a sector that is an even multiple of eight logical blocks.

Special I/O alignment considerations for Solaris LUNs

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

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

ESX boot LUNs report as misaligned

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

Related information

Guest VM file system partition/disk alignment for VMware vSphere, other virtual environments, and NetApp storage systems

Ways to address issues when LUNs go offline

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

If the	You can
Aggregate is full	Add more disks.
	• Use the volume modify command to shrink a volume that has available space.
	• If you have space-guarantee volumes that have available space, change the volume space guarantee to none with the volume modify command.

If the	You can
Volume is full but there is space available in the containing aggregate	• For space guarantee volumes, use the volume modify command to increase the size of your volume.
	• For thinly provisioned volumes, use the volume modify command to increase the maximum size of your volume.
	If volume autogrow is not enabled, use volume modify -autogrow-mode to enable it.
	• Delete snapshots manually with the volume snapshot delete command, or use the volume snapshot autodelete modify command to automatically delete snapshots.

Related information

Disk and local tier (aggregate) management

Logical storage management

Troubleshoot iSCSI LUNs not visible on the host

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

Configuration setting	What to do
Cabling	Verify that the cables between the host and storage system are properly connected.
Network connectivity	 Verify that there is TCP/IP connectivity between the host and storage system. From the storage system command line, ping the host interfaces that are being used for iSCSI: ping -node node_name -destination host_ip_address_for_iSCSI From the host command line, ping the storage system interfaces that are being used for iSCSI: ping -node node_name -destination host_ip_address_for_iSCSI
System requirements	Verify that the components of your configuration are qualified. Also, verify that you have the correct host operating system (OS) service pack level, initiator version, ONTAP version, and other system requirements. The Interoperability Matrix contains the most up-to-date system requirements.

Configuration setting	What to do
Jumbo frames	If you are using jumbo frames in your configuration, verify that jumbo frames are enabled on all devices in the network path: the host Ethernet NIC, the storage system, and any switches.
iSCSI service status	Verify that the iSCSI service is licensed and started on the storage system.
Initiator login	Verify that the initiator is logged in to the storage system. If the iscsi initiator show command output shows no initiators are logged in, check the initiator configuration on the host. Also verify that the storage system is configured as a target of the initiator.
iSCSI node names (IQNs)	Verify that you are using the correct initiator node names in the igroup configuration. On the host, you can use the initiator tools and commands to display the initiator node name. The initiator node names configured in the igroup and on the host must match.
LUN mappings	 Verify that the LUNs are mapped to an igroup. On the storage system console, you can use one of the following commands: lun mapping show displays all LUNs and the igroups to which they are mapped. lun mapping show -igroup displays the LUNs mapped to a specific igroup.
iSCSI LIFs enable	Verify that the iSCSI logical interfaces are enabled.

Related information

- NetApp Interoperability Matrix Tool
- lun mapping show

Manage igroups and portsets

Ways to limit LUN access with portsets and igroups

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

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

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



You can limit access to LUN1 by using a portset. In the following example, initiator1 can access LUN1 only through LIF1. However, initiator1 cannot access LUN1 through LIF2 because LIF2 is not in portset1.



Related information

- Selective LUN Map
- Create a portset and bind to an igroup

View and manage SAN initiators and igroups

You can use System Manager to view and manage initiator groups (igroups) and initiators.

About this task

- The initiator groups identify which hosts are able to access specific LUNs on the storage system.
- After an initiator and initiator groups are created, you can also edit them or delete them.
- To manage SAN initiators groups and initiators, you can perform the following tasks:
 - View and manage SAN initiator groups
 - View and manage SAN initiators

View and manage SAN initiator groups

You can use System Manager to view a list of initiator groups (igroups). From the list, you can perform additional operations.

Steps

1. In System Manager, click Hosts > SAN Initiator Groups.

The page displays a list of initiator groups (igroups). If the list is large, you can view additional pages of the list by clicking the page numbers at the lower right corner of the page.

The columns display various information about the igroups. Beginning with 9.11.1, the connection status of the igroup is also displayed. Hover over status alerts to view details.

2. (Optional): You can perform the following tasks by clicking the icons at the upper right corner of the list:

• Search

- Download the list.
- Show or Hide columns in the list.
- Filter the data in the list.
- 3. You can perform operations from the list:

[°] Click + Add to add an igroup.

• Click the igroup name to view the **Overview** page that shows details about the igroup.

On the **Overview** page, you can view the LUNs associated with the igroup, and you can initiate the operations to create LUNs and map the LUNs. Click **All SAN Initiators** to return to the main list.

- Hover over the igroup, then click i next to an igroup name to edit or delete the igroup.
- Hover over the area to the left of the igroup name, then check the check box. If you click **+Add to Initiator Group**, you can add that igroup to another igroup.
- In the Storage VM column, click the name of a storage VM to view details about it.

View and manage SAN initiators

You can use System Manager to view a list of initiators. From the list, you can perform additional operations.

Steps

1. In System Manager, click **Hosts > SAN Initiator Groups**.

The page displays a list of initiator groups (igroups).

- 2. To view initiators, perform the following:
 - Click the **FC Initiators** tab to view a list of FC initiators.
 - $\circ\,$ Click the <code>iSCSI Initiators</code> tab to view a list of iSCSI initiators.

The columns display various information about the initiators.

Beginning with 9.11.1, the connection status of the initiator is also displayed. Hover over status alerts to view details.

- 3. (Optional): You can perform the following tasks by clicking the icons at the upper right corner of the list:
 - **Search** the list for particular initiators.
 - Download the list.
 - Show or Hide columns in the list.
 - Filter the data in the list.

Create a nested igroup

Beginning with ONTAP 9.9.1, you can create an igroup that consists of other existing igroups.

- 1. In System Manager, click **Host > SAN Initiator Groups**, and then click **Add**.
- 2. Enter the igroup **Name** and **Description**.

The description serves as the igroup alias.

3. Select the Storage VM and Host Operating System.

The OS type of a nested igroup cannot be changed after the igroup is created.

4. Under Initiator Group Members select Existing initiator group.

You can use **Search** to find and select the initiator groups you want to add.

Map igroups to multiple LUNs

÷.

Beginning with ONTAP 9.9.1, you can map igroups to two or more LUNs simultaneously.

- 1. In System Manager, click **Storage > LUNs**.
- 2. Select the LUNs you want to map.
- 3. Click More, then click Map To Initiator Groups.



The selected igroups are added to the selected LUNs. The pre-existing mappings are not overwritten.

Create a portset and bind to an igroup

In addition to using Selective LUN Map (SLM), you can create a portset and bind the portset to an igroup to further limit which LIFs can be used by an initiator to access a LUN.

If you do not bind a portset to an igroup, then all of the initiators in the igroup can access mapped LUNs through all of the LIFs on the node owning the LUN and the owning node's HA partner.

Before you begin

You must have at least one LIF and one igroup.

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

About this task

It is advantageous to use portsets with SLM when you have more than two LIFs on a node and you want to restrict a certain initiator to a subset of LIFs. Without portsets, all targets on the node will be accessible by all of the initiators with access to the LUN through the node owning the LUN and the owning node's HA partner.

System Manager

Beginning with ONTAP 9.10.1, you can use System Manager to create portsets and bind them to igroups.

If you need to create a portset and bind it to an igroup in an ONTAP release earlier than 9.10.1 you must use the ONTAP CLI procedure.

Beginning with ONTAP 9.12.1, if you do not have an existing portset, you must create the first one using the ONTAP CLI procedure.

- 1. In System Manager, click **Network > Overview > Portsets**, and click **Add**.
- 2. Enter the information for the new portset and click Add.
- 3. Click Hosts > SAN Initiator Groups.
- 4. To bind the portset to a new igroup, click **Add**.

To bind the portset to an existing igroup, select the igroup, click **i**, and then click **Edit Initiator Group**.

Related information

View and manage initiators and igroups

CLI

1. Create a port set containing the appropriate LIFs:

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

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

2. Bind the igroup to the port set:

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

Learn more about lun igroup bind in the ONTAP command reference.

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

portset show -vserver vserver_name

Manage portsets

In addition to Selective LUN Map (SLM), you can use portsets to further limit which LIFs

can be used by an initiator to access a LUN.

Beginning with ONTAP 9.10.1, you can use System Manager to change the network interfaces associated with portsets and to delete portsets.

Change network interfaces associated with a portset

- 1. In System Manager, select **Network > Overview > Portsets**.
- 2. Select the portset you want to edit then 1, then select **Edit Portset**.

Delete a portset

- 1. In System Manager, click Network > Overview > Portsets.
- 2. To delete a single portset, select the portset, select **i** and then select **Delete Portsets**.

To delete multiple portsets, select the portsets, and click **Delete**.

Selective LUN Map overview

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

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

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

SLM is enabled by default on all new LUN maps.

Determine whether SLM is enabled on a LUN map

If your environment has a combination of LUNs created in an ONTAP 9 release and LUNs transitioned from previous versions, you might need to determine whether Selective LUN Map (SLM) is enabled on a specific LUN.

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

Learn more about lun mapping show in the ONTAP command reference.

Modify the SLM reporting-nodes list

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

Steps

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

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

If you have a consistent naming convention, you can modify multiple LUN mappings at the same time by using <code>igroup_prefix*</code> instead of <code>igroup_name</code>.

- 2. Rescan the host to discover the newly added paths.
- 3. If your OS requires it, add the new paths to your multipath network I/O (MPIO) configuration.
- 4. Run the command for the needed move operation and wait for the operation to finish.
- 5. Verify that I/O is being serviced through the Active/Optimized path:

lun mapping show -fields reporting-nodes

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

```
lun mapping remove-reporting-nodes -vserver <vserver_name> -path
<lun path> -igroup <igroup name> -remote-nodes
```

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

lun mapping show -fields reporting-nodes

- 8. Remove any stale device entries for the host OS.
- 9. Change any multipathing configuration files if required.
- Rescan the host to verify removal of old paths.
 See your host documentation for specific steps to rescan your hosts.

Manage iSCSI protocol

Configure your network for best performance

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

Steps

1. Connect the host and storage ports to the same network.

It is best to connect to the same switches. Routing should never be used.

2. Select the highest speed ports available, and dedicate them to iSCSI.

10 GbE ports are best. 1 GbE ports are the minimum.

3. Disable Ethernet flow control for all ports.

You should see Network management for using the CLI to configure Ethernet port flow control.

4. Enable jumbo frames (typically MTU of 9000).

All devices in the data path, including initiators, targets, and switches, must support jumbo frames. Otherwise, enabling jumbo frames actually reduces network performance substantially.

Configure an SVM for iSCSI

To configure a storage virtual machine (SVM) for iSCSI, you must create LIFs for the SVM and assign the iSCSI protocol to those LIFs.

About this task

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

Example 7. Steps

System Manager

Configure an storage VM for iSCSI with ONTAP System Manager (9.7 and later).

To configure iSCSI on a new storage VM	To configure iSCSI on an existing storage VM
 In System Manager, click Storage > Storage VMs and then click Add. 	 In System Manager, click Storage > Storage VMs.
2. Enter a name for the storage VM.	2. Click on the storage VM you want to configure.
3. Select iSCSI for the Access Protocol .	3. Click on the Settings tab, and then click 💠
 4. Click Enable iSCSI and enter the IP address and subnet mask for the network interface. + Each node should have at least two network interfaces. 5. Click Save. 	 4. Click Enable iSCSI protocol. 4. Click Enable iSCSI and enter the IP address and subnet mask for the network interface. + Each node should have at least two network interfaces. 5. Click Save.

CLI

Configure an storage VM for iSCSI with the ONTAP CLI.

1. Enable the SVMs to listen for iSCSI traffic:

vserver iscsi create -vserver vserver name -target-alias vserver name

- 2. Create a LIF for the SVMs on each node to use for iSCSI:
 - For ONTAP 9.6 and later:

network interface create -vserver vserver_name -lif lif_name -data
-protocol iscsi -service-policy default-data-iscsi -home-node node_name
-home-port port name -address ip address -netmask netmask

• For ONTAP 9.5 and earlier:

network interface create -vserver vserver_name -lif lif_name -role data -data-protocol iscsi -home-node node_name -home-port port_name -address ip address -netmask netmask

3. Verify that you set up your LIFs correctly:

network interface show -vserver vserver name

Learn more about network interface show in the ONTAP command reference.

4. Verify that iSCSI is up and running and the target IQN for that SVM:

vserver iscsi show -vserver vserver_name

5. From your host, create iSCSI sessions to your LIFs.

Related information

• NetApp Technical Report 4080: Best Practices for Modern SAN

Define a security policy method for an initiator

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

About this task

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



Not all initiators support hexadecimal CHAP secret passwords.

Steps

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

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

2. Follow the screen commands to add the passwords.

Creates a security policy method for initiator iqn.1991-05.com.microsoft:host1 with inbound and outbound CHAP user names and passwords.

Related information

- · How iSCSI authentication works
- CHAP authentication

Delete an iSCSI service for an SVM

You can delete an iSCSI service for a storage virtual machine (SVM) if it is no longer required.

Before you begin

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

Steps

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

vserver iscsi modify -vserver vsl -status-admin down

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

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

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

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

Get more details in iSCSI session error recoveries

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

About this task

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

Beginning with ONTAP 9.4, the max-error-recovery-level option is not supported in the iscsi show and iscsi modify commands.

Steps

1. Enter advanced mode:

set -privilege advanced

2. Verify the current setting by using the iscsi show command.

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

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

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

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

Register the SVM with an iSNS server

You can use the vserver iscsi isns command to configure the storage virtual machine (SVM) to register with an iSNS server.

About this task

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

Steps

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

```
network interface create -vserver SVM_name -lif lif_name -role data -data
-protocol none -home-node home_node_name -home-port home_port -address
IP address -netmask network mask
```

Learn more about network interface create in the ONTAP command reference.

3. Create an iSCSI service on your SVM if one does not already exist:

vserver iscsi create -vserver SVM_name

4. Verify that the iSCSI service was created successfully:

iscsi show -vserver SVM name

5. Verify that a default route exists for the SVM:

network route show -vserver SVM name

6. If a default route does not exist for the SVM, create a default route:

network route create -vserver SVM_name -destination destination -gateway
gateway

Learn more about network route create in the ONTAP command reference.

7. Configure the SVM to register with the iSNS service:

vserver iscsi isns create -vserver SVM name -address IP address

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

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

8. Verify that the iSNS service is running:

vserver iscsi isns show -vserver SVM name

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

vserver iscsi isns start -vserver SVM_name

Resolve iSCSI error messages on the storage system

There are a number of common iSCSI-related error messages that you can view with the event log show command. You need to know what these messages mean and what you can do to resolve the issues they identify.
The following table contains the most common error messages, and instructions for resolving them:

Message	Explanation	What to do
ISCSI: network interface identifier disabled for use; incoming connection discarded	The iSCSI service is not enabled on the interface.	You can use the iscsi interface enable command to enable the iSCSI service on the interface. For example: iscsi interface enable -vserver vs1 -lif lif1
ISCSI: Authentication failed for initiator nodename	CHAP is not configured correctly for the specified initiator.	 You should check the CHAP settings; you cannot use the same user name and password for inbound and outbound settings on the storage system: 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.

Learn more about event log show in the ONTAP command reference.

Enable or disable automatic iSCSI LIF failover

After you upgrade to ONTAP 9.11.1 or later, you should manually enable automatic LIF failover on all iSCSI LIFs created in ONTAP 9.10.1 or earlier.

Beginning with ONTAP 9.11.1, you can enable automatic LIF failover for iSCSI LIFs on All-flash SAN Array platforms. If a storage failover occurs, the iSCSI LIF is automatically migrated from its home node or port to its HA partner node or port and then back once the failover is complete. Or, if the port for iSCSI LIF becomes unhealthy, the LIF is automatically migrated to a healthy port in its current home node and then back to its original port once the port is healthy again. The enables SAN workloads running on iSCSI to resume I/O service faster after a failover is experienced.

In ONTAP 9.11.1 and later, by default, newly created iSCSI LIFs are enabled for automatic LIF failover if one of the following conditions is true:

- There are no iSCSI LIFs on the SVM
- · All iSCSI LIFs on the SVM are enabled for automatic LIF failover

Enable automatic iSCSI LIF failover

By default, iSCSI LIFs created in ONTAP 9.10.1 and earlier are not enabled for automatic LIF failover. If there are iSCSI LIFs on the SVM that are not enabled for automatic LIF failover, your newly created LIFs will not be enabled for automatic LIF failover either. If automatic LIF failover is not enabled and there is a failover event

your iSCSI LIFs will not migrate.

Learn more about LIF failover and giveback.

Step

1. Enable automatic failover for an iSCSI LIF:

```
network interface modify -vserver <SVM_name> -lif <iscsi_lif> -failover
-policy sfo-partner-only -auto-revert true
```

To update all iSCSI LIFs on the SVM, use -lif* instead of lif.

Disable automatic iSCSI LIF failover

If you previously enabled automatic iSCSI LIF failover on iSCSI LIFs created in ONTAP 9.10.1 or earlier, you have the option to disable it.

Step

1. Disable automatic failover for an iSCSI LIF:

```
network interface modify -vserver <SVM_name> -lif <iscsi_lif> -failover
-policy disabled -auto-revert false
```

To update all iSCSI LIFs on the SVM, use -lif* instead of lif.

Related Information

- Create a LIF
- Manually migrate a LIF
- Manually revert a LIF to its home port
- Configure failover settings on a LIF

Manage FC protocol

Configure an SVM for FC

To configure a storage virtual machine (SVM) for FC, you must create LIFs for the SVM and assign the FC protocol to those LIFs.

Before you begin

You must have an FC license (included with ONTAP One) and it must be enabled. If the FC license is not enabled, the LIFs and SVMs will appear to be online but the operational status will be down. The FC service must be enabled for your LIFs and SVMs to be operational. You must use single initiator zoning for all of the FC LIFs in the SVM to host the initiators.

About this task

NetApp supports a minimum of one FC LIF per node for each SVM serving data with the FC protocol. You

must use two LIFs per node and two fabrics, with one LIF per node attached. This provides for redundancy at the node layer and the fabric.

Example 8. Steps

System Manager

Configure an storage VM for iSCSI with ONTAP System Manager (9.7 and later).

To configure FC on a new storage VM	To configure FC on an existing storage VM
 In System Manager, click Storage > Storage VMs and then click Add. 	 In System Manager, click Storage > Storage VMs.
2. Enter a name for the storage VM.	2. Click on the storage VM you want to configure.
 Select FC for the Access Protocol. Click Enable FC 	 Click on the Settings tab, and then click next to the FC protocol.
 + The FC ports are automatically assigned. 5. Click Save. 	 4. Click Enable FC and enter the IP address and subnet mask for the network interface. + The FC ports are automatically assigned.
	5. Click Save .

CLI

1. Enable FC service on the SVM:

vserver fcp create -vserver vserver_name -status-admin up

- 2. Create two LIFs for the SVMs on each node serving FC:
 - For ONTAP 9.6 and later:

```
network interface create -vserver vserver_name -lif lif_name -data
-protocol fcp -service-policy default-data-fcp -home-node node_name
-home-port port_name -address ip_address -netmask netmask -status-admin
up
```

• For ONTAP 9.5 and earlier:

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

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

network interface show -vserver vserver name lif name

Learn more about network interface show in the ONTAP command reference.

Related information

- NetApp Support
- NetApp Interoperability Matrix Tool

· Considerations for LIFs in cluster SAN environments

Delete an FC service for an SVM

You can delete an FC service for a storage virtual machine (SVM) if it is no longer required.

Before you begin

The administration status must be "down" before you can delete a FC service for an SVM. You can set the administration status to down with either the <code>vserver fcp modify</code> command or the <code>vserver fcp stop</code> command.

Steps

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

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

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

vserver fcp delete -vserver vs_1

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

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

Recommended MTU configurations for FCoE jumbo frames

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

Manage NVMe protocol

Start the NVMe service for an SVM

Before you can use the NVMe protocol on your storage virtual machine (SVM), you must start the NVMe service on the SVM.

Before you begin

NVMe must be allowed as a protocol on your system.

The following NVMe protocols are supported:

Protocol	Beginning with	Allowed by
ТСР	ONTAP 9.10.1	Default
FCP	ONTAP 9.4	Default

Steps

1. Change the privilege setting to advanced:

set -privilege advanced

2. Verify that NVMe is allowed as a protocol:

vserver nvme show

3. Create the NVMe protocol service:

vserver nvme create

4. Start the NVMe protocol service on the SVM:

vserver nvme modify -status -admin up

Delete NVMe service from an SVM

If needed, you can delete the NVMe service from your storage virtual machine (SVM).

Steps

1. Change the privilege setting to advanced:

set -privilege advanced

2. Stop the NVMe service on the SVM:

vserver nvme modify -status -admin down

3. Delete the NVMe service:

vserver nvme delete

Resize a namespace

Beginning with ONTAP 9.10.1, you can use the ONTAP CLI to increase or decrease the size of a NVMe namespace. You can use System Manager to increase the size of a NVMe namespace.

Increase the size of a namespace

System Manager

- 1. Click Storage > NVMe Namespaces.
- 2. Hoover over the namespace you want to increase, click , and then click Edit.
- 3. Under CAPACITY, change the size of the namespace.

CLI

1. Enter the following command: vserver nvme namespace modify -vserver SVM_name -path path -size new_size_of_namespace

Decrease the size of a namespace

You must use the ONTAP CLI to decrease the size of a NVMe namespace.

1. Change the privilege setting to advanced:

set -privilege advanced

2. Decrease the size of the namespace:

vserver nvme namespace modify -vserver SVM_name -path namespace_path -size new_size_of_namespace

Convert a namespace into a LUN

Beginning with ONTAP 9.11.1, you can use the ONTAP CLI to convert in-place an existing NVMe namespace to a LUN.

Before you start

- Specified NVMe namespace should not have any existing maps to a Subsystem.
- Namespace should not be part of a snapshot or on the destination side of SnapMirror relationship as a read-only namespace.
- Since NVMe namespaces are only supported with specific platforms and network cards, this feature only
 works with specific hardware.

Steps

1. Enter the following command to convert an NVMe namespace to a LUN:

lun convert-from-namespace -vserver -namespace-path

Learn more about lun convert-from-namespace in the ONTAP command reference.

Set up in-band authentication over NVMe

Beginning with ONTAP 9.12.1 you can use the ONTAP command line interface (CLI) to configure in-band (secure), bidirectional and unidirectional authentication between an NVMe host and controller over the NVME/TCP and NVMe/FC protocols using DH-HMAC-

CHAP authentication. Beginning with ONTAP 9.14.1, in-band authentication can be configured in System Manager.

To set up in-band authentication, each host or controller must be associated with a DH-HMAC-CHAP key which is a combination of the NQN of the NVMe host or controller and an authentication secret configured by the administrator. For an NVMe host or controller to authenticate its peer, it must know the key associated with the peer.

In unidirectional authentication, a secret key is configured for the host, but not the controller. In bidirectional authentication, a secret key is configured for both the host and the controller.

SHA-256 is the default hash function and 2048-bit is the default DH group.

System Manager

Beginning with ONTAP 9.14.1, you can use System Manager to configure in-band authentication while creating or updating an NVMe subsystem, creating or cloning NVMe namespaces, or adding consistency groups with new NVMe namespaces.

Steps

- 1. In System Manager, click **Hosts > NVMe Subsystem** and then click **Add**.
- 2. Add the NVMe subsystem name, and select the storage VM and host operating system.
- 3. Enter the Host NQN.
- 4. Select Use in-band authentication next to the Host NQN.
- 5. Provide the host secret and controller secret.

The DH-HMAC-CHAP key is a combination of the NQN of the NVMe host or controller and an authentication secret configured by the administrator.

6. Select the preferred hash function and DH group for each host.

If you don't select a hash function and a DH group, SHA-256 is assigned as the default hash function and 2048-bit is assigned as the default DH group.

- 7. Optionally, click Add and repeat the steps as needed to add more host.
- 8. Click Save.
- To verify that in-band authentication is enabled, click System Manager > Hosts > NVMe Subsystem > Grid > Peek view.

A transparent key icon next to the host name indicates that unidirectional mode is enabled. An opaque key next to the host name indicates bidirectional mode is enabled.

CLI

Steps

1. Add DH-HMAC-CHAP authentication to your NVMe subsystem:

```
vserver nvme subsystem host add -vserver <svm_name> -subsystem
<subsystem> -host-nqn <host_nqn> -dhchap-host-secret
<authentication_host_secret> -dhchap-controller-secret
<authentication_controller_secret> -dhchap-hash-function <sha-
256|sha-512> -dhchap-group <none|2048-bit|3072-bit|4096-bit|6144-
bit|8192-bit>
```

Learn more about vserver nvme subsystem host add in the ONTAP command reference.

2. Verify that the DH-HMAC CHAP authentication protocol is added to your host:

vserver nvme subsystem host show

```
[ -dhchap-hash-function {sha-256|sha-512} ] Authentication Hash
Function
[ -dhchap-dh-group {none|2048-bit|3072-bit|4096-bit|6144-bit|8192-
bit} ]
Authentication
Diffie-Hellman
Group
[ -dhchap-mode {none|unidirectional|bidirectional} ]
Authentication Mode
```

Learn more about vserver nvme subsystem host show in the ONTAP command reference.

3. Verify that the DH-HMAC CHAP authentication was performed during NVMe controller creation:

vserver nvme subsystem controller show

```
[ -dhchap-hash-function {sha-256|sha-512} ] Authentication Hash
Function
[ -dhchap-dh-group {none|2048-bit|3072-bit|4096-bit|6144-bit|8192-
bit} ]
Diffie-Hellman
Group
[ -dhchap-mode {none|unidirectional|bidirectional} ]
```

[-dhchap-mode {none|unidirectional|bidirectional}]

Authentication Mode

Learn more about vserver nvme subsystem controller show in the ONTAP command reference.

Disable in-band authentication over NVMe

If you have configured in-band authentication over NVMe using DH-HMAC-CHAP, you can choose to disable it at any time.

If you are reverting from ONTAP 9.12.1 or later to ONTAP 9.12.0 or earlier, you must disable in-band authentication before you revert. If in-band authentication using DH-HMAC-CHAP is not disabled, revert will fail.

Steps

1. Remove the host from the subsystem to disable DH-HMAC-CHAP authentication:

```
vserver nvme subsystem host remove -vserver <svm_name> -subsystem
<subsystem> -host-nqn <host_nqn>
```

2. Verify that the DH-HMAC-CHAP authentication protocol is removed from the host:

vserver nvme subsystem host show

3. Add the host back to the subsystem without authentication:

```
vserver nvme subsystem host add vserver <svm_name> -subsystem
<subsystem> -host-nqn <host nqn>
```

Set up TLS secure channel for NVMe/TCP

Beginning with ONTAP 9.16.1, you can configure TLS secure channel for NVMe/TCP connections. You can use System Manager or the ONTAP CLI to either add a new NVMe subsystem with TLS enabled, or enable TLS for an existing NVMe subsystem.

System Manager

Beginning with ONTAP 9.16.1, you can use System Manager to configure TLS for NVMe/TCP connections while creating or updating an NVMe subsystem, creating or cloning NVMe namespaces, or adding consistency groups with new NVMe namespaces.

Steps

- 1. In System Manager, click **Hosts > NVMe Subsystem** and then click **Add**.
- 2. Add the NVMe subsystem name, and select the storage VM and host operating system.
- 3. Enter the Host NQN.
- 4. Select Require Transport Layer Security (TLS) next to the Host NQN.
- 5. Provide the pre-shared key (PSK).
- 6. Click Save.
- To verify that TLS secure channel is enabled, select System Manager > Hosts > NVMe Subsystem
 > Grid > Peek view.

CLI

Steps

1. Add an NVMe subsystem host that supports TLS secure channel. You can provide a pre-shared key (PSK) using the tls-configured-psk argument:

vserver nvme subsystem host add -vserver <svm_name> -subsystem
<subsystem> -host-nqn <host nqn> -tls-configured-psk <key text>

2. Verify that the NVMe subsystem host is configured for TLS secure channel. You can optionally use the tls-key-type argument to only display hosts that are using that key type:

vserver nvme subsystem host show -vserver <svm_name> -subsystem
<subsystem> -host-nqn <host_nqn> -tls-key-type {none|configured}

3. Verify that the NVMe subsystem host controller is configured for TLS secure channel. You can optionally use any of the tls-key-type, tls-identity, or tls-cipher arguments to only display the controllers that have those TLS attributes:

vserver nvme subsystem controller show -vserver <svm_name>
-subsystem <subsystem> -host-nqn <host_nqn> -tls-key-type
{none|configured} -tls-identity <text> -tls-cipher
{none|TLS_AES_128_GCM_SHA256|TLS_AES_256_GCM_SHA384}

Related information

vserver nvme subsystem

Disable TLS secure channel for NVMe/TCP

Beginning with ONTAP 9.16.1, you can configure TLS secure channel for NVMe/TCP connections. If you have configured TLS secure channel for NVMe/TCP connections, you can choose to disable it at any time.

Steps

1. Remove the host from the subsystem to disable TLS secure channel:

```
vserver nvme subsystem host remove -vserver <svm_name> -subsystem
<subsystem> -host-nqn <host_nqn>
```

2. Verify that TLS secure channel is removed from the host:

vserver nvme subsystem host show

3. Add the host back to the subsystem without TLS secure channel:

```
vserver nvme subsystem host add vserver <svm_name> -subsystem
<subsystem> -host-nqn <host nqn>
```

Related information

• vserver nvme subsystem host

Change NVMe host priority

Beginning with ONTAP 9.14.1, you can configure your NVMe subsystem to prioritize resource allocation for specific hosts. By default, when a host is added to the subsystem, it is assigned a regular priority. Hosts assigned a high priority are allocated larger I/O queue counts and queue-depths.

You can use the ONTAP command line interface (CLI) to manually change the default priority from regular to high. To change the priority assigned to a host, you must remove the host from the subsystem and then add it back.

Steps

1. Verify that the host priority is set to regular:

```
vserver nvme show-host-priority
```

Learn more about vserver nvme show-host-priority in the ONTAP command reference.

2. Remove the host from the subsystem:

```
vserver nvme subsystem host remove -vserver <svm_name> -subsystem
<subsystem> -host-nqn <host_nqn>
```

Learn more about vserver nvme subsystem host remove in the ONTAP command reference.

3. Verify that the host is removed from the subsystem:

vserver nvme subsystem host show

Learn more about vserver nvme subsystem host show in the ONTAP command reference.

4. Add the host back to the subsystem with high priority:

vserver nvme subsystem host add -vserver <SVM_name> -subsystem
<subsystem_name> -host-nqn <Host_NQN_:subsystem._subsystem_name>
-priority high

Learn more about vserver nvme subsystem host add in the ONTAP command reference.

Manage automated host discovery of NVMe/TCP controllers in ONTAP

Beginning with ONTAP 9.14.1, host discovery of controllers using the NVMe/TCP protocol is automated by default in IP-based fabrics.

Enable automated host discovery of NVMe/TCP controllers

If you previously disabled automated host discovery, but your needs have changed, you can re-enable it.

Steps

1. Enter advanced privilege mode:

set -privilege advanced

2. Enable automated discovery:

```
vserver nvme modify -vserver <vserver_name> -mdns-service-discovery
-enabled true
```

3. Verify automated discovery of NVMe/TCP controllers is enabled.

vserver nvme show -fields mdns-service-discovery-enabled

Disable automated host discovery of NVMe/TCP controllers

If you do not need NVMe/TCP controllers to be automatically discovered by your host and you detect unwanted multicast traffic on your network, you should disable this functionality.

Steps

1. Enter advanced privilege mode:

```
set -privilege advanced
```

2. Disable automated discovery:

```
vserver nvme modify -vserver <vserver_name> -mdns-service-discovery
-enabled false
```

3. Verify automated discovery of NVMe/TCP controllers is disabled.

vserver nvme show -fields mdns-service-discovery-enabled

Disable NVMe host virtual machine identifier in ONTAP

Beginning with ONTAP 9.14.1, by default, ONTAP supports the ability of NVMe/FC hosts to identify virtual machines by a unique identifier and for NVMe/FC hosts to monitor virtual machine resource utilization. This enhances host-side reporting and troubleshooting.

You can use the bootarg to disable this functionality.

Step

1. Disable the virtual machine identifier:

bootargs set fct_sli_appid_off <port>, <port>

The following example disables the VMID on port 0g and port 0i.

```
bootargs set fct sli appid off 0g,0i
```

```
fct_sli_appid_off == 0g,0i
```

Manage systems with FC adapters

Manage systems with FC adapters

Commands are available to manage onboard FC adapters and FC adapter cards. These

commands can be used to configure the adapter mode, display adapter information, and change the speed.

Most storage systems have onboard FC adapters that can be configured as initiators or targets. You can also use FC adapter cards configured as initiators or targets. Initiators connect to back-end disk shelves, and possibly foreign storage arrays. Targets connect only to FC switches. Both the FC target HBA ports and the switch port speed should be set to the same value and should not be set to auto.

Related information

SAN configuration

Commands for managing FC adapters

You can use FC commands to manage FC target adapters, FC initiator adapters, and onboard FC adapters for your storage controller. The same commands are used to manage FC adapters for the FC protocol and the FC-NVMe protocol.

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

Commands for managing FC target adapters

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

Commands for managing FC initiator adapters

If you want to	Use this command
Display information for all initiators and their adapters in a node	run -node <i>node_name</i> storage show adapter

If you want to	Use this command
Display adapter configuration and status	run -node <i>node_name</i> sysconfig -v <i>adapter</i>
Verify which expansion cards are installed and whether there are any configuration errors	run -node <i>node_name</i> sysconfig -ac

Commands for managing onboard FC adapters

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

Related information

• network fcp adapter

Configure FC adapters

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

Target mode is used to connect the ports to FC initiators. Initiator mode is used to connect the ports to tape drives, tape libraries, or third-party storage with Foreign LUN Import (FLI).

The same steps are used when configuring FC adapters for the FC protocol and the FC-NVMe protocol. However, only certain FC adapters support FC-NVMe. See the NetApp Hardware Universe for a list of adapters that support the FC-NVMe protocol.

Configure FC adapters for target mode

Steps

1. Take the adapter offline:

node run -node node_name storage disable adapter adapter_name

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

2. Change the adapter from initiator to target:

system hardware unified-connect modify -t target -node node_name adapter
adapter_name

- 3. Reboot the node hosting the adapter you changed.
- 4. Verify that the target port has the correct configuration:

network fcp adapter show -node node name

Learn more about network fcp adapter show in the ONTAP command reference.

5. Bring your adapter online:

network fcp adapter modify -node node name -adapter adapter port -state up

Configure FC adapters for initiator mode

Before you begin

- LIFs on the adapter must be removed from any port sets of which they are members.
- All LIF's from every storage virtual machine (SVM) using the physical port to be modified must be migrated or destroyed before changing the personality of the physical port from target to initiator.



NVMe/FC does support initiator mode.

Steps

1. Remove all LIFs from the adapter:

network interface delete -vserver SVM name -lif LIF name, LIF name

Learn more about network interface delete in the ONTAP command reference.

2. Take your adapter offline:

```
network fcp adapter modify -node node_name -adapter adapter_port -status-admin
down
```

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

3. Change the adapter from target to initiator:

system hardware unified-connect modify -t initiator adapter port

- 4. Reboot the node hosting the adapter you changed.
- 5. Verify that the FC ports are configured in the correct state for your configuration:

system hardware unified-connect show

6. Bring the adapter back online:

node run -node node_name storage enable adapter adapter_port

View adapter settings

You can use specific commands to view information about your FC/UTA adapters.

FC target adapter

Step

1. Use the network fcp adapter show command to display adapter information: network fcp adapter show -instance -node node1 -adapter 0a

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

Learn more about network fcp adapter show in the ONTAP command reference.

Unified Target Adapter (UTA) X1143A-R6

Steps

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

Change the UTA2 port from CNA mode to FC mode

You should change the UTA2 port from Converged Network Adapter (CNA) mode to Fibre Channel (FC) mode to support the FC initiator and FC target mode. You should change the personality from CNA mode to FC mode when you need to change the physical medium that connects the port to its network.

Steps

1. Take the adapter offline:

```
network fcp adapter modify -node node_name -adapter adapter_name -status-admin
down
```

2. Change the port mode:

ucadmin modify -node node_name -adapter adapter_name -mode fcp

3. Reboot the node, and then bring the adapter online:

```
network fcp adapter modify -node node_name -adapter adapter_name -status-admin
up
```

- 4. Notify your admin or VIF manager to delete or remove the port, as applicable:
 - If the port is used as a home port of a LIF, is a member of an interface group (ifgrp), or hosts VLANs, then an admin should do the following:
 - i. Move the LIFs, remove the port from the ifgrp, or delete the VLANs, respectively.
 - ii. Manually delete the port by running the network port delete command.

If the network port delete command fails, the admin should address the errors, and then run the command again.

Learn more about network port delete in the ONTAP command reference.

• If the port is not used as the home port of a LIF, is not a member of an ifgrp, and does not host VLANs, then the VIF manager should remove the port from its records at the time of reboot.

If the VIF manager does not remove the port, then the admin must remove it manually after the reboot by using the network port delete command.

```
net-f8040-34::> network port show
   Node: net-f8040-34-01
                                         Speed (Mbps)
Health
  Port
         IPspace Broadcast Domain Link MTU Admin/Oper
Status
   _____
   . . .
   e0i
         Default Default
                                down 1500 auto/10
  eOf Default Default down 1500 auto/10 -
   . . .
  net-f8040-34::> ucadmin show
                       Current Current Pending Pending
Admin
  Node
          Adapter Mode Type Mode
                                             Type
Status
                       _____
   _____
              _____
                              _____ ____
                                             _____
_____
  net-f8040-34-01 0e cna target -
offline
  net-f8040-34-01 Of cna target -
offline
   . . .
   net-f8040-34::> network interface create -vs net-f8040-34 -lif m
-role
node-mgmt-home-node net-f8040-34-01 -home-port e0e -address 10.1.1.1
-netmask 255.255.255.0
   net-f8040-34::> network interface show -fields home-port, curr-
port
   vserver lif
                         home-port curr-port
   _____ _ ____
   Cluster net-f8040-34-01 clus1 e0a
                                 e0a
```

```
Cluster net-f8040-34-01 clus2 e0b
                                             e0b
   Cluster net-f8040-34-01 clus3 e0c
                                             e0c
   Cluster net-f8040-34-01 clus4 e0d
                                             e0d
   net-f8040-34
           cluster mgmt
                                  еOМ
                                             e0M
   net-f8040-34
                                             e0i
            m
                                  e0e
   net-f8040-34
            net-f8040-34-01 mgmt1 e0M
                                             e0M
   7 entries were displayed.
   net-f8040-34::> ucadmin modify local 0e fc
   Warning: Mode on adapter 0e and also adapter 0f will be changed
to fc.
   Do you want to continue? {y|n}: y
   Any changes will take effect after rebooting the system. Use the
"system node reboot" command to reboot.
   net-f8040-34::> reboot local
     (system node reboot)
   Warning: Are you sure you want to reboot node "net-f8040-34-01"?
    \{y \mid n\}: y
```

Learn more about network port show in the ONTAP command reference.

5. Verify that you have the correct SFP+ installed:

network fcp adapter show -instance -node -adapter

For CNA, you should use a 10Gb Ethernet SFP. For FC, you should either use an 8 Gb SFP or a 16 Gb SFP, before changing the configuration on the node.

Learn more about network fcp adapter show in the ONTAP command reference.

Related information

network interface

Change the CNA/UTA2 target adapter optical modules

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

Steps

1. Verify the current SFP+ used in the card. Then, replace the current SFP+ with the appropriate SFP+ for the preferred personality (FC or CNA).

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

network fcp adapter show -instance -node -adapter

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

Related information

- NetApp Hardware Universe
- network fcp adapter show

Supported port configurations for X1143A-R6 adapters

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

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

Related information

NetApp Hardware Universe

SAN configuration

Configure the ports

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

Steps

- 1. Configure the ports as needed for Fibre Channel (FC) or Converged Network Adapter (CNA) using the system node hardware unified-connect modify command.
- 2. Attach the appropriate cables for FC or 10 Gb Ethernet.
- 3. Verify that you have the correct SFP+ installed:

network fcp adapter show -instance -node -adapter

For CNA, you should use a 10Gb Ethernet SFP. For FC, you should either use an 8 Gb SFP or a 16 Gb SFP, based on the FC fabric being connected to.

Learn more about network fcp adapter show in the ONTAP command reference.

Prevent loss of connectivity when using the X1133A-R6 adapter

You can prevent loss of connectivity during a port failure by configuring your system with

redundant paths to separate X1133A-R6 HBAs.

The X1133A-R6 HBA is a 4-port, 16 Gb FC adapter consisting of two 2-port pairs. The X1133A-R6 adapter can be configured as target mode or initiator mode. Each 2-port pair is supported by a single ASIC (for example, Port 1 and Port 2 on ASIC 1 and Port 3 and Port 4 on ASIC 2). Both ports on a single ASIC must be configured to operate in the same mode, either target mode or initiator mode. If an error occurs with the ASIC supporting a pair, both ports in the pair go offline.

To prevent this loss of connectivity, you configure your system with redundant paths to separate X1133A-R6 HBAs, or with redundant paths to ports supported by different ASICs on the HBA.

Manage LIFs for all SAN protocols

Manage LIFs for all SAN protocols

Initiators must use Multipath I/O (MPIO) and asymmetric logical unit access(ALUA) for failover capability for clusters in a SAN environment. If a node fails, LIFs do not migrate or assume the IP addresses of the failed partner node. Instead, the MPIO software, using ALUA on the host, is responsible for selecting the appropriate paths for LUN access through LIFs.

You need to create one or more iSCSI paths from each node in an HA pair, using logical interfaces (LIFs) to allow access to LUNs that are serviced by the HA pair. You should configure one management LIF for every storage virtual machine (SVM) supporting SAN.

Direct connect or the use of Ethernet switches is supported for connectivity. You must create LIFs for both types of connectivity.

• You should configure one management LIF for every storage virtual machine (SVM) supporting SAN. You can configure two LIFs per node, one for each fabric being used with FC and to separate Ethernet networks for iSCSI.

After LIFs are created, they can be removed from port sets, moved to different nodes within a storage virtual machine (SVM), and deleted.

Related information

- Configure LIFs overview
- Create a LIF

Configure an NVMe LIF in ONTAP

Certain requirements must be met when configuring NVMe LIFs.

Before you begin

NVMe must be supported by the FC adapter on which you create the LIF. Supported adapters are listed in Hardware Universe.

About this task

Beginning with ONTAP 9.12.1 and later, you can configure two NVMe LIFs per node on a maximum of 12 nodes. In ONTAP 9.11.1 and earlier, you can configure two NVMe LIFs per node on a maximum of two nodes.

The following rules apply when creating an NVMe LIF:

- NVMe can be the only data protocol on data LIFs.
- You should configure one management LIF for every SVM that supports SAN.
- For ONTAP 9.5 and later, you must configure an NVMe LIF on the node containing the namespace and on node's HA partner.
- For ONTAP 9.4 only:
 - NVMe LIFs and namespaces must be hosted on the same node.
 - Only one NVMe data LIF can be configured per SVM.

Steps

1. Create the LIF:

```
network interface create -vserver <SVM_name> -lif <LIF_name> -role
<LIF_role> -data-protocol {fc-nvme|nvme-tcp} -home-node <home_node>
-home-port <home_port>
```



NVME/TCP is available beginning with ONTAP 9.10.1 and later.

2. Verify that the LIF was created:

network interface show -vserver <SVM_name>

After creation, NVMe/TCP LIFs listen for discovery on port 8009.

Related information

network interface

What to know before moving a SAN LIF

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

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

Using LIF movement, you can nondisruptively do the following:

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

Remove a SAN LIF from a port set

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

About this task

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

Steps

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



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

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

Learn more about lun igroup unbind in the ONTAP command reference.

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

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

Learn more about lun portset remove in the ONTAP command reference.

Move a SAN LIF

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

Before you begin

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

About this task

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

Steps

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

network interface show -vserver vserver name

Learn more about network interface show in the ONTAP command reference.

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

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

Learn more about network interface modify in the ONTAP command reference.

3. Assign the LIF a new node and port:

network interface modify -vserver vserver_name -lif LIF_name -home-node
node name -home-port port name

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

network interface modify -vserver vserver name -lif LIF name -status-admin up

Learn more about up in the ONTAP command reference.

5. Verify your changes:

network interface show -vserver vserver_name

Delete a LIF in a SAN environment

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

Before you begin

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

System Manager

Delete a LIF with ONTAP System Manager (9.7 and later).

Steps

- 1. In System Manager, click **Network > Overview**, and then select **Network Interfaces**.
- 2. Select the storage VM from which you want to delete the LIF.
- 3. Click and select **Delete**.

CLI

Delete a LIF with the ONTAP CLI.

Steps

1. Verify the name of the LIF and current port to be deleted:

network interface show -vserver vserver_name

2. Delete the LIF:

network interface delete

network interface delete -vserver vs1 -lif lif1

Learn more about network interface delete in the ONTAP command reference.

3. Verify that you deleted the LIF:

network interface show

network interface show -vserver vs1

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

Learn more about network interface show in the ONTAP command reference.

SAN LIF requirements for adding nodes to a cluster

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

- You must create LIFs on the new nodes as appropriate before you create LUNs on those new nodes.
- You must discover those LIFs from the hosts as dictated by the host stack and protocol.
- You must create LIFs on the new nodes so that the LUN and volume movements are possible without using the cluster interconnect network.

Configure iSCSI LIFs to return FQDN to host iSCSI SendTargets Discovery Operation

Beginning with ONTAP 9, iSCSI LIFs can be configured to return a Fully Qualified Domain Name (FQDN) when a host OS sends an iSCSI SendTargets Discovery Operation. Returning a FQDN is useful when there is a Network Address Translation (NAT) device between the host OS and the storage service.

About this task

IP addresses on one side of the NAT device are meaningless on the other side, but FQDNs can have meaning on both sides.



The FQDN value interoperability limit is 128 characters on all host OS.

Steps

1. Change the privilege setting to advanced:

set -privilege advanced

2. Configure iSCSI LIFs to return FQDN:

```
vserver iscsi interface modify -vserver SVM_name -lif iscsi_LIF_name
-sendtargets fqdn FQDN
```

In the following example, the iSCSI LIFs are configured to return storagehost-005.example.com as the FQDN.

```
vserver iscsi interface modify -vserver vs1 -lif vs1_iscsi1 -sendtargets-fqdn
storagehost-005.example.com
```

3. Verify that sendtargets is the FQDN:

```
vserver iscsi interface show -vserver SVM name -fields sendtargets-fqdn
```

In this example, storagehost-005.example.com is displayed in the sendtargets-fqdn output field.

Related information

Enable ONTAP space allocation for SAN protocols

ONTAP space allocation helps you to prevent your LUNs or NVMe namespaces from being taken offline if they run out of space and enables your SAN hosts to reclaim space.

ONTAP support for space allocation is based upon your SAN protocol and your version of ONTAP. Beginning with ONTAP 9.16.1, space allocation is enabled by default for iSCSI, FC, and NVMe protocols for newly created LUNs and all namespaces.

ONTAP version	Protocols	Space allocation is
9.16.1 or later	iSCSIFCNVMe	Enabled by default for newly created LUNs and all namespaces
9.15.1	• iSCSI • FC	Enabled by default for newly created LUNs
	NVMe	Not supported
9.14.1 and earlier	• iSCSI • FC	Disabled by default for newly created LUNs
	NVMe	Not supported

When space allocation is enabled:

 If a LUN or namespace runs out of space, ONTAP communicates to the host that no free space is available for write operations. As a result, the LUN or namespace remains online and read operations continue to be serviced. Depending upon the host configuration, either the host retries write operations until it succeeds or the host filesystem is placed offline. Write operations resume when additional free space becomes available to the LUN or namespace.

If space allocation is not enabled, when a LUN or namespace runs out of space, all I/O operations fail and the LUN or namespace is taken offline; the space issue must be resolved to resume normal operations. Rescanning LUN devices might also be required on the host to restore paths and devices to an operational state.

• A host can perform SCSI or NVME UNMAP (sometimes called TRIM) operations. UNMAP operations allow a host to identify blocks of data that are no longer required because they no longer contain valid data. Identification normally happens after file deletion. The storage system can then deallocate those data blocks so that the space can be consumed elsewhere. This deallocation greatly improves overall storage efficiency, especially with filesystems that have data high turnover.

Before you begin

Enabling space allocation requires a host configuration that can correctly handle space allocation errors when a write cannot be completed. Leveraging SCSI or NVME UNMAP requires a configuration that can use logical block provisioning as defined in the SCSI SBC-3 standard.

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

- Citrix XenServer 6.5 and later
- VMware ESXi 5.0 and later
- Oracle Linux 6.2 UEK kernel and later
- Red Hat Enterprise Linux 6.2 and later
- SUSE Linux Enterprise Server 11 and later
- Solaris 11.1 and later
- Windows

About this task

When you upgrade your cluster to ONTAP 9.15.1 or later, the space allocation setting for all LUNs created prior to the software upgrade remains the same after the upgrade, regardless of host type. For example, if a LUN was created in ONTAP 9.13.1 for a VMware host with space allocation disabled, space allocation on that LUN remains disabled after upgrading to ONTAP 9.15.1.

Steps

1. Enable space allocation:

```
lun modify -vserver <vserver_name> -volume <volume_name> -lun <lun_name>
-space-allocation enabled
```

2. Verify that space allocation is enabled:

```
lun show -vserver <vserver_name> -volume <volume_name> -lun <lun_name>
-fields space-allocation
```

3. Verify that space allocation is enabled on the host OS.



Some host configurations, including some versions of VMware ESXi, can automatically recognize the setting change and do not require user intervention. Other configurations might require a device rescan. Some filesystems and volume managers might require additional specific settings to enable space reclamation using SCSI UNMAP. Remounting of filesystems or a full OS reboot might be required. Consult the documentation for your specific host for guidance.

Host configuration for VMware ESXi 8.x and later NVMe hosts

If you have a VMware host running ESXi 8.x or later with the NVMe protocol, after you have enabled space allocation in ONTAP, you should perform the following steps on the hosts.

Steps

1. On your ESXi host, verify that the DSM is disabled:

```
esxcfg-advcfg -g /SCSi/NVmeUseDsmTp4040
```

The expected value is 0.

2. Enable the NVMe DSM:

esxcfg-advcfg -s 1 /Scsi/NvmeUseDsmTp4040

3. Verify that the DSM is enabled:

esxcfg-advcfg -g /SCSi/NVmeUseDsmTp4040

The expected value is 1.

Related links

Learn more about NVMe-oF host configuration for ESXi 8.x with ONTAP.

Recommended volume and file or LUN configuration combinations

Recommended volume and file or LUN configuration combinations overview

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

The following volume and LUN configuration combinations are recommended:

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

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

Space-reserved files or LUNs with thick volume provisioning

Benefits:

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

Costs and limitations:

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

Non-space-reserved files or LUNs with thin volume provisioning

Benefits:

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

Costs and restrictions:

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

Space-reserved files or LUNs with semi-thick volume provisioning

Benefits:

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

Costs and restrictions:

• Write operations can fail with this option.

You can mitigate this risk by properly balancing free space in the volume against data volatility.

- You cannot rely on retention of data protection objects such as snapshots and FlexClone files and LUNs.
- You cannot use ONTAP block-sharing storage efficiency capabilities that cannot be automatically deleted, including deduplication, compression, and ODX/Copy Offload.

Determine the correct volume and LUN configuration combination for your environment

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

About this task

You can optimize your LUN and volume configurations for maximum storage utilization or for the security of write guarantees. Based on your requirements for storage utilization and your ability to monitor and replenish free space quickly, you must determine the FlexVol volume and LUN volumes appropriate for your installation.



You do not need a separate volume for each LUN.

Step

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



Calculate rate of data growth for LUNs

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

About this task

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

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

Steps

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

Make sure that your monitoring period is long enough to form a representative sample of regularly occurring increases in data growth. For instance, you might consistently have a large amount of data growth at the end of each month.

- 3. Each day, record in GB how much your data grows.
- 4. At the end of your monitoring period, add the totals for each day together, and then divide by the number of days in your monitoring period.

This calculation yields your average rate of growth.

Example

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

- Sunday: 20 GB
- Monday: 18 GB
- Tuesday: 17 GB
- Wednesday: 20 GB
- Thursday: 20 GB
- Friday: 23 GB
- Saturday: 22 GB

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

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

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

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

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

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

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

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

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

File or LUN setting	Value
Space reservation	Disabled

Additional considerations

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

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

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

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

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

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

Volume setting	Value
Autogrow	Optional; if enabled, aggregate free space must be actively monitored.
File or LUN setting	Value
Space reservation	Enabled

Technology restrictions

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

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

Additional considerations

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

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

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

SAN data protection

Learn about ONTAP data protection methods for SAN environments

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

SnapMirror active sync

Beginning with general availability in ONTAP 9.9.1, provides Zero Recovery Time Objective (Zero RTO) or Transparent Application Failover (TAF) to enable automatic failover of business-critical applications in SAN environments. SnapMirror active sync requires the installation of ONTAP Mediator 1.2 in a configuration with either two AFF clusters or two All-Flash SAN Array (ASA) clusters.

SnapMirror active sync

Snapshot

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

FlexClone LUNs (FlexClone license required)

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

SnapRestore (license required)

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

Data protection mirror copies (SnapMirror license required)

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

SnapVault backups (SnapMirror license required)

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

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

SnapDrive for Windows or UNIX (SnapDrive license required)

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

Native tape backup and recovery

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

Related information

NetApp Documentation: SnapDrive for UNIX NetApp Documentation: SnapDrive for Windows (current releases) Data protection using tape backup

Restore a single LUN from an ONTAP snapshot

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

Before you begin

- You must have enough space on your volume to complete the restore operation:
 - If you are restoring a space-reserved LUN where the fractional reserve is 0%, you require one times the size of the restored LUN.
 - If you are restoring a space-reserved LUN where the fractional reserve is 100%, you require two times the size of the restored LUN.
 - If you are restoring a non-space-reserved LUN, you only require the actual space used for the restored LUN.
- A snapshot of the destination LUN must have been created.

If the restore operation fails, the destination LUN might be truncated. In such cases, you can use the snapshot to prevent data loss.

• A snapshot of the source LUN must have been created.

In rare cases, the LUN restore can fail, leaving the source LUN unusable. If this occurs, you can use the snapshot to return the LUN to the state just before the restore attempt.

• The destination LUN and source LUN must have the same OS type.

If your destination LUN has a different OS type from your source LUN, your host can lose data access to the destination LUN after the restore operation.

Steps

- 1. From the host, stop all host access to the LUN.
- 2. Unmount the LUN on its host so that the host cannot access the LUN.
- 3. Unmap the LUN:

```
lun mapping delete -vserver <SVM_name> -volume <volume_name> -lun
<lun name> -igroup <igroup_name>
```

4. Determine the snapshot you want to restore your LUN to:

volume snapshot show -vserver <SVM_name> -volume <volume_name>

5. Create a snapshot of the LUN prior to restoring the LUN:

```
volume snapshot create -vserver <SVM_name> -volume <volume_name>
-snapshot <snapshot name>
```

6. Restore the specified LUN in a volume:

```
volume snapshot restore-file -vserver <SVM_name> -volume <volume_name>
-snapshot <snapshot name> -path <lun path>
```

- 7. Follow the steps on the screen.
- 8. If necessary, bring the LUN online:

lun modify -vserver <SVM_name> -path <lun_path> -state online

9. If necessary, remap the LUN:

```
lun mapping create -vserver <SVM_name> -volume <volume_name> -lun
<lun name> -igroup <igroup name>
```

- 10. From the host, remount the LUN.
- 11. From the host, restart access to the LUN.

Restore all LUNs in a volume from an ONTAP snapshot

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

Steps

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

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

- 2. Unmount the LUNs on that host so that the host cannot access the LUNs.
- 3. Unmap your LUNs:

```
lun mapping delete -vserver <SVM_name> -volume <volume_name> -lun
<lun name> -igroup <igroup name>
```

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

volume snapshot show -vserver <SVM name> -volume <volume name>

5. Change your privilege setting to advanced:

set -privilege advanced

6. Restore your data:

```
volume snapshot restore -vserver <SVM_name> -volume <volume_name>
-snapshot <snapshot_name>
```

- 7. Follow the instructions on the screen.
- 8. Remap your LUNs:

```
lun mapping create -vserver <SVM_name> -volume <volume_name> -lun
<lun name> -igroup <igroup name>
```

9. Verify that your LUNs are online:

```
lun show -vserver <SVM name> -path <lun path> -fields state
```

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

```
lun modify -vserver <SVM name> -path <lun path> -state online
```

11. Change your privilege setting to admin:

```
set -privilege admin
```

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

Protect your data with ONTAP FlexClone LUNs

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

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

Reasons to create FlexClone LUNs

- 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.

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

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

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

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



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

Related information

- Create a FlexClone LUN
- Configure a FlexVol volume to automatically delete FlexClone LUNs
- Prevent a FlexClone LUN from being automatically deleted

Configure and use SnapVault backups in a SAN environment

Learn about ONTAP SnapVault backups in a SAN environment

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

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

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

It is important to ensure that LUNs being backed up are in a consistent state before the snapshots are created and copied to the SnapVault secondary volume. Automating the snapshot creation with SnapCenter ensures that backed up LUNs are complete and usable by the original application.

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

• You can map a LUN directly from the SnapVault secondary volume and connect a host to the LUN to access the contents of the LUN.

The LUN is read-only and you can map only from the most recent snapshot in the SnapVault backup. Persistent reservations and other LUN metadata are lost. If desired, you can use a copy program on the host to copy the LUN contents back to the original LUN if it is still accessible.

The LUN has a different serial number from the source LUN.

• You can clone any snapshot in the SnapVault secondary volume to a new read-write volume.

You can then map any of the LUNs in the volume and connect a host to the LUN to access the contents of the LUN. If desired, you can use a copy program on the host to copy the LUN contents back to the original LUN if it is still accessible.

• You can restore the entire volume containing the LUN from any snapshot in the SnapVault secondary volume.

Restoring the entire volume replaces all of the LUNs, and any files, in the volume. Any new LUNs created since the snapshot was created are lost.

The LUNs retain their mapping, serial numbers, UUIDs, and persistent reservations.

Access a read-only LUN copy from an ONTAP SnapVault backup

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

Before you begin

- The SnapVault relationship must be initialized and the latest snapshot in the SnapVault secondary volume must contain the desired LUN.
- The storage virtual machine (SVM) containing the SnapVault backup must have one or more LIFs with the desired SAN protocol accessible from the host used to access the LUN copy.
- If you plan to access LUN copies directly from the SnapVault secondary volume, you must create your igroups on the SnapVault SVM in advance.

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

About this task

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

Steps

1. List the available LUNs in the SnapVault secondary volume.

```
lun show
```

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

```
cluster::> lun show

Vserver Path State Mapped Type Size

------- Vvol/srcvolA/lun_A online mapped windows 300.0GB

vserverA /vol/srcvolA/lun_C online mapped windows 300.0GB

vserverB /vol/dstvolB/lun_A online unmapped windows 300.0GB

vserverB /vol/dstvolB/lun_B online online unmapped windows 300.0GB

vserverB /vol/dstvolB/lun_C online online unmapped windows 300.0GB
```

Learn more about lun show in the ONTAP command reference.

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

```
igroup create -vserver <SVM_name> -igroup <igroup_name> -protocol
<protocol> -ostype <ostype> -initiator <initiator name>
```

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

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

3. Map the desired LUN copy to the igroup.

```
lun mapping create -vserver <SVM_name> -path <LUN_path> -igroup
<igroup_name>
```

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

Learn more about lun mapping create in the ONTAP command reference.

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

Restore a single LUN from an ONTAP SnapVault backup

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

Before you begin

- The SnapVault relationship must be initialized and the SnapVault secondary volume must contain an appropriate snapshot to restore.
- The storage virtual machine (SVM) containing the SnapVault secondary volume must have one or more LIFs with the desired SAN protocol that are accessible from the host used to access the LUN copy.
- The igroups must already exist on the SnapVault SVM.

About this task

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

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

Steps

1. Verify the secondary volume that contains the SnapVault backup.

```
snapmirror show
cluster::> snapmirror show
                 Mirror Relation Total
Source
          Dest
                                           Last
                       Status Progress Healthy Updated
Path
      Type Path
                 State
vserverA:srcvolA
      XDP vserverB:dstvolB
                 Snapmirrored
                       Idle
                                      true
```

2. Identify the snapshot that you want to restore the LUN from.

volume snapshot show

3. Create a read-write clone from the desired snapshot

```
volume clone create -vserver <SVM_name> -flexclone <flexclone_name>
-type <type> -parent-volume <parent_volume_name> -parent-snapshot
<snapshot_name>
```

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

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

4. List the LUNs in the volume clone.

lun show -vserver <SVM name> -volume <flexclone volume name>

```
cluster::> lun show -vserver vserverB -volume dstvolB_clone
Vserver Path State Mapped Type
vserverB /vol/dstvolB_clone/lun_A online unmapped windows
vserverB /vol/dstvolB_clone/lun_B online unmapped windows
vserverB /vol/dstvolB_clone/lun_C online unmapped windows
3 entries were displayed.
```

Learn more about lun show in the ONTAP command reference.

5. If the igroup for the desired host does not already exist on the SVM containing the SnapVault backup, create an igroup.

```
igroup create -vserver <SVM_name> -igroup <igroup_name> -protocol
<protocol> -ostype <os_type> -initiator <initiator_name>
```

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

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

6. Map the desired LUN copy to the igroup.

```
lun mapping create -vserver <SVM_name> -path <lun_path> -igroup
<igroup name>
```

cluster::> lun mapping create -vserver vserverB
-path /vol/dstvolB clone/lun C -igroup temp igroup

Learn more about lun mapping create in the ONTAP command reference.

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

The LUN is read-write and can be used in place of the original LUN. Because the LUN serial number is different, the host interprets it as a different LUN from the original.

8. Use a copy program on the host to copy the LUN contents back to the original LUN.

Related information

snapmirror show

Restore all LUNs in a volume from an ONTAP SnapVault backup

If one or more LUNs in a volume need to be restored from a SnapVault backup, you can restore the entire volume. Restoring the volume affects all LUNs in the volume.

Before you begin

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

About this task

Restoring an entire volume returns the volume to the state it was in when the snapshot was made. If a LUN was added to the volume after the snapshot, that LUN is removed during the restore process.

After restoring the volume, the LUNs remain mapped to the igroups they were mapped to just before the restore. The LUN mapping might be different from the mapping at the time of the snapshot. Persistent reservations on the LUNs from host clusters are retained.

Steps

- 1. Stop I/O to all LUNs in the volume.
- 2. Verify the secondary volume that contains the SnapVault secondary volume.

3. Identify the snapshot that you want to restore from.

volume snapshot show

4. Specify the snapshot to use.

```
snapmirror restore -destination-path <destination_path> -source-path
<source_path> -source-snapshot <snapshot_name>
```

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

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

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

Restoring a volume from a SnapVault backup

In the following example, the LUN named lun_D was added to the volume after the snapshot was created. After restoring the entire volume from the snapshot, lun_D no longer appears.

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

```
cluster::> lun show
Vserver Path
                          State
                                  Mapped Type Size
_____
                           ----- ----- ------
                                                      _____
vserverA /vol/srcvolA/lun A online mapped windows 300.0GB
vserverA /vol/srcvolA/lun B online mapped windows 300.0GB
vserverA /vol/srcvolA/lun_C online mapped windows 300.0GB
vserverA /vol/srcvolA/lun D online mapped windows 250.0GB
vserverB /vol/dstvolB/lun A online unmapped windows 300.0GB
vserverB /vol/dstvolB/lun B online unmapped windows 300.0GB
vserverB /vol/dstvolB/lun C online unmapped windows 300.0GB
7 entries were displayed.
cluster::>snapmirror restore -destination-path vserverA:srcvolA
 -source-path vserverB:dstvolB
 -source-snapshot daily.2013-02-10 0010
Warning: All data newer than snapshot hourly.2013-02-11 1205
on volume vserverA:src volA will be deleted.
Do you want to continue? \{y|n\}: y
[Job 98] Job is queued: snapmirror restore from source
"vserverB:dstvolB" for the snapshot daily.2013-02-10 0010.
cluster::> lun show
Vserver Path
                          State Mapped Type
                                                      Size
_____ ____
                                                      _____
vserverA /vol/srcvolA/lun A online mapped windows 300.0GB
vserverA /vol/srcvolA/lun B online mapped windows 300.0GB
vserverA /vol/srcvolA/lun C online mapped windows 300.0GB
vserverB /vol/dstvolB/lun A online unmapped windows 300.0GB
vserverB /vol/dstvolB/lun B online unmapped windows 300.0GB
vserverB /vol/dstvolB/lun C online unmapped windows 300.0GB
6 entries were displayed.
```

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

Related information

- snapmirror restore
- snapmirror show

Recommended configuration to connect a host backup system to ONTAP

You can back up SAN systems to tape through a separate backup host to avoid

performance degradation on the application host.

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

Use a host backup system to protect a LUN on your ONTAP storage system

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

Before you begin

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

Steps

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

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

2. Create a snapshot of the production LUN.

volume snapshot create -vserver <SVM_name> -volume <volume_name>
-snapshot <snapshot> -comment <comment> -foreground false

3. Create a clone of the production LUN.

```
volume file clone create -vserver <SMV_name> -volume <volume> -source
-path <path> -snapshot-name <snapshot> -destination-path
<destination_path>
```

4. Create an igroup that includes the WWPN of the backup server.

```
lun igroup create -vserver <SVM_name> -igroup <igroup> -protocol
<protocol> -ostype <os_type> -initiator <initiator>
```

5. Map the LUN clone you created in Step 3 to the backup host.

```
lun mapping create -vserver <SVM_name> -volume <volume_name> -lun
<lun name> -igroup <igroup>
```

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

6. From the host, discover the new LUN and make the file system available to the host.

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

- 7. Back up the data in the LUN clone from the backup host to tape by using your SAN backup application.
- 8. Take the LUN clone offline.

lun modify -vserver <SVM name> -path <path> -state offline

9. Remove the LUN clone.

```
lun delete -vserver <SVM name> -volume <volume> -lun <lun name>
```

10. Remove the snapshot.

```
volume snapshot delete -vserver <SVM_name> -volume <volume> -snapshot
<snapshot>
```

SAN configuration reference

Learn about ONTAP SAN configuration

A storage area network (SAN) consists of a storage solution connected to hosts over a SAN transport protocol such as iSCSI or FC. You can configure your SAN so that your storage solution attaches to your hosts through one or more switches. If you are using iSCSI, you can also configure your SAN so that your storage solution attaches directly to your host without using a switch.

In a SAN, multiple hosts, using different operating systems, such as Windows, Linux, or UNIX, can access the storage solution at the same time. You can use <u>Selective LUN mapping</u> and <u>portsets</u> to limit data access between the hosts and the storage.

For iSCSI, the network topology between the storage solution and the hosts is referred to as a network. For FC, FC/NVMe and FCoE the network topology between the storage solution and the hosts is referred to as a fabric. To create redundancy, which protects you against loss of data access, you should set up your SAN with HA pairs in a multi-network or multi-fabric configuration. Configurations using single nodes or single networks/fabrics are not fully redundant so are not recommended.

After your SAN is configured, you can provision storage for iSCSI or FC, or you can provision storage for FC/NVMe. Then you can connect to your hosts to begin servicing data.

SAN protocol support varies based on your version of ONTAP, your platform and your configuration. For details on your specific configuration, see the NetApp Interoperability Matrix Tool.

Related information

- SAN administration overview
- NVMe configuration, support and limitations

iSCSI configurations

Configure iSCSI networks with ONTAP systems

You should set up your iSCSI configuration with high-availability (HA) pairs that attach directly to your iSCSI SAN hosts or that connect to your hosts through one or more IP switches.

HA pairs are defined as the reporting nodes for the Active/Optimized and the Active/Unoptimized paths that will be used by the hosts to access the LUNs. Multiple hosts, using different operating systems, such as Windows, Linux, or UNIX, can access the storage at the same time. Hosts require that a supported multipathing solution that supports ALUA be installed and configured. Supported operating systems and multipathing solutions can be verified on the NetApp Interoperability Matrix Tool.

In a multi-network configuration, there are two or more switches connecting the hosts to the storage system. Multi-network configurations are recommended because they are fully redundant. In a single-network configuration, there is one switch connecting the hosts to the storage system. Single-network configurations are not fully redundant.



Single-node configurations are not recommended because they do not provide the redundancy needed to support fault tolerance and nondisruptive operations.

Related information

- Learn how Selective LUN mapping (SLM) limits the paths that are used to access the LUNs owned by an HA pair.
- Learn about SAN LIFs.
- Learn about the benefits of VLANs in iSCSI.

Multi-network iSCSI configurations

In multi-network HA pair configurations, two or more switches connect the HA pair to one or more hosts. Because there are multiple switches, this configuration is fully redundant.



Single-network iSCSI configurations

In single-network HA pair configurations, one switch connects the HA pair to one or more hosts. Because there is a single switch, this configuration is not fully redundant.



Direct-attachment iSCSI configuration

In a direct-attached configuration, one or more hosts are directly connected to the controllers.



Benefits of using VLANs with ONTAP systems in iSCSI configurations

A VLAN consists of a group of switch ports grouped together into a broadcast domain. A VLAN can be on a single switch or it can span multiple switch chassis. Static and dynamic VLANs enable you to increase security, isolate problems, and limit available paths within your IP network infrastructure.

When you implement VLANs in large IP network infrastructures, you derive the following benefits:

• Increased security.

VLANs enable you to leverage existing infrastructure while still providing enhanced security because they limit access between different nodes of an Ethernet network or an IP SAN.

- Improved Ethernet network and IP SAN reliability by isolating problems.
- Reduction of problem resolution time by limiting the problem space.
- Reduction of the number of available paths to a particular iSCSI target port.
- Reduction of the maximum number of paths used by a host.

Having too many paths slows reconnect times. If a host does not have a multipathing solution, you can use VLANs to allow only one path.

Dynamic VLANs

Dynamic VLANs are MAC address-based. You can define a VLAN by specifying the MAC address of the members you want to include.

Dynamic VLANs provide flexibility and do not require mapping to the physical ports where the device is physically connected to the switch. You can move a cable from one port to another without reconfiguring the VLAN.

Static VLANs

Static VLANs are port-based. The switch and switch port are used to define the VLAN and its members.

Static VLANs offer improved security because it is not possible to breach VLANs using media access control (MAC) spoofing. However, if someone has physical access to the switch, replacing a cable and reconfiguring the network address can allow access.

In some environments, it is easier to create and manage static VLANs than dynamic VLANs. This is because static VLANs require only the switch and port identifier to be specified, instead of the 48-bit MAC address. In addition, you can label switch port ranges with the VLAN identifier.

FC configurations

Configure FC or FC-NVME fabrics with ONTAP systems

It is recommended that you configure your FC and FC-NVMe SAN hosts using HA pairs and a minimum of two switches. This provides redundancy at the fabric and storage system layers to support fault tolerance and nondisruptive operations. You cannot directly attach FC or FC-NVMe SAN hosts to HA pairs without using a switch.

Cascade, partial mesh, full mesh, core-edge, and director fabrics are all industry-standard methods of connecting FC switches to a fabric, and all are supported. The use of heterogeneous FC switch fabrics is not supported, except in the case of embedded blade switches. Specific exceptions are listed on the Interoperability Matrix Tool. A fabric can consist of one or multiple switches, and the storage controllers can be connected to multiple switches.

Multiple hosts, using different operating systems, such as Windows, Linux, or UNIX, can access the storage controllers at the same time. Hosts require that a supported multipathing solution be installed and configured. Supported operating systems and multipathing solutions can be verified on the Interoperability Matrix Tool.

Multifabric FC and FC-NVMe configurations

In multifabric HA pair configurations, there are two or more switches connecting HA pairs to one or more hosts. For simplicity, the following multifabric HA pair figure shows only two fabrics, but you can have two or more fabrics in any multifabric configuration.

The FC target port numbers (0c, 0d, 1a, 1b) in the illustrations are examples. The actual port numbers vary depending on the model of your storage node and whether you are using expansion adapters.



Single-fabric FC and FC-NVMe configurations

In single-fabric HA pair configurations, there is one fabric connecting both controllers in the HA pair to one or more hosts. Because the hosts and controllers are connected through a single switch, single-fabric HA pair configurations are not fully redundant.

The FC target port numbers (0a, 0c) in the illustrations are examples. The actual port numbers vary depending on the model of your storage node and whether you are using expansion adapters.

All platforms that support FC configurations support single-fabric HA pair configurations.



 (\mathbf{i})

Single-node configurations are not recommended because they do not provide the redundancy needed to support fault tolerance and nondisruptive operations.

Related information

• Learn how Selective LUN mapping (SLM) limits the paths that are used to access the LUNs owned by an HA pair.

• Learn about SAN LIFs.

Best practices to configure FC switches with ONTAP systems

For best performance, you should consider certain best practices when configuring your FC switch.

A fixed link speed setting is the best practice for FC switch configurations, especially for large fabrics because it provides the best performance for fabric rebuilds and can significantly save time. Although autonegotiation provides the greatest flexibility, FC switch configuration does not always perform as expected, and it adds time to the overall fabric-build sequence.

All of the switches that are connected to the fabric must support N_Port ID virtualization (NPIV) and must have NPIV enabled. ONTAP uses NPIV to present FC targets to a fabric.

For details about which environments are supported, see the NetApp Interoperability Matrix Tool.

For FC and iSCSI best practices, see NetApp Technical Report 4080: Best Practices for Modern SAN.

Recommended FC target port configuration and speeds for ONTAP systems

FC target ports can be configured and used for the FC-NVMe protocol in the exact same way they are configured and used for the FC protocol. Support for the FC-NVMe protocol varies based upon your platform and your ONTAP version. Use NetApp Hardware Universe to verify support.

For best performance and highest availability, you should use the recommended target port configuration listed in NetApp Hardware Universe for your specific platform.

Configuration for FC target ports with shared ASICs

The following platforms have port pairs with shared application-specific integrated circuits (ASICs). If you use an expansion adapter with these platforms, you should configure your FC ports so that they do not use the same ASIC for connectivity.

Controller	Port pairs with shared ASIC	Number of target ports: Recommended ports
FAS8200AFF A300	0g+0h	1: 0g 2: 0g, 0h
FAS2720FAS2750AFF A220	0c+0d 0e+0f	1: 0c 2: 0c, 0e 3: 0c, 0e, 0d 4: 0c, 0e, 0d, 0f

FC target port supported speeds

FC target ports can be configured to run at different speeds. All target ports used by a given host should be set to the same speed. You should set the target port speed to match the speed of the device to which it connects. Do not use autonegotiation for your port speed. A port that is set to autonegotiation can take longer to reconnect after a takeover/giveback or other interruption.

You can configure onboard ports and expansion adapters to run at the following speeds. Each controller and expansion adapter port can be configured individually for different speeds as needed.

4 Gb ports	8 Gb ports	16 Gb ports	32 Gb ports
• 4 Gb	• 8 Gb	• 16 Gb	• 32 Gb
• 2 Gb	• 4 Gb	• 8 Gb	• 16 Gb
• 1 Gb	• 2 Gb	• 4 Gb	• 8 Gb

For a full list of supported adapters and their supported speeds, see the NetApp Hardware Universe.

Configure ONTAP FC adapter ports

Onboard FC adapters and some FC expansion adapter cards can be individually configured as either initiators or targets ports. Other FC expansion adapters are configured as initiators or targets at the factory and cannot be changed. Additional FC ports are also available through supported UTA2 cards configured with FC SFP+ adapters.

Initiator ports can be used to connect directly to back-end disk shelves, and possibly foreign storage arrays. Target ports can be used to connect only to FC switches.

The number of onboard ports and CNA/UTA2 ports configured for FC varies depending on the model of the controller. The supported target expansion adapters also varies depending on controller model. See NetApp Hardware Universe for a complete list of onboard FC ports and supported target expansion adapters for your controller model.

Configure FC adapters for initiator mode

Initiator mode is used to connect the ports to tape drives, tape libraries, or third-party storage with Foreign LUN Import (FLI).

Before you begin

- LIFs on the adapter must be removed from any port sets of which they are members.
- All LIF's from every storage virtual machine (SVM) using the physical port to be modified must be migrated or destroyed before changing the personality of the physical port from target to initiator.



NVMe/FC does support initiator mode.

Steps

1. Remove all LIFs from the adapter:

network interface delete -vserver _SVM_name_ -lif _lif_name_,_lif_name_

2. Take your adapter offline:

```
network fcp adapter modify -node _node_name_ -adapter _adapter_port_
-status-admin down
```

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

3. Change the adapter from target to initiator:

```
system hardware unified-connect modify -t initiator adapter port
```

- 4. Reboot the node hosting the adapter you changed.
- 5. Verify that the FC ports are configured in the correct state for your configuration:

system hardware unified-connect show

6. Bring the adapter back online:

```
node run -node node name storage enable adapter adapter port
```

Configure FC adapters for target mode

Target mode is used to connect the ports to FC initiators.

The same steps are used to configure FC adapters for the FC protocol and the FC-NVMe protocol. However, only certain FC adapters support FC-NVMe. See the NetApp Hardware Universe for a list of adapters that support the FC-NVMe protocol.

Steps

1. Take the adapter offline:

```
node run -node node name storage disable adapter adapter name
```

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

2. Change the adapter from initiator to target:

```
system node hardware unified-connect modify -t target -node _node_name_
adapter adapter name
```

- 3. Reboot the node hosting the adapter you changed.
- 4. Verify that the target port has the correct configuration:

5. Bring your adapter online:

```
network fcp adapter modify -node _node_name_ -adapter _adapter_port_
-state up
```

Configure FC adapter speed

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

About this task

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

Before you begin

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

Steps

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

```
network interface modify -vserver * -lif * { -home-node node1 -home-port
0c } -status-admin down
```

2. Take the adapter offline:

network fcp adapter modify -node nodel -adapter 0c -state down

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

3. Determine the maximum speed for the port adapter:

fcp adapter show -instance

You cannot modify the adapter speed beyond the maximum speed.

4. Change the adapter speed:

network fcp adapter modify -node nodel -adapter 0c -speed 16

5. Bring the adapter online:

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

6. Bring all of the LIFs on the adapter online:

```
network interface modify -vserver * -lif * { -home-node node1 -home-port
0c } -status-admin up
```

ONTAP commands for managing FC adapters

You can use FC commands to manage FC target adapters, FC initiator adapters, and onboard FC adapters for your storage controller. The same commands are used to manage FC adapters for the FC protocol and the FC-NVMe protocol.

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

Commands for managing FC target adapters

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

Commands for managing FC initiator adapters

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

Commands for managing onboard FC adapters

If you want to	Use this command
Display the status of the onboard FC ports	system node hardware unified-connect show

Related information

• network fcp adapter

Avoid connectivity loss to an ONTAP system using an X1133A-R6 adapter

You can prevent loss of connectivity during a port failure by configuring your system with redundant paths to separate X1133A-R6 HBAs.

The X1133A-R6 HBA is a 4-port, 16 Gb FC adapter consisting of two 2-port pairs. The X1133A-R6 adapter can be configured as target mode or initiator mode. Each 2-port pair is supported by a single ASIC (for example, Port 1 and Port 2 on ASIC 1 and Port 3 and Port 4 on ASIC 2). Both ports on a single ASIC must be configured to operate in the same mode, either target mode or initiator mode. If an error occurs with the ASIC supporting a pair, both ports in the pair go offline.

To prevent this loss of connectivity, you configure your system with redundant paths to separate X1133A-R6 HBAs, or with redundant paths to ports supported by different ASICs on the HBA.

FCoE configurations

Configure FCoE fabrics with ONTAP systems

FCoE can be configured in various ways using FCoE switches. Direct-attached configurations are not supported in FCoE.

All FCoE configurations are dual-fabric, fully redundant, and require host-side multipathing software. In all FCoE configurations, you can have multiple FCoE and FC switches in the path between the initiator and target, up to the maximum hop count limit. To connect switches to each other, the switches must run a firmware version that supports Ethernet ISLs. Each host in any FCoE configuration can be configured with a different operating system.

FCoE configurations require Ethernet switches that explicitly support FCoE features. FCoE configurations are validated through the same interoperability and quality assurance process as FC switches. Supported

configurations are listed in the Interoperability Matrix. Some of the parameters included in these supported configurations are the switch model, the number of switches that can be deployed in a single fabric, and the supported switch firmware version.

The FC target expansion adapter port numbers in the illustrations are examples. The actual port numbers might vary, depending on the expansion slots in which the FCoE target expansion adapters are installed.

FCoE initiator to FC target

Using FCoE initiators (CNAs), you can connect hosts to both controllers in an HA pair through FCoE switches to FC target ports. The FCoE switch must also have FC ports. The host FCoE initiator always connects to the FCoE switch. The FCoE switch can connect directly to the FC target or can connect to the FC target through FC switches.

The following illustration shows host CNAs connecting to an FCoE switch, and then to an FC switch before connecting to the HA pair:



FCoE initiator to FCoE target

Using host FCoE initiators (CNAs), you can connect hosts to both controllers in an HA pair to FCoE target ports (also called UTAs or UTA2s) through FCoE switches.



FCoE initiator to FCoE and FC targets

Using host FCoE initiators (CNAs), you can connect hosts to both controllers in an HA pair to FCoE and FC target ports (also called UTAs or UTA2s) through FCoE switches.



FCoE mixed with IP storage protocols

Using host FCoE initiators (CNAs), you can connect hosts to both controllers in an HA pair to FCoE target ports (also called UTAs or UTA2s) through FCoE switches. FCoE ports cannot use traditional link aggregation to a single switch. Cisco switches support a special type of link aggregation (Virtual Port Channel) that does support FCoE. A Virtual Port Channel aggregates individual links to two switches. You can also use Virtual Port Channels for other Ethernet traffic. Ports used for traffic other than FCoE, including NFS, SMB, iSCSI, and other Ethernet traffic, can use regular Ethernet ports on the FCoE switches.



ONTAP supported FCoE initiator and target port combinations

Certain combinations of FCoE and traditional FC initiators and targets are supported.

FCoE initiators

You can use FCoE initiators in host computers with both FCoE and traditional FC targets in storage controllers. The host FCoE initiator must connect to an FCoE DCB (data center bridging) switch; direct connection to a target is not supported.

The following table lists the supported combinations:

Initiator	Target	Supported?
FC	FC	Yes
FC	FCoE	Yes

Initiator	Target	Supported?
FCoE	FC	Yes
FCoE	FCoE	Yes

FCoE targets

You can mix FCoE target ports with 4-Gb, 8-Gb, or 16-Gb FC ports on the storage controller regardless of whether the FC ports are add-in target adapters or onboard ports. You can have both FCoE and FC target adapters in the same storage controller.



The rules for combining onboard and expansion FC ports still apply.

FC and FCoE zoning

Learn about FC and FCoE zoning with ONTAP systems

An FC, FC-NVMe or FCoE zone is a logical grouping of one or more ports within a fabric. For devices to be able see each other, connect, create sessions with one another, and communicate, both ports must be members of the same zone.

Zoning increases security by limiting access and connectivity to end-points that share a common zone. Ports that are not in the same zone cannot communicate with one another. This reduces or eliminates *crosstalk* between initiator HBAs. Should connectivity issues occur, zoning helps to isolate problems to a specific set of ports, thereby decreasing time to resolution.

Zoning reduces the number of available paths to a particular port and reduces the number of paths between a host and the storage system. For example, some host OS multipathing solutions have a limit on the number of paths they can manage. Zoning can reduce the number of paths visible to the host so that paths to the host do not exceed the maximum allowed by the host OS.

World Wide Name-based zoning

Zoning based on World Wide Name (WWN) specifies the WWN of the members to be included within the zone. Although World Wide Node Name (WWNN) zoning is possible with some switch vendors, when zoning in ONTAP, you must use World Wide Port Name (WWPN) zoning.

WWPN zoning is required to properly define a specific port and to use NPIV effectively. FC switches should be zoned using the WWPNs of the target's logical interfaces (LIFs), not the WWPNs of the physical ports on the node. The WWPNs of the physical ports start with "50" and the WWPNs of the LIFs start with "20".

WWPN zoning provides flexibility because access is not determined by where the device is physically connected to the fabric. You can move a cable from one port to another without reconfiguring zones.

Recommended FC and FCoE zoning configurations for ONTAP systems

You should create a zoning configuration if your host does not have a multipathing solution installed, if four or more hosts are connected to your SAN or if Selective LUN Mapping is not implemented on the nodes in your cluster.

In the recommended FC and FCoE zoning configuration, each zone includes one initiator port and one or more target LIFs. This configuration allows each host initiator to access any node, while preventing hosts accessing the same node from seeing see each other's ports

Add all LIFs from the storage virtual machine (SVM) to the zone with the host initiator. This allows you to move volumes or LUNs without editing your existing zones or creating new zones.

Dual-fabric zoning configurations

Dual-fabric zoning configurations are recommended because they provide protection against data loss due to a single component failure. In a dual-fabric configuration, each host initiator is connected to each node in the cluster using different switches. If one switch becomes unavailable, data access is maintained through the remaining switch. Multipathing software is required on the host to manage multiple paths.

In the following figure, the host has two initiators and is running multipathing software. There are two zones. Selective LUN Mapping (SLM) is configured so that all nodes are considered as reporting nodes.



The naming convention used in this figure is just a recommendation of one possible naming convention that you can choose to use for your ONTAP solution.

- Zone 1: HBA 0, LIF_1, LIF_3, LIF_5, and LIF_7
- Zone 2: HBA 1, LIF_2, LIF_4, LIF_6, and LIF_8

Each host initiator is zoned through a different switch. Zone 1 is accessed through Switch 1. Zone 2 is accessed through Switch 2.

Each host can access a LIF on every node. This enables the host to still access its LUNs if a node fails. SVMs have access to all iSCSI and FC LIFs on every node in the cluster based on your SLM reporting nodes configuration. You can use SLM, portsets, or FC switch zoning to reduce the number of paths from an SVM to the host and the number of paths from an SVM to a LUN.

If the configuration includes more nodes, the LIFs for the additional nodes are included in these zones.



The host operating system and multipathing software have to support the number of paths that is being used to access the LUNs on the nodes.

i.

Single-fabric zoning

In a single-fabric configuration, you connect each host initiator to each storage node through a single switch. Single-fabric zoning configurations are not recommended because they do not provide protection against data loss due to a single component failure. If you choose to configure single-fabric zoning, each host should have two initiators for multipathing to provide resiliency in the solution. Multipathing software is required on the host to manage multiple paths.

Each host initiator should have a minimum of one LIF from each node that the initiator can access. The zoning should allow at least one path from the host initiator to the HA pair of nodes in the cluster to provide a path for LUN connectivity. This means that each initiator on the host might only have one target LIF per node in its zone configuration. If there is a requirement for multipathing to the same node or multiple nodes in the cluster, then each node will have multiple LIFs per node in its zone configuration. This enables the host to still access its LUNs if a node fails or a volume containing the LUN is moved to a different node. This also requires the reporting nodes to be set appropriately.

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

In the following figure, the host has two initiators and is running multipathing software. There are two zones:



The naming convention used in this figure is just a recommendation of one possible naming convention that you can choose to use for your ONTAP solution.

- Zone 1: HBA 0, LIF_1, and LIF_3
- Zone 2: HBA 1, LIF_2, and LIF_4

If the configuration includes more nodes, the LIFs for the additional nodes are included in these zones.s.



In this example, you could also have all four LIFs in each zone. In that case, the zones would be as follows:

- Zone 1: HBA 0, LIF_1, LIF_2, LIF_3, and LIF_4
- Zone 2: HBA 1, LIF_1, LIF_2, LIF_3, and LIF_4



The host operating system and multipathing software have to support the number of supported paths that are being used to access the LUNs on the nodes. To determine the number of paths used to access the LUNs on nodes, see the SAN configuration limits section.

Zoning restrictions for Cisco FC and FCoE switches

When using Cisco FC and FCoE switches, certain restrictions apply to the use of physical ports and logical interfaces (LIFs) in zones.

Physical ports

- FC-NVMe and FC can share the same 32 Gb physical port
- FC-NVMe and FCoE cannot share the same physical port
- FC and FCoE can share the same physical port but their protocol LIFs must be in separate zones.

Logical Interfaces (LIFs)

• A zone can contain a LIF from every target port in the cluster.

Verify the SLM configuration so that you do not to exceed the maximum number of paths allowed to the host.

- Each LIF on a given port must be in a separate zone from other LIFs on that port
- LIFs on different physical ports can be in the same zone.

Requirements for SAN hosts connected to ONTAP and non-NetApp systems

Shared SAN configurations are defined as hosts that are attached to both ONTAP storage systems and other vendors' storage systems. Accessing ONTAP storage systems and other vendors' storage systems from a single host is supported as long as several requirements are met.

For all of the host operating systems, it is a best practice to use separate adapters to connect to each vendor's storage systems. Using separate adapters reduces the chances of conflicting drivers and settings. For connections to an ONTAP storage system, the adapter model, BIOS, firmware, and driver must be listed as supported in the NetApp Interoperability Matrix Tool.

You should set the required or recommended timeout values and other storage parameters for the host. You must always install the NetApp software or apply the NetApp settings last.

- For AIX, you should apply the values from the AIX Host Utilities version that is listed in the Interoperability Matrix Tool for your configuration.
- For ESX, you should apply host settings by using Virtual Storage Console for VMware vSphere.
- For HP-UX, you should use the HP-UX default storage settings.
- For Linux, you should apply the values from the Linux Host Utilities version that is listed in the Interoperability Matrix Tool for your configuration.
- For Solaris, you should apply the values from the Solaris Host Utilities version that is listed in the Interoperability Matrix Tool for your configuration.
- For Windows, you should install the Windows Host Utilities version that is listed in the Interoperability Matrix Tool for your configuration.

SAN configurations in a MetroCluster environment

Supported SAN configurations in an ONTAP MetroCluster environment

You must be aware of certain considerations when using SAN configurations in a MetroCluster environment.

- MetroCluster configurations do not support front-end FC fabric "routed" vSAN configurations.
- Beginning with ONTAP 9.15.1, four-node MetroCluster IP configurations are supported on NVMe/TCP.
- Beginning with ONTAP 9.12.1, four-node MetroCluster IP configurations are supported on NVMe/FC. MetroCluster configurations are not supported for front-end NVMe networks before ONTAP 9.12.1.
- Other SAN protocols such as iSCSI, FC, and FCoE are supported on MetroCluster configurations.
- When using SAN client configurations, you must check whether any special considerations for MetroCluster configurations are included in the notes that are provided in the NetApp Interoperability Matrix Tool (IMT).
- Operating systems and applications must provide an I/O resiliency of 120 seconds to support MetroCluster automatic unplanned switchover and Tiebreaker or Mediator-initiated switchover.
- MetroCluster configurations use the same WWNNs and WWPNs on both sides of the front-end FC fabric.

Related information

- Understanding MetroCluster data protection and disaster recovery
- Knowledge Base article: What are AIX Host support considerations in a MetroCluster configuration?
- Knowledge Base article: Solaris host support considerations in a MetroCluster configuration

Avoid port overlap during ONTAP MetroCluster switchover and switchback

In a SAN environment, you can configure the front-end switches to avoid overlap when the old port goes offline and the new port comes online.

During switchover, the FC port on the surviving site might log in to the fabric before the fabric has detected that the FC port on the disaster site is offline and has removed this port from the name and directory services.

If the FC port on the disaster is not yet removed, the fabric login attempt of the FC port at the surviving site might be rejected due to a duplicate WWPN. This behavior of the FC switches can be changed to honor the login of the previous device and not the existing one. You should verify the effects of this behavior on other fabric devices. Contact the switch vendor for more information.

Choose the correct procedure according to your switch type.

Example 9. Steps

Cisco switch

- 1. Connect to the switch and log in.
- 2. Enter configuration mode:

```
switch# config t
switch(config)#
```

3. Overwrite the first device entry in the name server database with the new device:

```
switch(config)# no fcns reject-duplicate-pwwn vsan 1
```

- 4. In switches that are running NX-OS 8.x, confirm that the flogi quiesce timeout is set to zero:
 - a. Display the quiesce timerval:

switch(config)# show flogi interval info \| i quiesce

Stats: fs flogi quiesce timerval: 0

b. If the output in the previous step does not indicate that the timerval is zero, then set it to zero:

switch(config)# flogi scale enable

```
switch(config)$ flogi quiesce timeout 0
```

Brocade switch

- 1. Connect to the switch and log in.
- 2. Enter the switchDisable command.
- 3. Enter the configure command, and press y at the prompt.

F-Port login parameters (yes, y, no, n): [no] y

4. Choose setting 1:

- 0: First login take precedence over the second login (default)
- 1: Second login overrides first login.
- 2: the port type determines the behavior
Enforce FLOGI/FDISC login: (0..2) [0] 1

5. Respond to the remaining prompts, or press Ctrl + D.

Related information

Performing switchover for tests or maintenance

ONTAP support for SAN host multipathing

ONTAP uses Asymmetric Logical Unit Access (ALUA) software for multipathing with both FC and iSCSI hosts.

Beginning with ONTAP 9.5 multipath high availability (HA) pair failover/giveback is supported for NVMe hosts using Asynchronous Namespace Access (ANA). In ONTAP 9.4, NVMe supports only one path from host to target, so the application host must manage path failover to its HA partner.

The multipathing software is required on your SAN host if it can access a LUN or NVMe namespace through more than one path. It presents a single disk to the operating system for all paths to a LUN or NVMe namespace. Without it, the operating system could treat each path as a separate disk, leading to data corruption.

Your solution is considered to have multiple paths if you have any of the following:

- A single initiator port in the host attaching to multiple SAN LIFs in the SVM
- · Multiple initiator ports attaching to a single SAN LIF in the SVM
- Multiple initiator ports attaching to multiple SAN LIFs in the SVM

Multipathing software, also known as MPIO (multipath I/O) software, is recommended in HA configurations. In addition to Selective LUN Map, using FC switch zoning or portsets to limit the paths used to access LUNs is also recommended.

For information about which specific host configurations support ALUA or ANA, see the NetApp Interoperability Matrix Tool and ONTAP SAN Host Configuration for your host operating system.

Recommended number of paths from host to nodes in cluster

You should not exceed more than eight paths from your host to each node in your cluster. You should also not exceed the total number of paths that can be supported for the host OS and the multipathing used on the host.

You should have a minimum of two paths per LUN connecting to each reporting node through Selective LUN Map (SLM) being used by the storage virtual machine (SVM) in your cluster. This eliminates single points of failure and enables the system to survive component failures.

If you have four or more nodes in your cluster or more than four target ports being used by the SVMs in any of your nodes, you can use the following methods to limit the number of paths that can be used to access LUNs on your nodes so that you do not exceed the recommended maximum of eight paths.

SLM

SLM reduces the number of paths from the host to LUN to only paths on the node owning the LUN and the owning node's HA partner. SLM is enabled by default.

• Portsets for iSCSI

- FC igroup mappings from your host
- FC switch zoning

Configuration limits

Determine the maximum supported nodes and SAN hosts per ONTAP cluster

The number of supported nodes per cluster varies depending on your version of ONTAP, your controller models, and the protocol of your cluster nodes. The maximum number of SAN hosts that can be connected to a cluster also varies based upon your specific configuration.

Determine the maximum supported nodes per cluster

If any node in the cluster is configured for FC, FC-NVMe, FCoE, or iSCSI, that cluster is limited to the SAN node limits. Node limits based on the controllers in your cluster are listed in the *Hardware Universe*.

Steps

- 1. Go to NetApp Hardware Universe.
- 2. In the upper left, next to Home, select Platforms; then select the platform type.
- 3. Select your version of ONTAP.

A new column is displayed for you to choose your platforms.

- 4. Select the platforms used in your solution.
- 5. Under Choose Your Specifications, deselect Select All.
- 6. Select Max Nodes per Cluster (NAS/SAN).
- 7. Click Show Results.

Results

The maximum nodes per cluster for your selected platforms is displayed.

Determine if your cluster can support more FC hosts

For FC and FC-NVMe configurations, you should use the number of initiator-target nexuses (ITNs) in your system to determine whether you can add more hosts to your cluster.

An ITN represents one path from the host's initiator to the storage system's target. The maximum number of ITNs per node in FC and FC-NVMe configurations is 2,048. If you are below the maximum number of ITNs, you can continue to add hosts to your cluster.

To determine the number of ITNs used in your cluster, perform the following steps for each node in the cluster.

Steps

- 1. Identify all the LIFs on a given node.
- 2. Run the following command for every LIF on the node:

fcp initiator show -fields wwpn, lif
The number of entries displayed at the bottom of the command output represents your number of ITNs for that LIF.

- 3. Record the number of ITNs displayed for each LIF.
- 4. Add the number of ITNs for each LIF on every node in your cluster.

This total represents the number of ITNs in your cluster.

Determine if your cluster can support more iSCSI hosts

The number of hosts that can be directly connected to a node or that can be connected through one or more switches depends on the number of available Ethernet ports. The number of available Ethernet ports is determined by the model of the controller and the number and type of adapters installed in the controller. The number of supported Ethernet ports for controllers and adapters is available in the *Hardware Universe*.

For all multi-node cluster configurations, you must determine the number of iSCSI sessions per node to know whether you can add more hosts to your cluster. As long as your cluster is below the maximum number of iSCSI sessions per node, you can continue to add hosts to your cluster. The maximum number of iSCSI sessions per node varies based on the types of controllers in your cluster.

Steps

- 1. Identify all of the target portal groups on the node.
- 2. Check the number of iSCSI sessions for every target portal group on the node:

iscsi session show -tpgroup _tpgroup_

The number of entries displayed at the bottom of the command output represents your number of iSCSI sessions for that target portal group.

- 3. Record the number of iSCSI sessions displayed for each target portal group.
- 4. Add the number of iSCSI sessions for each target portal group on the node.

The total represents the number of iSCSI sessions on your node.

All-Flash SAN Array configuration limits and support

All-Flash SAN Array (ASA) configuration limits and support varies by ONTAP version.

The most current details on supported configuration limits are available in NetApp Hardware Universe.



These limitations apply to ASA systems. If you have an ASA r2 system (ASA A1K, ASA A90, ASA A70, ASA A50, ASA A30, ASA A20, or ASA C30), see ASA r2 system storage limits.

SAN protocols and supported number of nodes per cluster

The supported SAN protocols and maximum number of nodes per cluster depends on whether you have a non-MetroCluster or MetroCluster configuration:

Non-MetroCluster configurations

The following table shows the ASA support for SAN protocols and the supported number of nodes per cluster in non-MetroCluster configurations:

Beginning with ONTAP	Protocol support	Maximum nodes per cluster
9.11.1	NVMe/TCPNVMe/FC	12
9.10.1	NVMe/TCP	2
9.9.1	NVMe/FC	2
	• FC • iSCSI	12
9.7	• FC • iSCSI	2

MetroCluster IP configurations

The following table shows the ASA support for SAN protocols and the supported number of nodes per cluster in MetroCluster IP configurations:

Beginning with ONTAP	Protocol support	Maximum nodes per cluster
9.15.1	NVMe/TCP	2 nodes per cluster in four-node MetroCluster IP configurations
9.12.1	• NVMe/FC	2 nodes per cluster in four-node MetroCluster IP configurations
9.9.1	• FC • iSCSI	4 nodes per cluster in eight-node MetroCluster IP configurations
9.7	• FC • iSCSI	2 nodes per cluster in four-node MetroCluster IP configurations

Support for persistent ports

Beginning with ONTAP 9.8, persistent ports are enabled by default on All-Flash SAN Arrays (ASAs) that are configured to use the FC protocol. Persistent ports are only available for FC and require zone membership identified by World Wide Port Name (WWPN).

Persistent ports reduce the impact of takeovers by creating a shadow LIF on the corresponding physical port of the high-availability (HA) partner. When a node is taken over, the shadow LIF on the partner node assumes the identity of the original LIF, including the WWPNe. Before the status of path to the taken over node is changed

to faulty, the shadow LIF appears as an Active/Optimized path to the host MPIO stack, and I/O is shifted. This reduces I/O disruption because the host always sees the same number of paths to the target, even during storage failover operations.

For persistent ports, the following FCP port characteristics should be identical within the HA pair:

- FCP port counts
- FCP port names
- · FCP port speeds
- FCP LIF WWPN-based zoning

If any of these characteristics are not identical within the HA pair, the following EMS message is generated:

EMS : scsiblade.lif.persistent.ports.fcp.init.error

For more information on persistent ports, see NetApp Technical Report 4080: Best Practices for Modern SAN.

Configuration limits for FC switches used with ONTAP systems

Fibre Channel switches have maximum configuration limits, including the number of logins supported per port, port group, blade, and switch. The switch vendors document their supported limits.

Each FC logical interface (LIF) logs into an FC switch port. The total number of logins from a single target on the node equals the number of LIFs plus one login for the underlying physical port. Do not exceed the switch vendor's configuration limits for logins or other configuration values. This also holds true for the initiators being used on the host side in virtualized environments with NPIV enabled. Do not exceed the switch vendor's configuration limits for logins for either the target or the initiators being used in the solution.

Brocade switch limits

You can find the configuration limits for Brocade switches in the Brocade Scalability Guidelines.

Cisco Systems switch limits

You can find the configuration limits for Cisco switches in the Cisco Configuration Limits guide for your version of Cisco switch software.

Maximum FC and FCoE hop count supported in ONTAP

The hop count is defined as the number of switches in the path between the initiator (host) and target (storage system). The maximum supported FC hop count between a host and storage system varies depending on the switch supplier.

Documentation from Cisco Systems also refers to this value as the diameter of the SAN fabric.

For FCoE, you can have FCoE switches connected to FC switches. For end-to-end FCoE connections, the FCoE switches must be running a firmware version that supports Ethernet inter-switch links (ISLs).

Switch supplier	Supported hop count
Brocade	 7 for FC 5 for FCoE
Cisco	 7 for FC Up to 3 of the switches can be FCoE switches.

Calculate queue depth for ONTAP FC hosts

You might need to tune your FC queue depth on the host to achieve the maximum values for ITNs per node and FC port fan-in. The maximum number of LUNs and the number of HBAs that can connect to an FC port are limited by the available queue depth on the FC target ports.

About this task

Queue depth is the number of I/O requests (SCSI commands) that can be queued at one time on a storage controller. Each I/O request from the host's initiator HBA to the storage controller's target adapter consumes a queue entry. Typically, a higher queue depth equates to better performance. However, if the storage controller's maximum queue depth is reached, that storage controller rejects incoming commands by returning a QFULL response to them. If a large number of hosts are accessing a storage controller, you should plan carefully to avoid QFULL conditions, which significantly degrade system performance and can lead to errors on some systems.

In a configuration with multiple initiators (hosts), all hosts should have similar queue depths. Because of the inequality in queue depth between hosts connected to the storage controller through the same target port, hosts with smaller queue depths are being deprived of access to resources by hosts with larger queue depths.

The following general recommendations can be made about "tuning" queue depths:

- For small to mid-size systems, use an HBA queue depth of 32.
- For large systems, use an HBA queue depth of 128.
- For exception cases or performance testing, use a queue depth of 256 to avoid possible queuing problems.
- All hosts should have the queue depths set to similar values to give equal access to all hosts.
- To avoid performance penalties or errors, the storage controller target FC port queue depth must not be exceeded.

Steps

- 1. Count the total number of FC initiators in all of the hosts that connect to one FC target port.
- 2. Multiply by 128.
 - If the result is less than 2,048, set the queue depth for all initiators to 128.
 You have 15 hosts with one initiator connected to each of two target ports on the storage controller. 15 × 128 = 1,920. Because 1,920 is less than the total queue depth limit of 2,048, you can set the queue depth for all of your initiators to 128.
 - If the result is greater than 2,048, go to step 3.
 You have 30 hosts with one initiator connected to each of two target ports on the storage controller. 30
 × 128 = 3,840. Because 3,840 is greater than the total queue depth limit of 2,048, you should choose

one of the options under step 3 for remediation.

- 3. Choose one of the following options to add more hosts to the storage controller.
 - Option 1:
 - i. Add more FC target ports.
 - ii. Redistribute your FC initiators.
 - iii. Repeat steps 1 and 2.

The desired queue depth of 3,840 exceeds the available queue depth per port. To remedy this, you can add a two-port FC target adapter to each controller, then rezone your FC switches so that 15 of your 30 hosts connect to one set of ports, and the remaining 15 hosts connect to a second set of ports. The queue depth per port is then reduced to $15 \times 128 = 1,920$.

- Option 2:
 - i. Designate each host as "large" or "small" based on its expected I/O need.
 - ii. Multiply the number of large initiators by 128.
 - iii. Multiply the number of small initiators by 32.
 - iv. Add the two results together.
 - v. If the result is less than 2,048, set the queue depth for large hosts to 128 and the queue depth for small hosts to 32.
 - vi. If the result is still greater than 2,048 per port, reduce the queue depth per initiator until the total queue depth is less than or equal to 2,048.

To estimate the queue depth needed to achieve a certain I/O per second throughput, use this formula:



Needed queue depth = (Number of I/O per second) × (Response time)

For example, if you need 40,000 I/O per second with a response time of 3 milliseconds, the needed queue depth = $40,000 \times (.003) = 120$.

The maximum number of hosts that you can connect to a target port is 64, if you decide to limit the queue depth to the basic recommendation of 32. However, if you decide to have a queue depth of 128, then you can have a maximum of 16 hosts connected to one target port. The larger the queue depth, the fewer hosts that a single target port can support. If your requirement is such that you cannot compromise on the queue depth, then you should get more target ports.

The desired queue depth of 3,840 exceeds the available queue depth per port. You have 10 "large" hosts that have high storage I/O needs, and 20 "small" hosts that have low I/O needs. Set the initiator queue depth on the large hosts to 128 and the initiator queue depth on the small hosts to 32.

Your resulting total queue depth is $(10 \times 128) + (20 \times 32) = 1,920$.

You can spread the available queue depth equally across each initiator.

Your resulting queue depth per initiator is $2,048 \div 30 = 68$.

Modify queue depths for ONTAP SAN hosts

You might need to change the queue depths on your host to achieve the maximum values

for ITNs per node and FC port fan-in. You can calculate the optimal queue depth for your environment.

AIX hosts

You can change the queue depth on AIX hosts using the chdev command. Changes made using the chdev command persist across reboots.

Examples:

• To change the queue depth for the hdisk7 device, use the following command:

```
chdev -1 hdisk7 -a queue_depth=32
```

• To change the queue depth for the fcs0 HBA, use the following command:

```
chdev -l fcs0 -a num cmd elems=128
```

The default value for num cmd elems is 200. The maximum value is 2,048.



It might be necessary to take the HBA offline to change <code>num_cmd_elems</code> and then bring it back online using the <code>rmdev -l fcs0 -R</code> and <code>makdev -l fcs0 -P</code> commands.

HP-UX hosts

You can change the LUN or device queue depth on HP-UX hosts using the kernel parameter scsi max qdepth. You can change the HBA queue depth using the kernel parameter max fcp reqs.

• The default value for scsi max qdepth is 8. The maximum value is 255.

scsi_max_qdepth can be dynamically changed on a running system using the -u option on the kmtune command. The change will be effective for all devices on the system. For example, use the following command to increase the LUN queue depth to 64:

kmtune -u -s scsi_max_qdepth=64

It is possible to change queue depth for individual device files using the scsictl command. Changes using the scsictl command are not persistent across system reboots. To view and change the queue depth for a particular device file, execute the following command:

```
scsictl -a /dev/rdsk/c2t2d0
scsictl -m queue depth=16 /dev/rdsk/c2t2d0
```

• The default value for max fcp reqs is 512. The maximum value is 1024.

The kernel must be rebuilt and the system must be rebooted for changes to max_fcp_reqs to take effect. To change the HBA queue depth to 256, for example, use the following command:

```
kmtune -u -s max_fcp_reqs=256
```

Solaris hosts

You can set the LUN and HBA queue depth for your Solaris hosts.

- For LUN queue depth: The number of LUNs in use on a host multiplied by the per-LUN throttle (lun-queue-depth) must be less than or equal to the tgt-queue-depth value on the host.
- For queue depth in a Sun stack: The native drivers do not allow for per LUN or per target max_throttle settings at the HBA level. The recommended method for setting the max_throttle value for native drivers is on a per-device type (VID_PID) level in the /kernel/drv/sd.conf and /kernel/drv/sd.conf files. The host utility sets this value to 64 for MPxIO configurations and 8 for Veritas DMP configurations.

Steps

- 1. # cd/kernel/drv
- 2. # vi lpfc.conf
- 3. Search for /tft-queue (/tgt-queue)

```
tgt-queue-depth=32
```



The default value is set to 32 at installation.

- 4. Set the desired value based on the configuration of your environment.
- 5. Save the file.
- 6. Reboot the host using the sync; sync; sync; reboot -- -r command.

VMware hosts for a QLogic HBA

Use the esxcfg-module command to change the HBA timeout settings. Manually updating the esx.conf file is not recommended.

Steps

- 1. Log on to the service console as the root user.
- 2. Use the #vmkload mod -1 command to verify which Qlogic HBA module is currently loaded.
- 3. For a single instance of a Qlogic HBA, run the following command:

#esxcfg-module -s ql2xmaxqdepth=64 qla2300_707



This example uses $qla2300_707$ module. Use the appropriate module based on the output of $vmkload_mod -l$.

4. Save your changes using the following command:

#/usr/sbin/esxcfg-boot -b

5. Reboot the server using the following command:

#reboot

6. Confirm the changes using the following commands:

a. #esxcfg-module -g qla2300 707

b. qla2300_707 enabled = 1 options = 'ql2xmaxqdepth=64'

VMware hosts for an Emulex HBA

Use the <code>esxcfg-module</code> command to change the HBA timeout settings. Manually updating the <code>esx.conf</code> file is not recommended.

Steps

- 1. Log on to the service console as the root user.
- 2. Use the #vmkload mod -1 grep lpfc command to verify which Emulex HBA is currently loaded.
- 3. For a single instance of an Emulex HBA, enter the following command:

#esxcfg-module -s lpfc0 lun queue depth=16 lpfcdd 7xx



Depending on the model of the HBA, the module can be either lpfcdd_7xx or lpfcdd_732. The above command uses the lpfcdd_7xx module. You should use the appropriate module based on the outcome of vmkload_mod -1.

Running this command will set the LUN queue depth to 16 for the HBA represented by lpfc0.

4. For multiple instances of an Emulex HBA, run the following command:

```
a esxcfg-module -s "lpfc0_lun_queue_depth=16 lpfc1_lun_queue_depth=16" lpfcdd 7xx
```

The LUN queue depth for lpfc0 and the LUN queue depth for lpfc1 is set to 16.

5. Enter the following command:

#esxcfg-boot -b

6. Reboot using #reboot.

Windows hosts for an Emulex HBA

On Windows hosts, you can use the LPUTILNT utility to update the queue depth for Emulex HBAs.

Steps

- 1. Run the LPUTILNT utility located in the C:\WINNT\system32 directory.
- 2. Select Drive Parameters from the menu on the right side.
- 3. Scroll down and double-click QueueDepth.



If you are setting **QueueDepth** greater than 150, the following Windows Registry value also need to be increased appropriately:

HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\lpxnds\Paramete
rs\Device\NumberOfRequests

Windows hosts for a Qlogic HBA

On Windows hosts, you can use theand the SANsurfer HBA manager utility to update the queue depths for Qlogic HBAs.

Steps

- 1. Run the SANsurfer HBA manager utility.
- 2. Click on **HBA port** > **Settings**.
- 3. Click Advanced HBA port settings in the list box.
- 4. Update the Execution Throttle parameter.

Linux hosts for Emulex HBA

You can update the queue depths of an Emulex HBA on a Linux host. To make the updates persistent across reboots, you must then create a new RAM disk image and reboot the host.

Steps

1. Identify the queue depth parameters to be modified:

```
modinfo lpfc|grep queue_depth
```

The list of queue depth parameters with their description is displayed. Depending on your operating system version, you can modify one or more of the following queue depth parameters:

- ° lpfc_lun_queue_depth: Maximum number of FC commands that can be queued to a specific LUN
 (uint)
- ° lpfc_hba_queue_depth: Maximum number of FC commands that can be queued to an lpfc HBA
 (uint)
- o lpfc_tgt_queue_depth: Maximum number of FC commands that can be queued to a specific target
 port (uint)

The lpfc_tgt_queue_depth parameter is applicable only for Red Hat Enterprise Linux 7.x systems, SUSE Linux Enterprise Server 11 SP4 systems and 12.x systems.

2. Update the queue depths by adding the queue depth parameters to the /etc/modprobe.conf file for a Red Hat Enterprise Linux 5.x system and to the /etc/modprobe.d/scsi.conf file for a Red Hat Enterprise Linux 6.x or 7.x system, or a SUSE Linux Enterprise Server 11.x or 12.x system.

Depending on your operating system version, you can add one or more of the following commands:

- ° options lpfc lpfc_hba_queue_depth=new_queue_depth
- ° options lpfc lpfc_lun_queue_depth=new_queue_depth
- ° options lpfc tgt queue depth=new queue depth
- 3. Create a new RAM disk image, and then reboot the host to make the updates persistent across reboots.

For more information, see the System administration for your version of Linux operating system.

4. Verify that the queue depth values are updated for each of the queue depth parameter that you have modified:

```
cat /sys/class/scsi_host/host_number/lpfc_lun_queue_depthcat
/sys/class/scsi_host/host_number/lpfc_tgt_queue_depthcat
/sys/class/scsi_host/host_number/lpfc_hba_queue_depth
```

```
root@localhost ~]#cat /sys/class/scsi_host/host5/lpfc_lun_queue_depth
30
```

The current value of the queue depth is displayed.

Linux hosts for QLogic HBA

You can update the device queue depth of a QLogic driver on a Linux host. To make the updates persistent across reboots, you must then create a new RAM disk image and reboot the host. You can use the QLogic HBA management GUI or command-line interface (CLI) to modify the QLogic HBA queue depth.

This task shows how to use the QLogic HBA CLI to modify the QLogic HBA queue depth

Steps

1. Identify the device queue depth parameter to be modified:

modinfo qla2xxx | grep ql2xmaxqdepth

You can modify only the <code>ql2xmaxqdepth</code> queue depth parameter, which denotes the maximum queue depth that can be set for each LUN. The default value is 64 for RHEL 7.5 and later. The default value is 32 for RHEL 7.4 and earlier.

```
root@localhost ~]# modinfo qla2xxx|grep ql2xmaxqdepth
parm: ql2xmaxqdepth:Maximum queue depth to set for each LUN.
Default is 64. (int)
```

- 2. Update the device queue depth value:
 - If you want to make the modifications persistent, perform the following steps:
 - i. Update the queue depths by adding the queue depth parameter to the /etc/modprobe.conf file for a Red Hat Enterprise Linux 5.x system and to the /etc/modprobe.d/scsi.conf file for a Red Hat Enterprise Linux 6.x or 7.x system, or a SUSE Linux Enterprise Server 11.x or 12.x system: options qla2xxx ql2xmaxqdepth=new queue depth
 - ii. Create a new RAM disk image, and then reboot the host to make the updates persistent across reboots.

For more information, see the System administration for your version of Linux operating system.

• If you want to modify the parameter only for the current session, run the following command:

echo new queue depth > /sys/module/qla2xxx/parameters/ql2xmaxqdepth

In the following example, the queue depth is set to 128.

echo 128 > /sys/module/qla2xxx/parameters/ql2xmaxqdepth

3. Verify that the queue depth values are updated:

```
cat /sys/module/qla2xxx/parameters/ql2xmaxqdepth
```

The current value of the queue depth is displayed.

- 4. Modify the QLogic HBA queue depth by updating the firmware parameter Execution Throttle from the QLogic HBA BIOS.
 - a. Log in to the QLogic HBA management CLI:

```
/opt/QLogic Corporation/QConvergeConsoleCLI/qaucli
```

b. From the main menu, select the Adapter Configuration option.

```
[root@localhost ~]#
/opt/QLogic Corporation/QConvergeConsoleCLI/qaucli
Using config file:
/opt/QLogic Corporation/QConvergeConsoleCLI/qaucli.cfg
Installation directory: /opt/QLogic Corporation/QConvergeConsoleCLI
Working dir: /root
QConvergeConsole
        CLI - Version 2.2.0 (Build 15)
    Main Menu
    1: Adapter Information
    **2: Adapter Configuration**
    3: Adapter Updates
    4: Adapter Diagnostics
    5: Monitoring
    6: FabricCache CLI
    7: Refresh
    8: Help
    9: Exit
        Please Enter Selection: 2
```

c. From the list of adapter configuration parameters, select the HBA Parameters option.

```
1: Adapter Alias

2: Adapter Port Alias

**3: HBA Parameters**

4: Persistent Names (udev)

5: Boot Devices Configuration

6: Virtual Ports (NPIV)

7: Target Link Speed (iiDMA)

8: Export (Save) Configuration

9: Generate Reports

10: Personality

11: FEC

(p or 0: Previous Menu; m or 98: Main Menu; ex or 99: Quit)

Please Enter Selection: 3
```

d. From the list of HBA ports, select the required HBA port.

```
Fibre Channel Adapter Configuration
HBA Model QLE2562 SN: BFD1524C78510
1: Port 1: WWPN: 21-00-00-24-FF-8D-98-E0 Online
2: Port 2: WWPN: 21-00-00-24-FF-8D-98-E1 Online
HBA Model QLE2672 SN: RFE1241G81915
3: Port 1: WWPN: 21-00-00-0E-1E-09-B7-62 Online
4: Port 2: WWPN: 21-00-00-0E-1E-09-B7-63 Online
(p or 0: Previous Menu; m or 98: Main Menu; ex or 99: Quit)
Please Enter Selection: 1
```

The details of the HBA port are displayed.

e. From the HBA Parameters menu, select the Display HBA Parameters option to view the current value of the Execution Throttle option.

The default value of the Execution Throttle option is 65535.

```
HBA Parameters Menu

HBA : 2 Port: 1

SN : BFD1524C78510

HBA Model : QLE2562

HBA Desc. : QLE2562 PCI Express to 8Gb FC Dual Channel

FW Version : 8.01.02
```

: 21-00-00-24-FF-8D-98-E0 WWPN : 20-00-00-24-FF-8D-98-E0 WWNN Link : Online 1: Display HBA Parameters 2: Configure HBA Parameters 3: Restore Defaults (p or 0: Previous Menu; m or 98: Main Menu; x or 99: Quit) Please Enter Selection: 1 _____ HBA Instance 2: QLE2562 Port 1 WWPN 21-00-00-24-FF-8D-98-E0 PortID 03-07-00 Link: Online _____ _____ Connection Options : 2 - Loop Preferred, Otherwise Pointto-Point Data Rate : Auto : 2048 Frame Size Hard Loop ID : 0 Loop Reset Delay (seconds) : 5 Enable Host HBA BIOS : Enabled Enable Hard Loop ID : Disabled Enable FC Tape Support : Enabled Operation Mode : 0 - Interrupt for every I/O completion Interrupt Delay Timer (100us) : 0 **Execution Throttle : 65535** Login Retry Count : 8 Port Down Retry Count : 30 : Enabled Enable LIP Full Login Link Down Timeout (seconds) : 30 Enable Target Reset : Enabled : 128 LUNs Per Target Out Of Order Frame Assembly : Disabled Enable LR Ext. Credits : Disabled Enable Fabric Assigned WWN : N/A Press <Enter> to continue:

f. Press **Enter** to continue.

- g. From the HBA Parameters menu, select the Configure HBA Parameters option to modify the HBA parameters.
- h. From the Configure Parameters menu, select the Execute Throttle option and update the value of this parameter.

```
Configure Parameters Menu
: 2 Port: 1
HBA
SN
           : BFD1524C78510
HBA Model : QLE2562
HBA Desc. : QLE2562 PCI Express to 8Gb FC Dual Channel
FW Version : 8.01.02
           : 21-00-00-24-FF-8D-98-E0
WWPN
WWNN
           : 20-00-00-24-FF-8D-98-E0
Link
           : Online
_____
   1: Connection Options
   2: Data Rate
   3: Frame Size
   4: Enable HBA Hard Loop ID
   5: Hard Loop ID
   6: Loop Reset Delay (seconds)
   7: Enable BIOS
   8: Enable Fibre Channel Tape Support
   9: Operation Mode
  10: Interrupt Delay Timer (100 microseconds)
  11: Execution Throttle
  12: Login Retry Count
  13: Port Down Retry Count
  14: Enable LIP Full Login
  15: Link Down Timeout (seconds)
  16: Enable Target Reset
  17: LUNs per Target
  18: Enable Receive Out Of Order Frame
  19: Enable LR Ext. Credits
  20: Commit Changes
  21: Abort Changes
       (p or 0: Previous Menu; m or 98: Main Menu; x or 99: Quit)
       Please Enter Selection: 11
Enter Execution Throttle [1-65535] [65535]: 65500
```

- i. Press Enter to continue.
- j. From the Configure Parameters menu, select the Commit Changes option to save the changes.
- k. Exit the menu.

Copyright information

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

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

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

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

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

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

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

Trademark information

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