



SAN administration

ONTAP 9

NetApp
February 20, 2026

Table of Contents

- SAN administration 1
 - SAN provisioning 1
 - SAN management overview 1
 - Learn about All-Flash SAN Array configurations 2
 - Configure switches for FCoE 3
 - System Requirements 4
 - What to know before you create a LUN 4
 - Verify and add your protocol FC or iSCSI license 5
 - Provision SAN storage 6
 - NVMe provisioning 11
 - NVMe Overview 11
 - NVMe license requirements 12
 - NVMe configuration, support, and limitations 13
 - Configure a storage VM for NVMe 15
 - Provision NVMe storage 19
 - Map an NVMe namespace to a subsystem 21
- Manage LUNs 22
 - Edit LUN QoS policy group 22
 - Convert a LUN into a namespace 23
 - Take a LUN offline 23
 - Resize a LUN in ONTAP 24
 - Move a LUN 27
 - Delete LUNs 28
 - What to know before copying LUNs 29
 - Examine configured and used space of a LUN 30
 - Control and monitor I/O performance to LUNs by using Storage QoS 30
 - Tools available to effectively monitor your LUNs 31
 - Capabilities and restrictions of transitioned LUNs 31
 - I/O misalignments on properly aligned LUNs overview 32
 - Ways to address issues when LUNs go offline 34
 - Troubleshoot iSCSI LUNs not visible on the host 35
- Manage igroups and portsets 36
 - Ways to limit LUN access with portsets and igroups 36
 - View and manage SAN initiators and igroups 37
 - Create a nested igroup 38
 - Map igroups to multiple LUNs 39
 - Create a portset and bind to an igroup 39
 - Manage portsets 41
 - Selective LUN Map overview 41
- Manage iSCSI protocol 42
 - Configure your network for best performance 42
 - Configure an SVM for iSCSI 43
 - Define a security policy method for an initiator 45

| | |
|---|----|
| Delete an iSCSI service for an SVM | 45 |
| Get more details in iSCSI session error recoveries | 46 |
| Register the SVM with an iSNS server | 46 |
| Resolve iSCSI error messages on the storage system | 47 |
| Enable or disable automatic iSCSI LIF failover | 48 |
| Manage FC protocol | 49 |
| Configure an SVM for FC | 49 |
| Delete an FC service for an SVM | 51 |
| Recommended MTU configurations for FCoE jumbo frames | 51 |
| Manage NVMe protocol | 51 |
| Start the NVMe service for an SVM | 51 |
| Delete NVMe service from an SVM | 52 |
| Resize a namespace | 52 |
| Convert a namespace into a LUN | 53 |
| Set up in-band authentication over NVMe | 53 |
| Disable in-band authentication over NVMe | 56 |
| Set up TLS secure channel for NVMe/TCP | 57 |
| Disable TLS secure channel for NVMe/TCP | 59 |
| Change NVMe host priority | 59 |
| Manage automated host discovery of NVMe/TCP controllers in ONTAP | 60 |
| Disable NVMe host virtual machine identifier in ONTAP | 61 |
| Manage systems with FC adapters | 61 |
| Manage systems with FC adapters | 61 |
| Commands for managing FC adapters | 62 |
| Configure FC adapters | 63 |
| View adapter settings | 64 |
| Change the UTA2 port from CNA mode to FC mode | 65 |
| Change the CNA/UTA2 target adapter optical modules | 67 |
| Supported port configurations for X1143A-R6 adapters | 68 |
| Configure the ports | 68 |
| Prevent loss of connectivity when using the X1133A-R6 adapter | 68 |
| Manage LIFs for all SAN protocols | 69 |
| Manage LIFs for all SAN protocols | 69 |
| Configure an NVMe LIF in ONTAP | 69 |
| What to know before moving a SAN LIF | 70 |
| Remove a SAN LIF from a port set | 70 |
| Move a SAN LIF | 71 |
| Delete a LIF in a SAN environment | 72 |
| SAN LIF requirements for adding nodes to a cluster | 73 |
| Configure iSCSI LIFs to return FQDN to host iSCSI SendTargets Discovery Operation | 74 |
| Enable ONTAP space allocation for SAN protocols | 75 |
| Host configuration for VMware ESXi 8.x and later NVMe hosts | 76 |
| Recommended volume and file or LUN configuration combinations | 77 |
| Recommended volume and file or LUN configuration combinations overview | 77 |
| Determine the correct volume and LUN configuration combination for your environment | 78 |

| | |
|---|----|
| Calculate rate of data growth for LUNs | 79 |
| Configuration settings for space-reserved files or LUNs with thick-provisioned volumes | 80 |
| Configuration settings for non-space-reserved files or LUNs with thin-provisioned volumes | 80 |
| Configuration settings for space-reserved files or LUNs with semi-thick volume provisioning | 81 |

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 world-wide 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 bus 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 node1 -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`.

```
fcv 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](#).

Example 1. Steps

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 | Type | Description | Expiration |
|---------|------|----------------------|------------|
| Base | site | Cluster Base License | - |
| NFS | site | NFS License | - |
| CIFS | site | CIFS License | - |
| iSCSI | site | iSCSI License | - |
| FCP | site | FCP License | - |

2. If you do not have a 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.

Example 2. Steps

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 |
|--|--------------------------|
| <ul style="list-style-type: none">• Assign QoS policy to LUNs instead of parent volume<ul style="list-style-type: none">◦ 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.<p>By default, performance limits are applied at the volume level.</p> | ONTAP 9.10.1 |
| <ul style="list-style-type: none">• Create a new initiator group using existing initiator groups<ul style="list-style-type: none">◦ More Options > HOST INFORMATION◦ Select New initiator group using existing initiator groups.<div data-bbox="337 1297 396 1356"></div><p>The OS type for an igroup containing other igroups cannot be changed after it has been created.</p> | ONTAP 9.9.1 |
| <ul style="list-style-type: none">• Add a description to your igroup or host initiator<p>The description serves as an alias for the igroup or host initiator.</p><ul style="list-style-type: none">◦ More Options > HOST INFORMATION | ONTAP 9.9.1 |
| <ul style="list-style-type: none">• Create your LUN on an existing volume<p>By default, a new LUN is created in a new volume.</p><ul style="list-style-type: none">◦ More Options > Add LUNs◦ Select Group related LUNs. | ONTAP 9.9.1 |

- Disable QoS or choose a custom QoS policy
 - **More Options > Storage and Optimization**
 - Select **Performance Service Level**.

ONTAP 9.8



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 | Type | Description | Expiration |
|---------|------|----------------------|------------|
| Base | site | Cluster Base License | - |
| NFS | site | NFS License | - |
| CIFS | site | CIFS License | - |
| iSCSI | site | iSCSI License | - |
| FCP | site | FCP License | - |

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

```
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
- NVMe 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 bottlenecks 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 to 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 supports |
|-------------------------|---|
| 9.17.1 | <ul style="list-style-type: none">• SnapMirror active sync NVMe/FC and NVMe/TCP host access for VMware workloads. |
| 9.15.1 | <ul style="list-style-type: none">• Four-node MetroCluster IP configurations on NVMe/TCP |
| 9.14.1 | <ul style="list-style-type: none">• Setting the host priority at the subsystem (host-level QoS) |

| | |
|--------|--|
| 9.12.1 | <ul style="list-style-type: none"> • 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 | <ul style="list-style-type: none"> • Namespaces and LUNs coexistence on the same volume |
| 9.8 | <ul style="list-style-type: none"> • Protocol co-existence <p>SCSI, NAS and NVMe protocols can exist on the same storage virtual machine (SVM).</p> <p>Prior to ONTAP 9.8, NVMe can be the only protocol on the SVM.</p> |
| 9.6 | <ul style="list-style-type: none"> • 512 byte blocks and 4096 byte blocks for namespaces <p>4096 is the default value. 512 should only be used if the host operating system does not support 4096 byte blocks.</p> <ul style="list-style-type: none"> • Volume move with mapped namespaces |
| 9.5 | <ul style="list-style-type: none"> • 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 |
|---|--|
| <ol style="list-style-type: none">1. In System Manager, click Storage > Storage VMs and then click Add.2. Enter a name for the storage VM.3. Select NVMe for the Access Protocol.4. Select Enable NVMe/FC or Enable NVMe/TCP and Save. | <ol style="list-style-type: none">1. In System Manager, click Storage > Storage VMs.2. Click on the storage VM you want to configure.3. Click on the Settings tab, and then click  next to the NVMe protocol.4. 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>
```

If a warning message is displayed about the auto efficiency policy, it can be safely ignored.

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 `nvme` 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 | <ol style="list-style-type: none">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 | <ul style="list-style-type: none">• 128 TB for All-Flash SAN Array (ASA) platforms• 16 TB for non-ASA platforms |
| ONTAP 9.5, 9.6, 9.7 | 16TB |

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

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  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>
```



ONTAP operations round down the actual maximum size of the LUN so it is slightly less than the expected value. Also, actual LUN size might vary slightly based on the OS type of the LUN. To obtain the exact resized value, run the following commands in advanced mode:

```
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.

Example 5. Steps

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 shows 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      25.00GB
vs3      /vol/vol0/lun1_backup    50.01GB disabled      32.15GB
vs3      /vol/vol0/lun2           75.00GB disabled         0B
vs3      /vol/volospace/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 `qos 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 512-byte data block used by the host is referred to as a logical block. The 4-KB data block used by the LUN to store data is referred to as a physical block. This means that there are eight 512-byte logical blocks in each 4-KB physical block.



The host operating system can begin a read or write I/O operation at any logical block. I/O operations are only considered aligned when they begin at the first logical block in the physical block. If an I/O operation begins at a logical block that is not also the start of a physical block, the I/O is considered misaligned. 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 [NetApp Knowledge Base: 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 [NetApp Knowledge Base: 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 | <ul style="list-style-type: none">• Add more disks.• Use the <code>volume modify</code> command to shrink a volume that has available space.• If you have space-guarantee volumes that have available space, change the volume space guarantee to <code>none</code> with the <code>volume modify</code> command. |

| If the... | You can... |
|---|--|
| Volume is full but there is space available in the containing aggregate | <ul style="list-style-type: none"> • For space guarantee volumes, use the <code>volume modify</code> command to increase the size of your volume. • For thinly provisioned volumes, use the <code>volume modify</code> command to increase the maximum size of your volume. <p>If volume autogrow is not enabled, use <code>volume modify -autogrow-mode</code> to enable it.</p> <ul style="list-style-type: none"> • Delete snapshots manually with the <code>volume snapshot delete</code> command, or use the <code>volume snapshot autodelete modify</code> 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 | <p>Verify that there is TCP/IP connectivity between the host and storage system.</p> <ul style="list-style-type: none"> • From the storage system command line, ping the host interfaces that are being used for iSCSI: <pre>ping -node node_name -destination host_ip_address_for_iSCSI</pre> • From the host command line, ping the storage system interfaces that are being used for iSCSI: <pre>ping -node node_name -destination host_ip_address_for_iSCSI</pre> |

| Configuration setting | What to do |
|-------------------------|---|
| 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. |
| 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 <code>iscsi initiator show</code> 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: <ul style="list-style-type: none"> • <code>lun mapping show</code> displays all LUNs and the igroups to which they are mapped. • <code>lun mapping show -igroup</code> displays the LUNs mapped to a specific igroup. |
| iSCSI LIFs enable | Verify that the iSCSI logical interfaces are enabled. |

Related information

- [NetApp Interoperability Matrix Tool](#)
- [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, host1 does not have a portset. Without a portset, host1 can access LUN1 through both LIF1 and LIF2.



You can limit access to LUN1 by using a portset. In the following example, host1 can access LUN1 only through LIF1. However, host1 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  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  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.
- Click the **iSCSI Initiators** 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.

Example 6. Steps

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 , 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 `fc`. 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
```

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

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 , 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  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 reporting-nodes, 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 `igroup_prefix*` instead of `igroup_name`.

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.
10. 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 |
|---|--|
| <ol style="list-style-type: none">1. In System Manager, click Storage > Storage VMs and then click Add.2. Enter a name for the storage VM.3. Select iSCSI for the Access 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. | <ol style="list-style-type: none">1. In System Manager, click Storage > Storage VMs.2. Click on the storage VM you want to configure.3. Click on the Settings tab, and then click  next to the 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 vs1 -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 an SVM management LIF with an IPv4 address to an iSNS server with an IPv6 address.

8. Verify that the iSNS service is running:

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

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

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

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 <code>iscsi interface enable</code> command to enable the iSCSI service on the interface. For example: <pre>iscsi interface enable -vserver vs1 -lif lif1</pre> |
| 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: <ul style="list-style-type: none">• Inbound credentials on the storage system must match outbound credentials on the initiator.• Outbound credentials on the storage system must match inbound credentials on the initiator. |

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. This 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 |
|--|--|
| <ol style="list-style-type: none">1. In System Manager, click Storage > Storage VMs and then click Add.2. Enter a name for the storage VM.3. Select FC for the Access Protocol.4. Click Enable FC. + The FC ports are automatically assigned.5. Click Save. | <ol style="list-style-type: none">1. In System Manager, click Storage > Storage VMs.2. Click on the storage VM you want to configure.3. Click on the Settings tab, and then click  next to the FC protocol.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 `vserver fcp modify` command or the `vserver fcp stop` command.

Steps

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

```
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... |
|----------|--------------------|---------------|
| TCP | 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**.
9. 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} ]
Diffie-Hellman Authentication
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 Authentication
Group
[ -dhchap-mode {none|unidirectional|bidirectional} ]
Authentication Mode

```

Related information

- [vserver nvme subsystem controller show](#)

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. ONTAP does not support TLS hardware offload.

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**.
7. 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. See the [NetApp Knowledge Base: How to disable NVMe host virtual machine identifier in ONTAP](#).

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

Commands for managing FC initiator adapters

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

Commands for managing onboard FC adapters

| If you want to... | Use this command... |
|--|--|
| Display the status of the onboard FC ports | <pre>run -node <i>node_name</i> 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 | - |
| e0f | Default | Default | | down | 1500 | auto/10 | - |
| ... | | | | | | | |

net-f8040-34::> ucadmin show

| Admin | Node | Adapter | Current Mode | Current Type | Pending Mode | Pending Type |
|-------|------|---------|--------------|--------------|--------------|--------------|
|-------|------|---------|--------------|--------------|--------------|--------------|

| | | | | | | |
|---------|-----------------|----|-----|--------|---|---|
| offline | net-f8040-34-01 | 0e | cna | target | - | - |
| offline | net-f8040-34-01 | 0f | cna | target | - | - |
| ... | | | | | | |

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 | e0M | e0M |
| net-f8040-34 | m | e0e | e0i |
| net-f8040-34 | net-f8040-34-01_mgmt1 | e0M | e0M |

```
7 entries were displayed.
```

```
net-f8040-34::> uadmin 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|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 nodes in the cluster

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 vs1
```

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

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.

```
cluster::vserver*> vserver iscsi interface show -vserver vs1 -fields
sendtargets-fqdn
vserver lif          sendtargets-fqdn
-----
vs1      vs1_iscsi1 storagehost-005.example.com
vs1      vs1_iscsi2 storagehost-006.example.com
```

Related information

[ONTAP command reference](#)

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 | <ul style="list-style-type: none">• iSCSI• FC• NVMe | Enabled by default for newly created LUNs and all namespaces |
| 9.15.1 | <ul style="list-style-type: none">• iSCSI• FC | Enabled by default for newly created LUNs |
| | NVMe | Not supported |
| 9.14.1 and earlier | <ul style="list-style-type: none">• 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 $(20+18+17+20+20+23+22) / 7 = 20$ 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 |

| Volume setting | Value |
|-----------------------|---|
| 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. |
| 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.

Copyright information

Copyright © 2026 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.