



SAN configuration reference

ONTAP 9

NetApp
February 02, 2026

This PDF was generated from <https://docs.netapp.com/us-en/ontap/san-config/index.html> on February 02, 2026. Always check docs.netapp.com for the latest.

Table of Contents

SAN configuration reference	1
Learn about ONTAP SAN configuration	1
iSCSI configurations	1
Configure iSCSI networks with ONTAP systems	1
Benefits of using VLANs with ONTAP systems in iSCSI configurations	3
FC configurations	4
Configure FC or FC-NVME fabrics with ONTAP systems	4
Best practices to configure FC switches with ONTAP systems	6
Recommended FC target port configuration and speeds for ONTAP systems	6
Configure ONTAP FC adapter ports	7
ONTAP commands for managing FC adapters	10
Avoid connectivity loss to an ONTAP system using an X1133A-R6 adapter	11
FCoE configurations	11
Configure FCoE fabrics with ONTAP systems	12
ONTAP supported FCoE initiator and target port combinations	15
FC and FCoE zoning	16
Learn about FC and FCoE zoning with ONTAP systems	16
Recommended FC and FCoE zoning configurations for ONTAP systems	16
Requirements for SAN hosts connected to ONTAP and non-NetApp systems	19
SAN configurations in a MetroCluster environment	19
Supported SAN configurations in an ONTAP MetroCluster environment	20
Avoid port overlap during ONTAP MetroCluster switchover and switchback	20
ONTAP support for SAN host multipathing	22
Recommended number of paths from host to nodes in cluster	22
Configuration limits	23
Determine the maximum supported nodes and SAN hosts per ONTAP cluster	23
All-Flash SAN Array configuration limits and support	24
Configuration limits for FC switches used with ONTAP systems	26
Maximum FC and FCoE hop count supported in ONTAP	26
Calculate queue depth for ONTAP FC hosts	27
Modify queue depths for ONTAP SAN hosts	29

SAN configuration reference

Learn about ONTAP SAN configuration

A storage area network (SAN) consists of a storage solution connected to hosts over a SAN transport protocol such as iSCSI or FC. You can configure your SAN so that your storage solution attaches to your hosts through one or more switches. If you are using iSCSI, you can also configure your SAN so that your storage solution attaches directly to your host without using a switch.

In a SAN, multiple hosts, using different operating systems, such as Windows, Linux, or UNIX, can access the storage solution at the same time. You can use [Selective LUN mapping](#) and [portsets](#) to limit data access between the hosts and the storage.

For iSCSI, the network topology between the storage solution and the hosts is referred to as a network. For FC, FC/NVMe and FCoE the network topology between the storage solution and the hosts is referred to as a fabric. To create redundancy, which protects you against loss of data access, you should set up your SAN with HA pairs in a multi-network or multi-fabric configuration. Configurations using single nodes or single networks/fabrics are not fully redundant so are not recommended.

After your SAN is configured, you can [provision storage for iSCSI or FC](#), or you can [provision storage for FC/NVMe](#). Then you can connect to your hosts to begin servicing data.

SAN protocol support varies based on your version of ONTAP, your platform and your configuration. For details on your specific configuration, see the [NetApp Interoperability Matrix Tool](#).

Related information

- [SAN administration overview](#)
- [NVMe configuration, support and limitations](#)

iSCSI configurations

Configure iSCSI networks with ONTAP systems

You should set up your iSCSI configuration with high-availability (HA) pairs that attach directly to your iSCSI SAN hosts or that connect to your hosts through one or more IP switches.

[HA pairs](#) are defined as the reporting nodes for the Active/Optimized and the Active/Unoptimized paths that will be used by the hosts to access the LUNs. Multiple hosts, using different operating systems, such as Windows, Linux, or UNIX, can access the storage at the same time. Hosts require that a supported multipathing solution that supports ALUA be installed and configured. Supported operating systems and multipathing solutions can be verified on the [NetApp Interoperability Matrix Tool](#).

In a multi-network configuration, there are two or more switches connecting the hosts to the storage system. Multi-network configurations are recommended because they are fully redundant. In a single-network configuration, there is one switch connecting the hosts to the storage system. Single-network configurations are not fully redundant.




[Single-node configurations](#) are not recommended because they do not provide the redundancy needed to support fault tolerance and nondisruptive operations.

Related information

- Learn how [Selective LUN mapping \(SLM\)](#) limits the paths that are used to access the LUNs owned by an HA pair.
- Learn about [SAN LIFs](#).
- Learn about the [benefits of VLANs in iSCSI](#).


Multi-network iSCSI configurations

In multi-network HA pair configurations, two or more switches connect the HA pair to one or more hosts. Because there are multiple switches, this configuration is fully redundant.




Single-network iSCSI configurations

In single-network HA pair configurations, one switch connects the HA pair to one or more hosts. Because there is a single switch, this configuration is not fully redundant.



Direct-attachment iSCSI configuration

In a direct-attached configuration, one or more hosts are directly connected to the controllers.



Benefits of using VLANs with ONTAP systems in iSCSI configurations

A VLAN consists of a group of switch ports grouped together into a broadcast domain. A VLAN can be on a single switch or it can span multiple switch chassis. Static and dynamic VLANs enable you to increase security, isolate problems, and limit available paths within your IP network infrastructure.

When you implement VLANs in large IP network infrastructures, you derive the following benefits:

- Increased security.

VLANs enable you to leverage existing infrastructure while still providing enhanced security because they limit access between different nodes of an Ethernet network or an IP SAN.

- Improved Ethernet network and IP SAN reliability by isolating problems.
- Reduction of problem resolution time by limiting the problem space.

- Reduction of the number of available paths to a particular iSCSI target port.
- Reduction of the maximum number of paths used by a host.

Having too many paths slows reconnect times. If a host does not have a multipathing solution, you can use VLANs to allow only one path.

Dynamic VLANs

Dynamic VLANs are MAC address-based. You can define a VLAN by specifying the MAC address of the members you want to include.

Dynamic VLANs provide flexibility and do not require mapping to the physical ports where the device is physically connected to the switch. You can move a cable from one port to another without reconfiguring the VLAN.

Static VLANs

Static VLANs are port-based. The switch and switch port are used to define the VLAN and its members.

Static VLANs offer improved security because it is not possible to breach VLANs using media access control (MAC) spoofing. However, if someone has physical access to the switch, replacing a cable and reconfiguring the network address can allow access.

In some environments, it is easier to create and manage static VLANs than dynamic VLANs. This is because static VLANs require only the switch and port identifier to be specified, instead of the 48-bit MAC address. In addition, you can label switch port ranges with the VLAN identifier.

FC configurations

Configure FC or FC-NVME fabrics with ONTAP systems

It is recommended that you configure your FC and FC-NVMe SAN hosts using HA pairs and a minimum of two switches. This provides redundancy at the fabric and storage system layers to support fault tolerance and nondisruptive operations. You cannot directly attach FC or FC-NVMe SAN hosts to HA pairs without using a switch.


Cascade, partial mesh, full mesh, core-edge, and director fabrics are all industry-standard methods of connecting FC switches to a fabric, and all are supported. The use of heterogeneous FC switch fabrics is not supported, except in the case of embedded blade switches. Specific exceptions are listed on the [Interoperability Matrix Tool](#). A fabric can consist of one or multiple switches, and the storage controllers can be connected to multiple switches.

Multiple hosts, using different operating systems, such as Windows, Linux, or UNIX, can access the storage controllers at the same time. Hosts require that a supported multipathing solution be installed and configured. Supported operating systems and multipathing solutions can be verified on the Interoperability Matrix Tool.

Multifabric FC and FC-NVMe configurations

In multifabric HA pair configurations, there are two or more switches connecting HA pairs to one or more hosts. For simplicity, the following multifabric HA pair figure shows only two fabrics, but you can have two or more fabrics in any multifabric configuration.

The FC target port numbers (0c, 0d, 1a, 1b) in the illustrations are examples. The actual port numbers vary depending on the model of your storage node and whether you are using expansion adapters.




Single-fabric FC and FC-NVMe configurations

In single-fabric HA pair configurations, there is one fabric connecting both controllers in the HA pair to one or more hosts. Because the hosts and controllers are connected through a single switch, single-fabric HA pair configurations are not fully redundant.

The FC target port numbers (0a, 0c) in the illustrations are examples. The actual port numbers vary depending on the model of your storage node and whether you are using expansion adapters.

All platforms that support FC configurations support single-fabric HA pair configurations.



Single-node configurations are not recommended because they do not provide the redundancy needed to support fault tolerance and nondisruptive operations.

Related information

- Learn how [Selective LUN mapping \(SLM\)](#) limits the paths that are used to access the LUNs owned by an HA pair.
- Learn about [SAN LIFs](#).

Best practices to configure FC switches with ONTAP systems

For best performance, you should consider certain best practices when configuring your FC switch.

A fixed link speed setting is the best practice for FC switch configurations, especially for large fabrics because it provides the best performance for fabric rebuilds and can significantly save time. Although autonegotiation provides the greatest flexibility, FC switch configuration does not always perform as expected, and it adds time to the overall fabric-build sequence.

All of the switches that are connected to the fabric must support N_Port ID virtualization (NPIV) and must have NPIV enabled. ONTAP uses NPIV to present FC targets to a fabric.

For details about which environments are supported, see the [NetApp Interoperability Matrix Tool](#).

For FC and iSCSI best practices, see [NetApp Technical Report 4080: Best Practices for Modern SAN](#).

Recommended FC target port configuration and speeds for ONTAP systems

FC target ports can be configured and used for the FC-NVMe protocol in the exact same way they are configured and used for the FC protocol. Support for the FC-NVMe protocol varies based upon your platform and your ONTAP version. Use NetApp Hardware Universe to verify support.

For best performance and highest availability, you should use the recommended target port configuration listed in [NetApp Hardware Universe](#) for your specific platform.

Configuration for FC target ports with shared ASICs

The following platforms have port pairs with shared application-specific integrated circuits (ASICs). If you use an expansion adapter with these platforms, you should configure your FC ports so that they do not use the same ASIC for connectivity.

Controller	Port pairs with shared ASIC	Number of target ports: Recommended ports
<ul style="list-style-type: none">• FAS8200• AFF A300	0g+0h	1: 0g 2: 0g, 0h
<ul style="list-style-type: none">• FAS2720• FAS2750• AFF A220	0c+0d 0e+0f	1: 0c 2: 0c, 0e 3: 0c, 0e, 0d 4: 0c, 0e, 0d, 0f

FC target port supported speeds

FC target ports can be configured to run at different speeds. All target ports used by a given host should be set to the same speed. You should set the target port speed to match the speed of the device to which it connects. Do not use autonegotiation for your port speed. A port that is set to autonegotiation can take longer to reconnect after a takeover/giveback or other interruption.

You can configure onboard ports and expansion adapters to run at the following speeds. Each controller and expansion adapter port can be configured individually for different speeds as needed.

4 Gb ports	8 Gb ports	16 Gb ports	32 Gb ports
<ul style="list-style-type: none">• 4 Gb• 2 Gb• 1 Gb	<ul style="list-style-type: none">• 8 Gb• 4 Gb• 2 Gb	<ul style="list-style-type: none">• 16 Gb• 8 Gb• 4 Gb	<ul style="list-style-type: none">• 32 Gb• 16 Gb• 8 Gb

For a full list of supported adapters and their supported speeds, see the [NetApp Hardware Universe](#).

Configure ONTAP FC adapter ports

Onboard FC adapters and some FC expansion adapter cards can be individually configured as either initiators or targets ports. Other FC expansion adapters are configured as initiators or targets at the factory and cannot be changed. Additional FC ports are also available through supported UTA2 cards configured with FC SFP+ adapters.

Initiator ports can be used to connect directly to back-end disk shelves, and possibly foreign storage arrays. Target ports can be used to connect only to FC switches.

The number of onboard ports and CNA/UTA2 ports configured for FC varies depending on the model of the controller. The supported target expansion adapters also varies depending on controller model. See [NetApp Hardware Universe](#) for a complete list of onboard FC ports and supported target expansion adapters for your controller model.

Configure FC adapters for initiator mode

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

Before you begin

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



NVMe/FC does support initiator mode.

Steps

1. Remove all LIFs from the adapter:

```
network interface delete -vserver _SVM_name_ -lif _lif_name_,_lif_name_
```

2. Take your adapter offline:

```
network fcp adapter modify -node _node_name_ -adapter _adapter_port_
-status-admin down
```

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

3. Change the adapter from target to initiator:

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

4. Reboot the node hosting the adapter you changed.

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

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

6. Bring the adapter back online:

```
node run -node _node_name_ storage enable adapter _adapter_port_
```

Configure FC adapters for target mode

Target mode is used to connect the ports to FC initiators.

The same steps are used to configure FC adapters for the FC protocol and the FC-NVMe protocol. However, only certain FC adapters support FC-NVMe. See the [NetApp Hardware Universe](#) for a list of adapters that support the FC-NVMe protocol.

Steps

1. Take the adapter offline:

```
node run -node _node_name_ storage disable adapter _adapter_name_
```

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

2. Change the adapter from initiator to target:

```
system node hardware unified-connect modify -t target -node _node_name_
adapter _adapter_name_
```

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

```
network fcp adapter show -node _node_name_
```

5. Bring your adapter online:

```
network fcp adapter modify -node _node_name_ -adapter _adapter_port_
-state up
```

Configure FC adapter speed

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

About this task

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

Before you begin

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

Steps

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

```
network interface modify -vserver * -lif * { -home-node node1 -home-port
0c } -status-admin down
```

2. Take the adapter offline:

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

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

3. Determine the maximum speed for the port adapter:

```
fcp adapter show -instance
```

You cannot modify the adapter speed beyond the maximum speed.

4. Change the adapter speed:

```
network fcp adapter modify -node node1 -adapter 0c -speed 16
```

5. Bring the adapter online:

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

6. Bring all of the LIFs on the adapter online:

```
network interface modify -vserver * -lif * { -home-node node1 -home-port 0c } -status-admin up
```

ONTAP commands for managing FC adapters

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

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

Commands for managing FC target adapters

If you want to...	Use this command...
Display FC adapter information on a node	<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>

If you want to...	Use this command...
Verify which expansion cards are installed and whether there are any configuration errors	run -node <i>node_name</i> sysconfig -ac
View a man page for a command	man <i>command_name</i>

Commands for managing FC initiator adapters

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

Commands for managing onboard FC adapters

If you want to...	Use this command...
Display the status of the onboard FC ports	system node hardware unified-connect show

Related information

- [network fcp adapter](#)

Avoid connectivity loss to an ONTAP system using an X1133A-R6 adapter

You can prevent loss of connectivity during a port failure by configuring your system with redundant paths to separate X1133A-R6 HBAs.

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

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

FCoE configurations

Configure FCoE fabrics with ONTAP systems

FCoE can be configured in various ways using FCoE switches. Direct-attached configurations are not supported in FCoE.

All FCoE configurations are dual-fabric, fully redundant, and require host-side multipathing software. In all FCoE configurations, you can have multiple FCoE and FC switches in the path between the initiator and target, up to the maximum hop count limit. To connect switches to each other, the switches must run a firmware version that supports Ethernet ISLs. Each host in any FCoE configuration can be configured with a different operating system.


FCoE configurations require Ethernet switches that explicitly support FCoE features. FCoE configurations are validated through the same interoperability and quality assurance process as FC switches. Supported configurations are listed in the Interoperability Matrix. Some of the parameters included in these supported configurations are the switch model, the number of switches that can be deployed in a single fabric, and the supported switch firmware version.

The FC target expansion adapter port numbers in the illustrations are examples. The actual port numbers might vary, depending on the expansion slots in which the FCoE target expansion adapters are installed.

FCoE initiator to FC target


Using FCoE initiators (CNAs), you can connect hosts to both controllers in an HA pair through FCoE switches to FC target ports. The FCoE switch must also have FC ports. The host FCoE initiator always connects to the FCoE switch. The FCoE switch can connect directly to the FC target or can connect to the FC target through FC switches.

The following illustration shows host CNAs connecting to an FCoE switch, and then to an FC switch before connecting to the HA pair:




FCoE initiator to FCoE target

Using host FCoE initiators (CNAs), you can connect hosts to both controllers in an HA pair to FCoE target ports (also called UTAs or UTA2s) through FCoE switches.




FCoE initiator to FCoE and FC targets

Using host FCoE initiators (CNAs), you can connect hosts to both controllers in an HA pair to FCoE and FC target ports (also called UTAs or UTA2s) through FCoE switches.



FCoE mixed with IP storage protocols

Using host FCoE initiators (CNAs), you can connect hosts to both controllers in an HA pair to FCoE target ports (also called UTAs or UTA2s) through FCoE switches. FCoE ports cannot use traditional link aggregation to a single switch. Cisco switches support a special type of link aggregation (Virtual Port Channel) that does support FCoE. A Virtual Port Channel aggregates individual links to two switches. You can also use Virtual Port Channels for other Ethernet traffic. Ports used for traffic other than FCoE, including NFS, SMB, iSCSI, and other Ethernet traffic, can use regular Ethernet ports on the FCoE switches.



ONTAP supported FCoE initiator and target port combinations

Certain combinations of FCoE and traditional FC initiators and targets are supported.

FCoE initiators

You can use FCoE initiators in host computers with both FCoE and traditional FC targets in storage controllers. The host FCoE initiator must connect to an FCoE DCB (data center bridging) switch; direct connection to a target is not supported.

The following table lists the supported combinations:

Initiator	Target	Supported?
FC	FC	Yes
FC	FCoE	Yes
FCoE	FC	Yes
FCoE	FCoE	Yes

FCoE targets

You can mix FCoE target ports with 4-Gb, 8-Gb, or 16-Gb FC ports on the storage controller regardless of whether the FC ports are add-in target adapters or onboard ports. You can have both FCoE and FC target

adapters in the same storage controller.



The rules for combining onboard and expansion FC ports still apply.

FC and FCoE zoning

Learn about FC and FCoE zoning with ONTAP systems

An FC, FC-NVMe or FCoE zone is a logical grouping of one or more ports within a fabric. For devices to be able see each other, connect, create sessions with one another, and communicate, both ports must be members of the same zone.

Zoning increases security by limiting access and connectivity to end-points that share a common zone. Ports that are not in the same zone cannot communicate with one another. This reduces or eliminates *crosstalk* between initiator HBAs. Should connectivity issues occur, zoning helps to isolate problems to a specific set of ports, thereby decreasing time to resolution.

Zoning reduces the number of available paths to a particular port and reduces the number of paths between a host and the storage system. For example, some host OS multipathing solutions have a limit on the number of paths they can manage. Zoning can reduce the number of paths visible to the host so that paths to the host do not exceed the maximum allowed by the host OS.

World Wide Name-based zoning

Zoning based on World Wide Name (WWN) specifies the WWN of the members to be included within the zone. Although World Wide Node Name (WWNN) zoning is possible with some switch vendors, when zoning in ONTAP, you must use World Wide Port Name (WWPN) zoning.

WWPN zoning is required to properly define a specific port and to use NPIV effectively. FC switches should be zoned using the WWPNs of the target's logical interfaces (LIFs), not the WWPNs of the physical ports on the node. The WWPNs of the physical ports start with "50" and the WWPNs of the LIFs start with "20".

WWPN zoning provides flexibility because access is not determined by where the device is physically connected to the fabric. You can move a cable from one port to another without reconfiguring zones.

Recommended FC and FCoE zoning configurations for ONTAP systems

You should create a zoning configuration if your host does not have a multipathing solution installed, if four or more hosts are connected to your SAN or if Selective LUN Mapping is not implemented on the nodes in your cluster.

In the recommended FC and FCoE zoning configuration, each zone includes one initiator port and one or more target LIFs. This configuration allows each host initiator to access any node, while preventing hosts accessing the same node from seeing each other's ports.

Add all LIFs from the storage virtual machine (SVM) to the zone with the host initiator. This allows you to move volumes or LUNs without editing your existing zones or creating new zones.

Dual-fabric zoning configurations

Dual-fabric zoning configurations are recommended because they provide protection against data loss due to a single component failure. In a dual-fabric configuration, each host initiator is connected to each node in the

cluster using different switches. If one switch becomes unavailable, data access is maintained through the remaining switch. Multipathing software is required on the host to manage multiple paths.

In the following figure, the host has two initiators and is running multipathing software. There are two zones. [Selective LUN Mapping \(SLM\)](#) is configured so that all nodes are considered as reporting nodes.




The naming convention used in this figure is just a recommendation of one possible naming convention that you can choose to use for your ONTAP solution.

- Zone 1: HBA 0, LIF_1, LIF_3, LIF_5, and LIF_7
- Zone 2: HBA 1, LIF_2, LIF_4, LIF_6, and LIF_8

Each host initiator is zoned through a different switch. Zone 1 is accessed through Switch 1. Zone 2 is accessed through Switch 2.

Each host can access a LIF on every node. This enables the host to still access its LUNs if a node fails. SVMs have access to all iSCSI and FC LIFs on every node in the cluster based on your SLM reporting nodes configuration. You can use SLM, portsets, or FC switch zoning to reduce the number of paths from an SVM to the host and the number of paths from an SVM to a LUN.

If the configuration includes more nodes, the LIFs for the additional nodes are included in these zones..



The host operating system and multipathing software have to support the number of paths that is being used to access the LUNs on the nodes.

Single-fabric zoning

In a single-fabric configuration, you connect each host initiator to each storage node through a single switch. Single-fabric zoning configurations are not recommended because they do not provide protection against data loss due to a single component failure. If you choose to configure single-fabric zoning, each host should have two initiators for multipathing to provide resiliency in the solution. Multipathing software is required on the host to manage multiple paths.

Each host initiator should have a minimum of one LIF from each node that the initiator can access. The zoning should allow at least one path from the host initiator to the HA pair of nodes in the cluster to provide a path for

LUN connectivity. This means that each initiator on the host might only have one target LIF per node in its zone configuration. If there is a requirement for multipathing to the same node or multiple nodes in the cluster, then each node will have multiple LIFs per node in its zone configuration. This enables the host to still access its LUNs if a node fails or a volume containing the LUN is moved to a different node. This also requires the reporting nodes to be set appropriately.

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


In the following figure, the host has two initiators and is running multipathing software. There are two zones:



The naming convention used in this figure is just a recommendation of one possible naming convention that you can choose to use for your ONTAP solution.

- Zone 1: HBA 0, LIF_1, and LIF_3
- Zone 2: HBA 1, LIF_2, and LIF_4

If the configuration includes more nodes, the LIFs for the additional nodes are included in these zones.s.



In this example, you could also have all four LIFs in each zone. In that case, the zones would be as follows:

- Zone 1: HBA 0, LIF_1, LIF_2, LIF_3, and LIF_4
- Zone 2: HBA 1, LIF_1, LIF_2, LIF_3, and LIF_4



The host operating system and multipathing software have to support the number of supported paths that are being used to access the LUNs on the nodes. To determine the number of paths used to access the LUNs on nodes, see the SAN configuration limits section.

Zoning restrictions for Cisco FC and FCoE switches

When using Cisco FC and FCoE switches, certain restrictions apply to the use of physical ports and logical interfaces (LIFs) in zones.

Physical ports

- FC-NVMe and FC can share the same 32 Gb physical port
- FC-NVMe and FCoE cannot share the same physical port
- FC and FCoE can share the same physical port but their protocol LIFs must be in separate zones.

Logical Interfaces (LIFs)

- A zone can contain a LIF from every target port in the cluster.

Verify the SLM configuration so that you do not exceed the maximum number of paths allowed to the host.

- Each LIF on a given port must be in a separate zone from other LIFs on that port
- LIFs on different physical ports can be in the same zone.

Requirements for SAN hosts connected to ONTAP and non-NetApp systems

Shared SAN configurations are defined as hosts that are attached to both ONTAP storage systems and other vendors' storage systems. Accessing ONTAP storage systems and other vendors' storage systems from a single host is supported as long as several requirements are met.

For all of the host operating systems, it is a best practice to use separate adapters to connect to each vendor's storage systems. Using separate adapters reduces the chances of conflicting drivers and settings. For connections to an ONTAP storage system, the adapter model, BIOS, firmware, and driver must be listed as supported in the NetApp Interoperability Matrix Tool.

You should set the required or recommended timeout values and other storage parameters for the host. You must always install the NetApp software or apply the NetApp settings last.

- For AIX, you should apply the values from the AIX Host Utilities version that is listed in the Interoperability Matrix Tool for your configuration.
- For ESX, you should apply host settings by using Virtual Storage Console for VMware vSphere.
- For HP-UX, you should use the HP-UX default storage settings.
- For Linux, you should apply the values from the Linux Host Utilities version that is listed in the Interoperability Matrix Tool for your configuration.
- For Solaris, you should apply the values from the Solaris Host Utilities version that is listed in the Interoperability Matrix Tool for your configuration.
- For Windows, you should install the Windows Host Utilities version that is listed in the Interoperability Matrix Tool for your configuration.

Related information

[NetApp Interoperability Matrix Tool](#)

SAN configurations in a MetroCluster environment

Supported SAN configurations in an ONTAP MetroCluster environment

You must be aware of certain considerations when using SAN configurations in a MetroCluster environment.

- MetroCluster configurations do not support front-end FC fabric “routed” vSAN configurations.
- Beginning with ONTAP 9.15.1, four-node MetroCluster IP configurations are supported on NVMe/TCP.
- Beginning with ONTAP 9.12.1, four-node MetroCluster IP configurations are supported on NVMe/FC. MetroCluster configurations are not supported for front-end NVMe networks before ONTAP 9.12.1.
- Other SAN protocols such as iSCSI, FC, and FCoE are supported on MetroCluster configurations.
- When using SAN client configurations, you must check whether any special considerations for MetroCluster configurations are included in the notes that are provided in the [NetApp Interoperability Matrix Tool \(IMT\)](#).
- Operating systems and applications must provide an I/O resiliency of 120 seconds to support MetroCluster automatic unplanned switchover and Tiebreaker or Mediator-initiated switchover.
- MetroCluster configurations use the same WWNNs and WWPNs on both sides of the front-end FC fabric.

Related information

- [Understanding MetroCluster data protection and disaster recovery](#)
- [NetApp Knowledge Base: What are AIX Host support considerations in a MetroCluster configuration?](#)
- [NetApp Knowledge Base: Solaris host support considerations in a MetroCluster configuration](#)

Avoid port overlap during ONTAP MetroCluster switchover and switchback

In a SAN environment, you can configure the front-end switches to avoid overlap when the old port goes offline and the new port comes online.

During switchover, the FC port on the surviving site might log in to the fabric before the fabric has detected that the FC port on the disaster site is offline and has removed this port from the name and directory services.

If the FC port on the disaster is not yet removed, the fabric login attempt of the FC port at the surviving site might be rejected due to a duplicate WWPN. This behavior of the FC switches can be changed to honor the login of the previous device and not the existing one. You should verify the effects of this behavior on other fabric devices. Contact the switch vendor for more information.

Choose the correct procedure according to your switch type.

Example 1. Steps

Cisco switch

1. Connect to the switch and log in.
2. Enter configuration mode:

```
switch# config t  
switch(config) #
```

3. Overwrite the first device entry in the name server database with the new device:

```
switch(config) # no fcns reject-duplicate-pwwn vsan 1
```

4. In switches that are running NX-OS 8.x, confirm that the flogi quiesce timeout is set to zero:

- a. Display the quiesce timerval:

```
switch(config) # show flogi interval info \| i quiesce
```

```
Stats: fs flogi quiesce timerval: 0
```

- b. If the output in the previous step does not indicate that the timerval is zero, then set it to zero:

```
switch(config) # flogi scale enable
```

```
switch(config) $ flogi quiesce timeout 0
```

Brocade switch

1. Connect to the switch and log in.
2. Enter the `switchDisable` command.
3. Enter the `configure` command, and press **y** at the prompt.

```
F-Port login parameters (yes, y, no, n): [no] y
```

4. Choose setting 1:

```
- 0: First login take precedence over the second login (default)  
- 1: Second login overrides first login.  
- 2: the port type determines the behavior  
Enforce FLOGI/FDISC login: (0..2) [0] 1
```

5. Respond to the remaining prompts, or press **Ctrl + D**.

6. Enter the `switchEnable` command.

Related information

[Performing switchover for tests or maintenance](#)

ONTAP support for SAN host multipathing

ONTAP uses Asymmetric Logical Unit Access (ALUA) software for multipathing with both FC and iSCSI hosts.

Beginning with ONTAP 9.5 multipath high availability (HA) pair failover/giveback is supported for NVMe hosts using Asynchronous Namespace Access (ANA). In ONTAP 9.4, NVMe supports only one path from host to target, so the application host must manage path failover to its HA partner.

The multipathing software is required on your SAN host if it can access a LUN or NVMe namespace through more than one path. It presents a single disk to the operating system for all paths to a LUN or NVMe namespace. Without it, the operating system could treat each path as a separate disk, leading to data corruption.

Your solution is considered to have multiple paths if you have any of the following:

- A single initiator port in the host attaching to multiple SAN LIFs in the SVM
- Multiple initiator ports attaching to a single SAN LIF in the SVM
- Multiple initiator ports attaching to multiple SAN LIFs in the SVM

Multipathing software, also known as MPIO (multipath I/O) software, is recommended in HA configurations. In addition to Selective LUN Map, using FC switch zoning or portsets to limit the paths used to access LUNs is also recommended.

For information about which specific host configurations support ALUA or ANA, see the [NetApp Interoperability Matrix Tool](#) and [ONTAP SAN Host Configuration](#) for your host operating system.

Recommended number of paths from host to nodes in cluster

You should not exceed more than eight paths from your host to each node in your cluster. You should also not exceed the total number of paths that can be supported for the host OS and the multipathing used on the host.

You should have a minimum of two paths per LUN connecting to each reporting node through [Selective LUN Map \(SLM\)](#) being used by the storage virtual machine (SVM) in your cluster. This eliminates single points of failure and enables the system to survive component failures.

If you have four or more nodes in your cluster or more than four target ports being used by the SVMs in any of your nodes, you can use the following methods to limit the number of paths that can be used to access LUNs on your nodes so that you do not exceed the recommended maximum of eight paths.

- SLM

SLM reduces the number of paths from the host to LUN to only paths on the node owning the LUN and the owning node's HA partner. SLM is enabled by default.

- [Portsets for iSCSI](#)

- FC igroup mappings from your host
- FC switch zoning

Configuration limits

Determine the maximum supported nodes and SAN hosts per ONTAP cluster

The number of supported nodes per cluster varies depending on your version of ONTAP, your controller models, and the protocol of your cluster nodes. The maximum number of SAN hosts that can be connected to a cluster also varies based upon your specific configuration.

Determine the maximum supported nodes per cluster

If any node in the cluster is configured for FC, FC-NVMe, FCoE, or iSCSI, that cluster is limited to the SAN node limits. Node limits based on the controllers in your cluster are listed in the *Hardware Universe*.

Steps

1. Go to [NetApp Hardware Universe](#).
2. In the upper left, next to **Home**, select **Platforms**; then select the platform type.
3. Select your version of ONTAP.

A new column is displayed for you to choose your platforms.

4. Select the platforms used in your solution.
5. Under **Choose Your Specifications**, deselect **Select All**.
6. Select **Max Nodes per Cluster (NAS/SAN)**.
7. Click **Show Results**.

Results

The maximum nodes per cluster for your selected platforms is displayed.

Determine if your cluster can support more FC hosts

For FC and FC-NVMe configurations, you should use the number of initiator-target nexuses (ITNs) in your system to determine whether you can add more hosts to your cluster.

An ITN represents one path from the host's initiator to the storage system's target. The maximum number of ITNs per node in FC and FC-NVMe configurations is 2,048. If you are below the maximum number of ITNs, you can continue to add hosts to your cluster.

To determine the number of ITNs used in your cluster, perform the following steps for each node in the cluster.

Steps

1. Identify all the LIFs on a given node.
2. Run the following command for every LIF on the node:

```
fcp initiator show -fields wwpn, lif
```

The number of entries displayed at the bottom of the command output represents your number of ITNs for that LIF.

3. Record the number of ITNs displayed for each LIF.
4. Add the number of ITNs for each LIF on every node in your cluster.

This total represents the number of ITNs in your cluster.

Determine if your cluster can support more iSCSI hosts

The number of hosts that can be directly connected to a node or that can be connected through one or more switches depends on the number of available Ethernet ports. The number of available Ethernet ports is determined by the model of the controller and the number and type of adapters installed in the controller. The number of supported Ethernet ports for controllers and adapters is available in the *Hardware Universe*.

For all multi-node cluster configurations, you must determine the number of iSCSI sessions per node to know whether you can add more hosts to your cluster. As long as your cluster is below the maximum number of iSCSI sessions per node, you can continue to add hosts to your cluster. The maximum number of iSCSI sessions per node varies based on the types of controllers in your cluster.

Steps

1. Identify all of the target portal groups on the node.
2. Check the number of iSCSI sessions for every target portal group on the node:

```
iscsi session show -tpgroup _tpgroup_
```

The number of entries displayed at the bottom of the command output represents your number of iSCSI sessions for that target portal group.

3. Record the number of iSCSI sessions displayed for each target portal group.
4. Add the number of iSCSI sessions for each target portal group on the node.

The total represents the number of iSCSI sessions on your node.

All-Flash SAN Array configuration limits and support

All-Flash SAN Array (ASA) configuration limits and support varies by ONTAP version.

The most current details on supported configuration limits are available in [NetApp Hardware Universe](#).



These limitations apply to ASA systems. If you have an ASA r2 system (ASA A1K, ASA A90, ASA A70, ASA A50, ASA A30, ASA A20, or ASA C30), see [ASA r2 system storage limits](#).

SAN protocols and supported number of nodes per cluster

The supported SAN protocols and maximum number of nodes per cluster depends on whether you have a non-MetroCluster or MetroCluster configuration:

Non-MetroCluster configurations

The following table shows the ASA support for SAN protocols and the supported number of nodes per cluster in non-MetroCluster configurations:

Beginning with ONTAP...	Protocol support	Maximum nodes per cluster
9.11.1	<ul style="list-style-type: none">• NVMe/TCP• NVMe/FC	12
9.10.1	<ul style="list-style-type: none">• NVMe/TCP	2
9.9.1	<ul style="list-style-type: none">• NVMe/FC	2
	<ul style="list-style-type: none">• FC• iSCSI	12
9.7	<ul style="list-style-type: none">• FC• iSCSI	2

MetroCluster IP configurations

The following table shows the ASA support for SAN protocols and the supported number of nodes per cluster in MetroCluster IP configurations:

Beginning with ONTAP...	Protocol support	Maximum nodes per cluster
9.15.1	<ul style="list-style-type: none">• NVMe/TCP	2 nodes per cluster in four-node MetroCluster IP configurations
9.12.1	<ul style="list-style-type: none">• NVMe/FC	2 nodes per cluster in four-node MetroCluster IP configurations
9.9.1	<ul style="list-style-type: none">• FC• iSCSI	4 nodes per cluster in eight-node MetroCluster IP configurations
9.7	<ul style="list-style-type: none">• FC• iSCSI	2 nodes per cluster in four-node MetroCluster IP configurations

Support for persistent ports

Beginning with ONTAP 9.8, persistent ports are enabled by default on All-Flash SAN Arrays (ASAs) that are configured to use the FC protocol. Persistent ports are only available for FC and require zone membership

identified by World Wide Port Name (WWPN).

Persistent ports reduce the impact of takeovers by creating a shadow LIF on the corresponding physical port of the high-availability (HA) partner. When a node is taken over, the shadow LIF on the partner node assumes the identity of the original LIF, including the WWPN. Before the status of path to the taken over node is changed to faulty, the shadow LIF appears as an Active/Optimized path to the host MPIO stack, and I/O is shifted. This reduces I/O disruption because the host always sees the same number of paths to the target, even during storage failover operations.

For persistent ports, the following FCP port characteristics should be identical within the HA pair:

- FCP port counts
- FCP port names
- FCP port speeds
- FCP LIF WWPN-based zoning

If any of these characteristics are not identical within the HA pair, the following EMS message is generated:

```
EMS : scsiblade.lif.persistent.ports.fcp.init.error
```

For more information on persistent ports, see [NetApp Technical Report 4080: Best Practices for Modern SAN](#).

Configuration limits for FC switches used with ONTAP systems

Fibre Channel switches have maximum configuration limits, including the number of logins supported per port, port group, blade, and switch. The switch vendors document their supported limits.

Each FC logical interface (LIF) logs into an FC switch port. The total number of logins from a single target on the node equals the number of LIFs plus one login for the underlying physical port. Do not exceed the switch vendor's configuration limits for logins or other configuration values. This also holds true for the initiators being used on the host side in virtualized environments with NPIV enabled. Do not exceed the switch vendor's configuration limits for logins for either the target or the initiators being used in the solution.

Brocade switch limits

You can find the configuration limits for Brocade switches in the *Brocade Scalability Guidelines*.

Cisco Systems switch limits

You can find the configuration limits for Cisco switches in the [Cisco Configuration Limits](#) guide for your version of Cisco switch software.

Maximum FC and FCoE hop count supported in ONTAP

The hop count is defined as the number of switches in the path between the initiator (host) and target (storage system). The maximum supported FC hop count between a host and storage system varies depending on the switch supplier.

Documentation from Cisco Systems also refers to this value as the *diameter of the SAN fabric*.

For FCoE, you can have FCoE switches connected to FC switches. For end-to-end FCoE connections, the

FCoE switches must be running a firmware version that supports Ethernet inter-switch links (ISLs).

Switch supplier	Supported hop count
Brocade	<ul style="list-style-type: none">• 7 for FC• 5 for FCoE
Cisco	<ul style="list-style-type: none">• 7 for FC• Up to 3 of the switches can be FCoE switches.

Calculate queue depth for ONTAP FC hosts

You might need to tune your FC queue depth on the host to achieve the maximum values for ITNs per node and FC port fan-in. The maximum number of LUNs and the number of HBAs that can connect to an FC port are limited by the available queue depth on the FC target ports.

About this task

Queue depth is the number of I/O requests (SCSI commands) that can be queued at one time on a storage controller. Each I/O request from the host's initiator HBA to the storage controller's target adapter consumes a queue entry. Typically, a higher queue depth equates to better performance. However, if the storage controller's maximum queue depth is reached, that storage controller rejects incoming commands by returning a QFULL response to them. If a large number of hosts are accessing a storage controller, you should plan carefully to avoid QFULL conditions, which significantly degrade system performance and can lead to errors on some systems.

In a configuration with multiple initiators (hosts), all hosts should have similar queue depths. Because of the inequality in queue depth between hosts connected to the storage controller through the same target port, hosts with smaller queue depths are being deprived of access to resources by hosts with larger queue depths.

The following general recommendations can be made about "tuning" queue depths:

- For small to mid-size systems, use an HBA queue depth of 32.
- For large systems, use an HBA queue depth of 128.
- For exception cases or performance testing, use a queue depth of 256 to avoid possible queuing problems.
- All hosts should have the queue depths set to similar values to give equal access to all hosts.
- To avoid performance penalties or errors, the storage controller target FC port queue depth must not be exceeded.

Steps

1. Count the total number of FC initiators in all of the hosts that connect to one FC target port.
2. Multiply by 128.
 - If the result is less than 2,048, set the queue depth for all initiators to 128.
You have 15 hosts with one initiator connected to each of two target ports on the storage controller. $15 \times 128 = 1,920$. Because 1,920 is less than the total queue depth limit of 2,048, you can set the queue depth for all of your initiators to 128.
 - If the result is greater than 2,048, go to step 3.

You have 30 hosts with one initiator connected to each of two target ports on the storage controller. $30 \times 128 = 3,840$. Because 3,840 is greater than the total queue depth limit of 2,048, you should choose one of the options under step 3 for remediation.

3. Choose one of the following options to add more hosts to the storage controller.

◦ Option 1:

- i. Add more FC target ports.
- ii. Redistribute your FC initiators.
- iii. Repeat steps 1 and 2.

The desired queue depth of 3,840 exceeds the available queue depth per port. To remedy this, you can add a two-port FC target adapter to each controller, then rezone your FC switches so that 15 of your 30 hosts connect to one set of ports, and the remaining 15 hosts connect to a second set of ports. The queue depth per port is then reduced to $15 \times 128 = 1,920$.

◦ Option 2:

- i. Designate each host as “large” or “small” based on its expected I/O need.
- ii. Multiply the number of large initiators by 128.
- iii. Multiply the number of small initiators by 32.
- iv. Add the two results together.
- v. If the result is less than 2,048, set the queue depth for large hosts to 128 and the queue depth for small hosts to 32.
- vi. If the result is still greater than 2,048 per port, reduce the queue depth per initiator until the total queue depth is less than or equal to 2,048.

To estimate the queue depth needed to achieve a certain I/O per second throughput, use this formula:



Needed queue depth = (Number of I/O per second) × (Response time)

For example, if you need 40,000 I/O per second with a response time of 3 milliseconds, the needed queue depth = $40,000 \times (.003) = 120$.

The maximum number of hosts that you can connect to a target port is 64, if you decide to limit the queue depth to the basic recommendation of 32. However, if you decide to have a queue depth of 128, then you can have a maximum of 16 hosts connected to one target port. The larger the queue depth, the fewer hosts that a single target port can support. If your requirement is such that you cannot compromise on the queue depth, then you should get more target ports.

The desired queue depth of 3,840 exceeds the available queue depth per port. You have 10 “large” hosts that have high storage I/O needs, and 20 “small” hosts that have low I/O needs. Set the initiator queue depth on the large hosts to 128 and the initiator queue depth on the small hosts to 32.

Your resulting total queue depth is $(10 \times 128) + (20 \times 32) = 1,920$.

You can spread the available queue depth equally across each initiator.

Your resulting queue depth per initiator is $2,048 \div 30 = 68$.

Modify queue depths for ONTAP SAN hosts

You might need to change the queue depths on your host to achieve the maximum values for ITNs per node and FC port fan-in. You can [calculate the optimal queue depth](#) for your environment.

AIX hosts

You can change the queue depth on AIX hosts using the `chdev` command. Changes made using the `chdev` command persist across reboots.

Examples:

- To change the queue depth for the `hdisk7` device, use the following command:

```
chdev -l hdisk7 -a queue_depth=32
```

- To change the queue depth for the `fcs0` HBA, use the following command:

```
chdev -l fcs0 -a num_cmd_elems=128
```

The default value for `num_cmd_elems` is 200. The maximum value is 2,048.



It might be necessary to take the HBA offline to change `num_cmd_elems` and then bring it back online using the `rmdev -l fcs0 -R` and `makdev -l fcs0 -P` commands.

HP-UX hosts

You can change the LUN or device queue depth on HP-UX hosts using the kernel parameter `scsi_max_qdepth`. You can change the HBA queue depth using the kernel parameter `max_fcp_reqs`.

- The default value for `scsi_max_qdepth` is 8. The maximum value is 255.

`scsi_max_qdepth` can be dynamically changed on a running system using the `-u` option on the `kmtune` command. The change will be effective for all devices on the system. For example, use the following command to increase the LUN queue depth to 64:

```
kmtune -u -s scsi_max_qdepth=64
```

It is possible to change queue depth for individual device files using the `scsictl` command. Changes using the `scsictl` command are not persistent across system reboots. To view and change the queue depth for a particular device file, execute the following command:

```
scsictl -a /dev/rdsck/c2t2d0
```

```
scsictl -m queue_depth=16 /dev/rdsck/c2t2d0
```

- The default value for `max_fcp_reqs` is 512. The maximum value is 1024.

The kernel must be rebuilt and the system must be rebooted for changes to `max_fcp_reqs` to take effect. To change the HBA queue depth to 256, for example, use the following command:

```
kmtune -u -s max_fcp_reqs=256
```

Solaris hosts

You can set the LUN and HBA queue depth for your Solaris hosts.

- For LUN queue depth: The number of LUNs in use on a host multiplied by the per-LUN throttle (lun-queue-depth) must be less than or equal to the tgt-queue-depth value on the host.
- For queue depth in a Sun stack: The native drivers do not allow for per LUN or per target max_throttle settings at the HBA level. The recommended method for setting the max_throttle value for native drivers is on a per-device type (VID_PID) level in the /kernel/drv/sd.conf and /kernel/drv/ssd.conf files. The host utility sets this value to 64 for MPxIO configurations and 8 for Veritas DMP configurations.

Steps

1. # cd/kernel/drv
2. # vi lpfc.conf
3. Search for /tft-queue (/tgt-queue)

```
tgt-queue-depth=32
```



The default value is set to 32 at installation.

4. Set the desired value based on the configuration of your environment.
5. Save the file.
6. Reboot the host using the sync; sync; sync; reboot -- -r command.

VMware hosts for a QLogic HBA

Use the esxcfg-module command to change the HBA timeout settings. Manually updating the esx.conf file is not recommended.

Steps

1. Log on to the service console as the root user.
2. Use the #vmkload_mod -l command to verify which Qlogic HBA module is currently loaded.
3. For a single instance of a Qlogic HBA, run the following command:

```
#esxcfg-module -s ql2xmaxqdepth=64 qla2300_707
```



This example uses qla2300_707 module. Use the appropriate module based on the output of vmkload_mod -l.

4. Save your changes using the following command:

```
#/usr/sbin/esxcfg-boot -b
```

5. Reboot the server using the following command:

```
#reboot
```

6. Confirm the changes using the following commands:

- a. #esxcfg-module -g qla2300_707
- b. qla2300_707 enabled = 1 options = 'ql2xmaxqdepth=64'

VMware hosts for an Emulex HBA

Use the `esxcfg-module` command to change the HBA timeout settings. Manually updating the `esx.conf` file is not recommended.

Steps

1. Log on to the service console as the root user.
2. Use the `#vmkload_mod -l grep lpfc` command to verify which Emulex HBA is currently loaded.
3. For a single instance of an Emulex HBA, enter the following command:

```
#esxcfg-module -s lpfc0_lun_queue_depth=16 lpfcdd_7xx
```



Depending on the model of the HBA, the module can be either `lpfcdd_7xx` or `lpfcdd_732`. The above command uses the `lpfcdd_7xx` module. You should use the appropriate module based on the outcome of `vmkload_mod -l`.

Running this command will set the LUN queue depth to 16 for the HBA represented by `lpfc0`.

4. For multiple instances of an Emulex HBA, run the following command:

```
a esxcfg-module -s "lpfc0_lun_queue_depth=16 lpfc1_lun_queue_depth=16"  
lpfcdd_7xx
```

The LUN queue depth for `lpfc0` and the LUN queue depth for `lpfc1` is set to 16.

5. Enter the following command:

```
#esxcfg-boot -b
```

6. Reboot using `#reboot`.

Windows hosts for an Emulex HBA

On Windows hosts, you can use the `LPUTILNT` utility to update the queue depth for Emulex HBAs.

Steps

1. Run the `LPUTILNT` utility located in the `C:\WINNT\system32` directory.
2. Select **Drive Parameters** from the menu on the right side.
3. Scroll down and double-click **QueueDepth**.

If you are setting **QueueDepth** greater than 150, the following Windows Registry value also need to be increased appropriately:



HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\lpxnds\Parameters\Device\NumberOfRequests

Windows hosts for a Qlogic HBA

On Windows hosts, you can use the **and the SANsurfer HBA manager utility** to update the queue depths for Qlogic HBAs.

Steps

1. Run the SANsurfer HBA manager utility.
2. Click on **HBA port > Settings**.
3. Click **Advanced HBA port settings** in the list box.
4. Update the **Execution Throttle** parameter.

Linux hosts for Emulex HBA

You can update the queue depths of an Emulex HBA on a Linux host. To make the updates persistent across reboots, you must then create a new RAM disk image and reboot the host.

Steps

1. Identify the queue depth parameters to be modified:

```
modinfo lpfc|grep queue_depth
```

The list of queue depth parameters with their description is displayed. Depending on your operating system version, you can modify one or more of the following queue depth parameters:

- **lpfc_lun_queue_depth**: Maximum number of FC commands that can be queued to a specific LUN (uint)
- **lpfc_hba_queue_depth**: Maximum number of FC commands that can be queued to an lpfc HBA (uint)
- **lpfc_tgt_queue_depth**: Maximum number of FC commands that can be queued to a specific target port (uint)

The **lpfc_tgt_queue_depth** parameter is applicable only for Red Hat Enterprise Linux 7.x systems, SUSE Linux Enterprise Server 11 SP4 systems and 12.x systems.

2. Update the queue depths by adding the queue depth parameters to the **/etc/modprobe.conf** file for a Red Hat Enterprise Linux 5.x system and to the **/etc/modprobe.d/scsi.conf** file for a Red Hat Enterprise Linux 6.x or 7.x system, or a SUSE Linux Enterprise Server 11.x or 12.x system.

Depending on your operating system version, you can add one or more of the following commands:

- `options lpfc lpfc_hba_queue_depth=new_queue_depth`
- `options lpfc lpfc_lun_queue_depth=new_queue_depth`

- options lpfc_tgt_queue_depth=new_queue_depth

3. Create a new RAM disk image, and then reboot the host to make the updates persistent across reboots.

For more information, see the [System administration](#) for your version of Linux operating system.

4. Verify that the queue depth values are updated for each of the queue depth parameter that you have modified:

```
cat /sys/class/scsi_host/host_number/lpfc_lun_queue_depth
cat /sys/class/scsi_host/host_number/lpfc_tgt_queue_depth
cat /sys/class/scsi_host/host_number/lpfc_hba_queue_depth
```

```
root@localhost ~]#cat /sys/class/scsi_host/host5/lpfc_lun_queue_depth
30
```

The current value of the queue depth is displayed.

Linux hosts for QLogic HBA

You can update the device queue depth of a QLogic driver on a Linux host. To make the updates persistent across reboots, you must then create a new RAM disk image and reboot the host. You can use the QLogic HBA management GUI or command-line interface (CLI) to modify the QLogic HBA queue depth.

This task shows how to use the QLogic HBA CLI to modify the QLogic HBA queue depth

Steps

1. Identify the device queue depth parameter to be modified:

```
modinfo qla2xxx | grep ql2xmaxqdepth
```

You can modify only the ql2xmaxqdepth queue depth parameter, which denotes the maximum queue depth that can be set for each LUN. The default value is 64 for RHEL 7.5 and later. The default value is 32 for RHEL 7.4 and earlier.

```
root@localhost ~]# modinfo qla2xxx|grep ql2xmaxqdepth
parm: ql2xmaxqdepth:Maximum queue depth to set for each LUN.
Default is 64. (int)
```

2. Update the device queue depth value:

- If you want to make the modifications persistent, perform the following steps:
 - Update the queue depths by adding the queue depth parameter to the /etc/modprobe.conf file for a Red Hat Enterprise Linux 5.x system and to the /etc/modprobe.d/scsi.conf file for a Red Hat Enterprise Linux 6.x or 7.x system, or a SUSE Linux Enterprise Server 11.x or 12.x system: options qla2xxx ql2xmaxqdepth=new_queue_depth
 - Create a new RAM disk image, and then reboot the host to make the updates persistent across reboots.

For more information, see the [System administration](#) for your version of Linux operating system.

- If you want to modify the parameter only for the current session, run the following command:

```
echo new_queue_depth > /sys/module/qla2xxx/parameters/ql2xmaxqdepth
```

In the following example, the queue depth is set to 128.

```
echo 128 > /sys/module/qla2xxx/parameters/ql2xmaxqdepth
```

3. Verify that the queue depth values are updated:

```
cat /sys/module/qla2xxx/parameters/ql2xmaxqdepth
```

The current value of the queue depth is displayed.

4. Modify the QLogic HBA queue depth by updating the firmware parameter `Execution Throttle` from the QLogic HBA BIOS.

- a. Log in to the QLogic HBA management CLI:

```
/opt/QLogic_Corporation/QConvergeConsoleCLI/qaucli
```

- b. From the main menu, select the `Adapter Configuration` option.

```

[root@localhost ~]#
/opt/QLogic_Corporation/QConvergeConsoleCLI/qaucli
Using config file:
/opt/QLogic_Corporation/QConvergeConsoleCLI/qaucli.cfg
Installation directory: /opt/QLogic_Corporation/QConvergeConsoleCLI
Working dir: /root

QConvergeConsole

      CLI - Version 2.2.0 (Build 15)

Main Menu

1: Adapter Information
**2: Adapter Configuration**
3: Adapter Updates
4: Adapter Diagnostics
5: Monitoring
6: FabricCache CLI
7: Refresh
8: Help
9: Exit

Please Enter Selection: 2

```

c. From the list of adapter configuration parameters, select the HBA Parameters option.

```

1: Adapter Alias
2: Adapter Port Alias
**3: HBA Parameters**
4: Persistent Names (udev)
5: Boot Devices Configuration
6: Virtual Ports (NPIV)
7: Target Link Speed (iidMA)
8: Export (Save) Configuration
9: Generate Reports
10: Personality
11: FEC
(p or 0: Previous Menu; m or 98: Main Menu; ex or 99: Quit)

Please Enter Selection: 3

```

d. From the list of HBA ports, select the required HBA port.

Fibre Channel Adapter Configuration

```
HBA Model QLE2562 SN: BFD1524C78510
  1: Port    1: WWPN: 21-00-00-24-FF-8D-98-E0 Online
  2: Port    2: WWPN: 21-00-00-24-FF-8D-98-E1 Online
HBA Model QLE2672 SN: RFE1241G81915
  3: Port    1: WWPN: 21-00-00-0E-1E-09-B7-62 Online
  4: Port    2: WWPN: 21-00-00-0E-1E-09-B7-63 Online
```

```
(p or 0: Previous Menu; m or 98: Main Menu; ex or 99: Quit)
Please Enter Selection: 1
```

The details of the HBA port are displayed.

e. From the HBA Parameters menu, select the Display HBA Parameters option to view the current value of the Execution Throttle option.

The default value of the Execution Throttle option is 65535.

HBA Parameters Menu

```
=====
HBA          : 2 Port: 1
SN           : BFD1524C78510
HBA Model    : QLE2562
HBA Desc.    : QLE2562 PCI Express to 8Gb FC Dual Channel
FW Version   : 8.01.02
WWPN         : 21-00-00-24-FF-8D-98-E0
WWNN         : 20-00-00-24-FF-8D-98-E0
Link         : Online
=====
```

```
1: Display HBA Parameters
2: Configure HBA Parameters
3: Restore Defaults
```

```
(p or 0: Previous Menu; m or 98: Main Menu; x or 99: Quit)
Please Enter Selection: 1
```

```
-----
HBA Instance 2: QLE2562 Port 1 WWPN 21-00-00-24-FF-8D-98-E0 PortID
03-07-00
Link: Online
```

```
-----  
Connection Options : 2 - Loop Preferred, Otherwise Point-  
to-Point  
Data Rate : Auto  
Frame Size : 2048  
Hard Loop ID : 0  
Loop Reset Delay (seconds) : 5  
Enable Host HBA BIOS : Enabled  
Enable Hard Loop ID : Disabled  
Enable FC Tape Support : Enabled  
Operation Mode : 0 - Interrupt for every I/O  
completion  
Interrupt Delay Timer (100us) : 0  
**Execution Throttle : 65535**  
Login Retry Count : 8  
Port Down Retry Count : 30  
Enable LIP Full Login : Enabled  
Link Down Timeout (seconds) : 30  
Enable Target Reset : Enabled  
LUNs Per Target : 128  
Out Of Order Frame Assembly : Disabled  
Enable LR Ext. Credits : Disabled  
Enable Fabric Assigned WWN : N/A
```

Press <Enter> to continue:

- f. Press **Enter** to continue.
- g. From the HBA Parameters menu, select the **Configure HBA Parameters** option to modify the HBA parameters.
- h. From the **Configure Parameters** menu, select the **Execute Throttle** option and update the value of this parameter.

Configure Parameters Menu

```
=====
HBA          : 2 Port: 1
SN           : BFD1524C78510
HBA Model    : QLE2562
HBA Desc.    : QLE2562 PCI Express to 8Gb FC Dual Channel
FW Version   : 8.01.02
WWPN         : 21-00-00-24-FF-8D-98-E0
WWNN         : 20-00-00-24-FF-8D-98-E0
Link         : Online
=====
```

- 1: Connection Options
- 2: Data Rate
- 3: Frame Size
- 4: Enable HBA Hard Loop ID
- 5: Hard Loop ID
- 6: Loop Reset Delay (seconds)
- 7: Enable BIOS
- 8: Enable Fibre Channel Tape Support
- 9: Operation Mode
- 10: Interrupt Delay Timer (100 microseconds)
- 11: Execution Throttle
- 12: Login Retry Count
- 13: Port Down Retry Count
- 14: Enable LIP Full Login
- 15: Link Down Timeout (seconds)
- 16: Enable Target Reset
- 17: LUNs per Target
- 18: Enable Receive Out Of Order Frame
- 19: Enable LR Ext. Credits
- 20: Commit Changes
- 21: Abort Changes

(p or 0: Previous Menu; m or 98: Main Menu; x or 99: Quit)

Please Enter Selection: 11

Enter Execution Throttle [1-65535] [65535]: 65500

- i. Press **Enter** to continue.
- j. From the Configure Parameters menu, select the **Commit Changes** option to save the changes.
- k. Exit the menu.

Copyright information

Copyright © 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.