

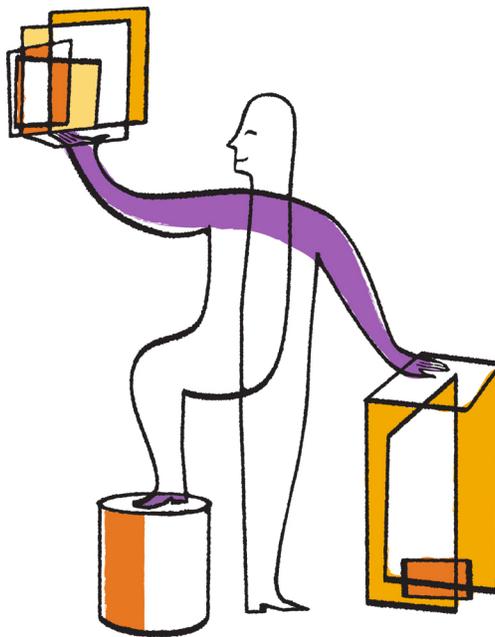


**NetApp**<sup>®</sup>

Updated for 8.3.1

## Clustered Data ONTAP<sup>®</sup> 8.3

### SAN Configuration Guide



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

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

Part number: 215-10114\_A0  
June 2015



# Contents

<b>Considerations for iSCSI configurations .....</b>	<b>5</b>
Ways to configure iSCSI SAN hosts with single nodes .....	5
Ways to configure iSCSI SAN hosts with HA pairs .....	7
Benefits of using VLANs in iSCSI configurations .....	8
Static VLANs .....	8
Dynamic VLANs .....	9
<b>Considerations for FC configurations .....</b>	<b>10</b>
Ways to configure FC SAN hosts with single nodes .....	10
Ways to configure FC with HA pairs .....	12
FC switch configuration best practices .....	13
Supported number of FC hop counts .....	13
Supported FC ports .....	14
How to prevent loss of connectivity when using the X1133A-R6 adapter ...	14
Port configuration options for the X1143A-R6 adapter .....	15
FC supported speeds .....	15
FC Target port configuration recommendations .....	15
<b>Ways to Configure FCoE .....</b>	<b>17</b>
FCoE initiator and target combinations .....	20
FCoE supported hop count .....	21
<b>Fibre Channel and FCoE zoning .....</b>	<b>22</b>
World Wide Name-based zoning .....	22
Individual zones .....	23
Single-fabric zoning .....	23
Dual-fabric HA pair zoning .....	24
FC and FCoE LIFs on the same port need to be in separate zones .....	25
<b>Requirements for shared SAN configurations .....</b>	<b>26</b>
<b>ALUA Configurations .....</b>	<b>27</b>
When host multipathing software is required .....	27
Recommended number of paths from host to nodes in cluster .....	27
<b>Configuration limits for FC, FCoE, and iSCSI configurations .....</b>	<b>29</b>
SAN configuration requirement for FlexVol volumes .....	29
SAN configuration limit parameters and definitions .....	29
Determining the number of supported nodes for SAN configurations .....	31
Determining the number of supported hosts per cluster in FC configurations .....	31
Determining the supported number of hosts in iSCSI configurations .....	32
Host operating system limits for SAN configurations .....	33
SAN configuration limits .....	34
SAN configuration limits for Data ONTAP-v platforms .....	42
FC switch configuration limits .....	42
Calculating queue depth .....	43
Setting queue depths on AIX hosts .....	45

Setting queue depths on HP-UX hosts .....	45
Setting queue depths on Solaris hosts .....	45
Setting queue depths on VMware hosts .....	46
Setting queue depths on Windows hosts .....	47
<b>Copyright information .....</b>	<b>48</b>
<b>Trademark information .....</b>	<b>49</b>
<b>How to send comments about documentation and receive update notifications .....</b>	<b>50</b>
<b>Index .....</b>	<b>51</b>

## Considerations for iSCSI configurations

You should be mindful of several things when setting up your iSCSI configuration.

- You can set up your iSCSI configuration with single-nodes or with HA pairs.
- You need to create one or more iSCSI paths from each storage controller, using logical interfaces (LIFs) to access a given LUN.  
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.
- VLANs offer specific benefits, such as increased security and improved network reliability that you might want to leverage in iSCSI.

### Related tasks

[Determining the number of supported nodes for SAN configurations](#) on page 31

[Determining the supported number of hosts in iSCSI configurations](#) on page 32

### Related references

[FCoE supported hop count](#) on page 21

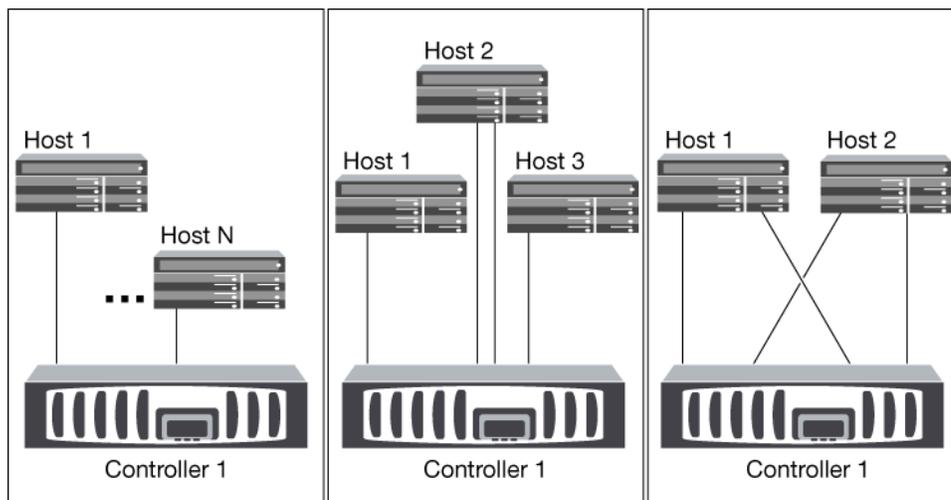
## Ways to configure iSCSI SAN hosts with single nodes

You can attach iSCSI SAN hosts to a single node through one or more IP switches, or you can attach iSCSI SAN hosts directly to the node without using a switch.

You can configure iSCSI SAN hosts with single nodes in the following ways. If there are multiple hosts connecting to the node, each host can be configured with a different operating system. For single and multi-network configurations, the node can have multiple iSCSI connections to the switch, but multipathing software is required.

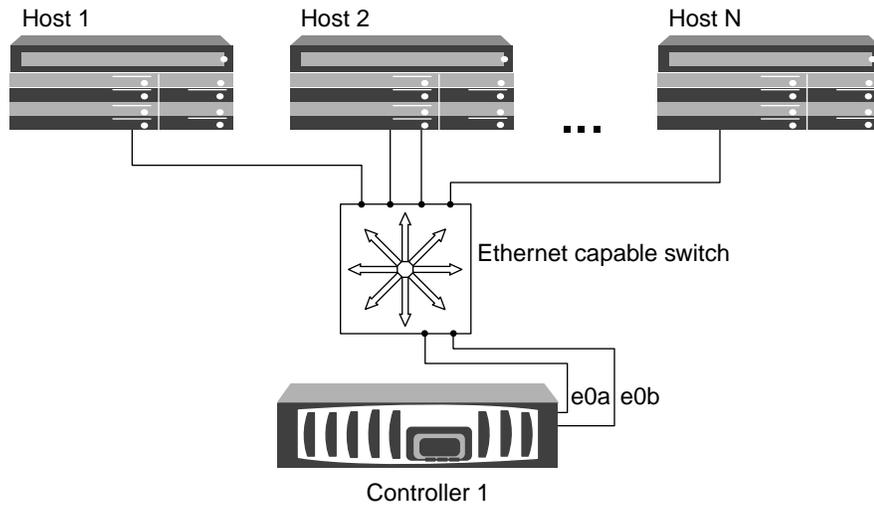
### Direct-attached single-node configurations

In direct-attached configurations, one or more hosts are directly connected to the node. Direct-attached configurations are only supported using single nodes.



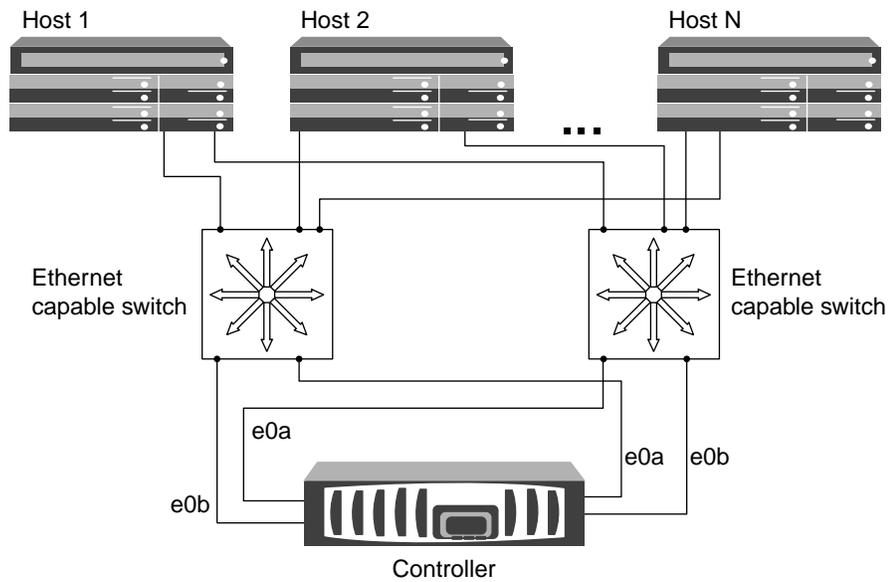
### Single-network single-node configurations

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



### Multi-network single-node configurations

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



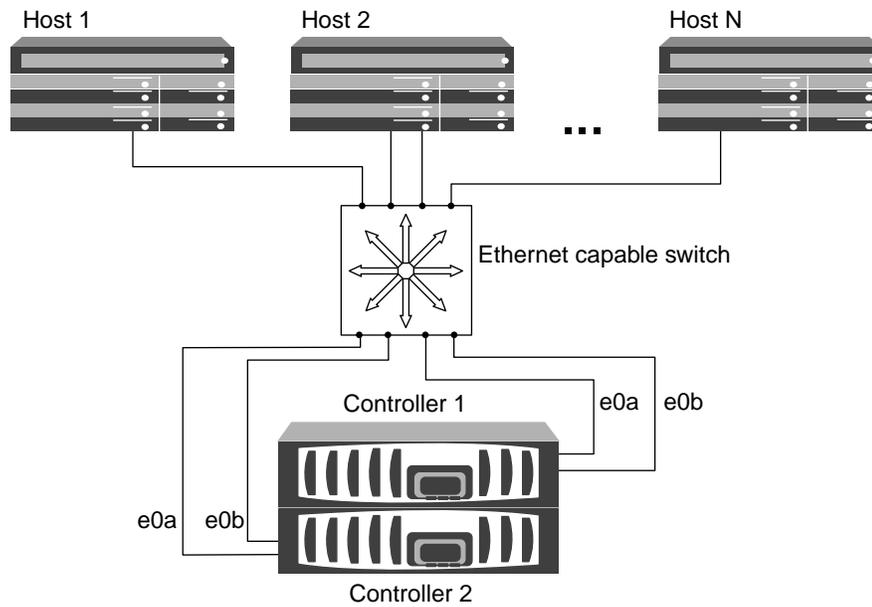
## Ways to configure iSCSI SAN hosts with HA pairs

You can configure iSCSI SAN hosts to connect to HA pairs through one or more IP switches. You cannot directly attach iSCSI SAN hosts to HA pairs without using a switch.

You can configure iSCSI SAN hosts with single network HA pairs or multi-network HA pairs. HA pairs can have multiple iSCSI connections to each switch, but multipathing software is required. If there are multiple hosts, you can configure each one with a different operating system.

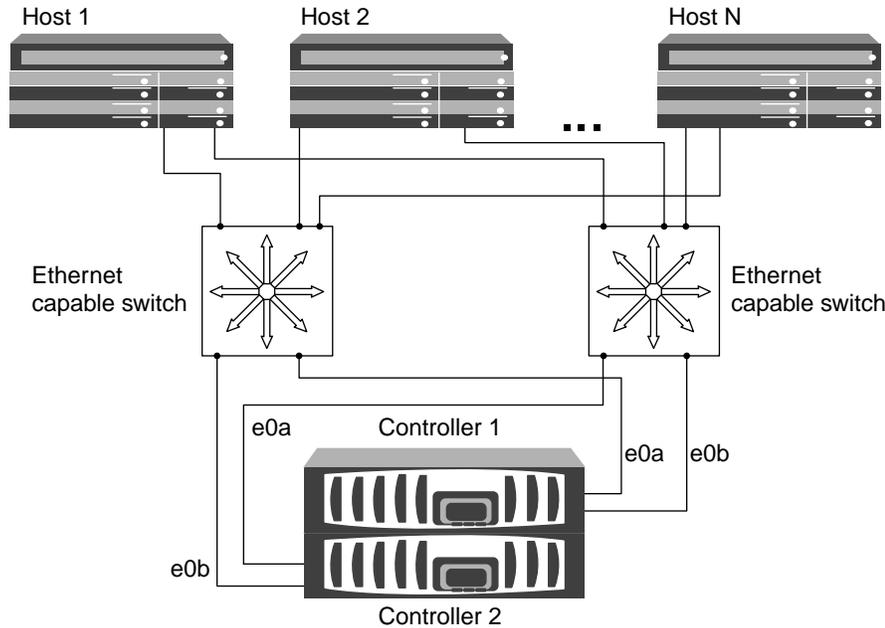
### Single network HA pairs

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.



### Multi-network HA pairs

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.



## Benefits of using VLANs 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.

### 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.

## **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.

## Considerations for FC configurations

---

You should be aware of several things when setting up your FC configuration.

- You can set up your FC configuration with single nodes or HA pairs using a single fabric or multifabric.
- Multiple hosts, using different operating systems, such as Windows, Linux, or UNIX, can access the storage solution 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.
- HA pairs with multiple, physically independent storage fabrics (minimum of two) are recommended for SAN solutions.  
This provides redundancy at the fabric and storage system layers. Redundancy is particularly important because these layers typically support many hosts.
- 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.
- Cascade, mesh, and core-edge 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.

### Related information

[NetApp Interoperability](#)

## Ways to configure FC SAN hosts with single nodes

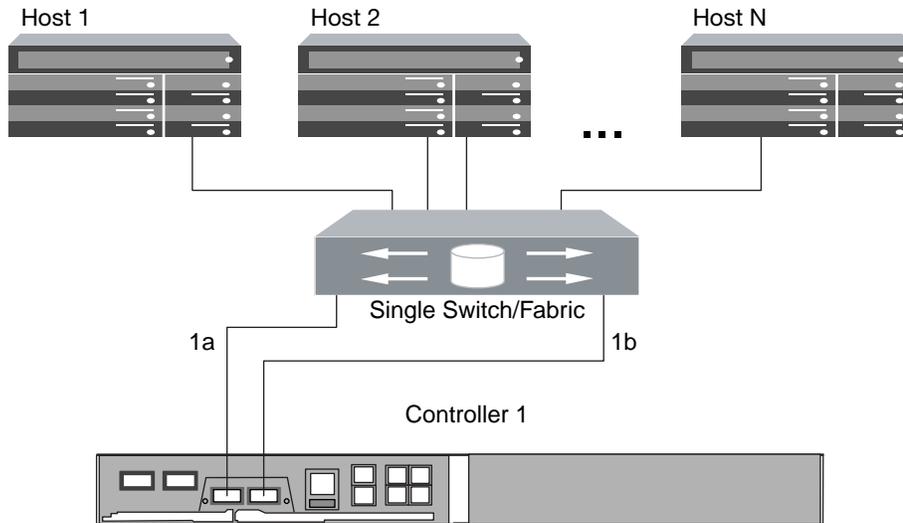
You can configure FC SAN hosts with single nodes through one or more fabrics. You cannot directly attach FC SAN hosts to single nodes without using a switch.

You can configure FC SAN hosts with single nodes through a single fabric or multifabrics. The FC target ports (0a, 0c, 0b, 0d) 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 single-node configurations

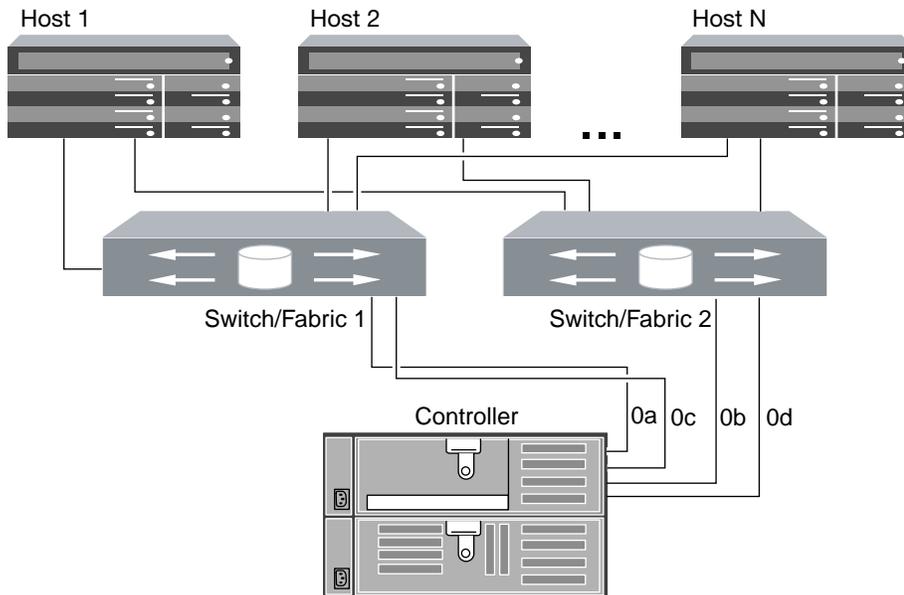
In single-fabric single-node configurations, there is one switch connecting a single-node to one or more hosts. Because there is a singular switch, this configuration is not fully redundant. All hardware platforms that support FC support single-fabric single-node configurations. However, the FAS2240 platform requires the X1150A-R6 expansion adapter to support a single-fabric single-node configuration.

The following figure shows a FAS2240 single-fabric single node configuration. It shows the storage controllers side by side, which is how they are mounted in the FAS2240-2. For the FAS2240-4, the controllers are mounted one above the other. There is no difference in the SAN configuration for the two models.



**Multifabric single-node configurations**

In multifabric single-node configurations, there are two or more switches connecting a single-node to one or more hosts. For simplicity, the following illustration shows a multifabric single-node configuration with only two fabrics, but you can have two or more fabrics in any multifabric configuration. In this illustration, the storage controller is mounted in the top chassis and the bottom chassis is empty.



## Ways to configure FC with HA pairs

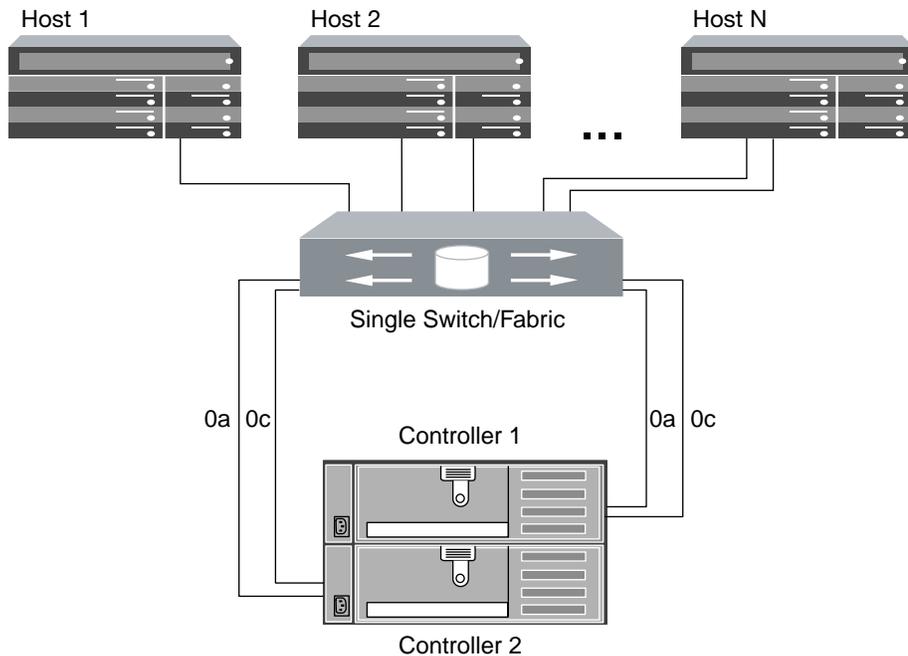
You can configure FC SAN hosts to connect to HA pairs through one or more fabrics. You cannot directly attach FC SAN hosts to HA pairs without using a switch.

You can configure FC SAN hosts with single fabric HA pairs or with multifabric HA pairs. The FC target port numbers (0a, 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 HA pairs

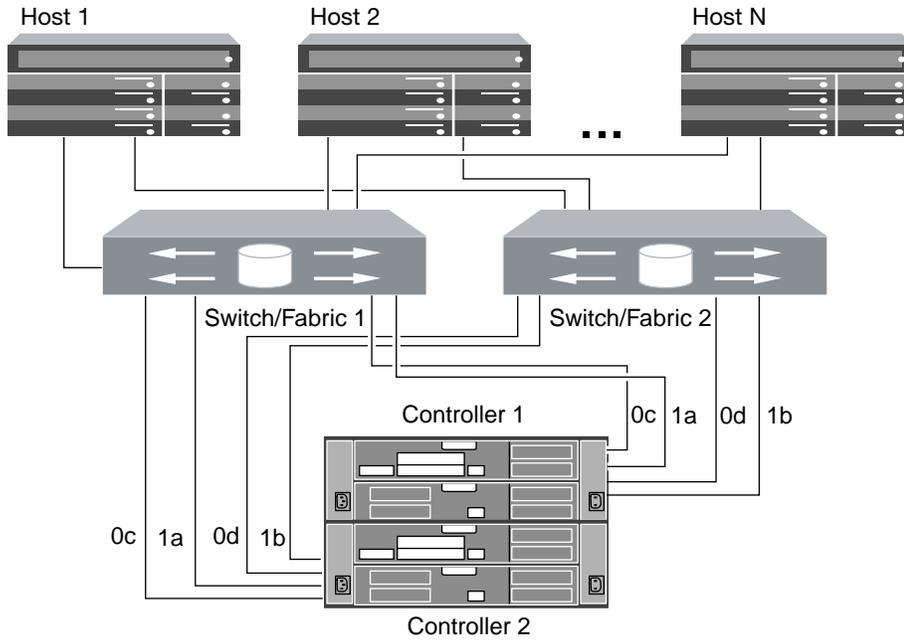
In single-fabric HA pair configurations, there is one fabric connecting both controllers in the HA pair to one or more hosts. Because the hosts and controllers are connected through a single switch, single-fabric HA pairs are not fully redundant.

All platforms that support FC configurations, support single-fabric HA pair configurations, except the FAS2240 platform. The FAS2240 platform only supports single-fabric single-node configurations.



### Multifabric HA pairs

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



## FC switch configuration best practices

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 create significant time savings. Although autonegotiation provides the greatest flexibility, it does not always perform as expected and it adds time to the overall fabric-build sequence .

All switches connected to the fabric have to support N\_Port ID virtualization (NPIV) and have it enabled. Clustered Data ONTAP uses NPIV to present FC targets to a fabric. For details about which environments are supported, see the Interoperability Matrix. For FC and iSCSI best practices, see TR-4080: *Best Practices for Scalable SAN in Clustered Data ONTAP 8.2*.

**Note:** Where supported, it works best to set the switch port topology to F (point-to-point).

### Related information

[NetApp Interoperability](#)

[NetApp Technical Report 4080: Best Practices for Scalable SAN in Clustered Data ONTAP 8.2](#)

## Supported number of FC hop counts

The maximum supported FC hop count between a host and storage system depends on the switch supplier and storage system support for FC configurations.

The hop count is defined as the number of switches in the path between the initiator (host) and target (storage system). Cisco also refers to this value as the *diameter of the SAN fabric*.

Switch supplier	Supported hop count
Brocade	7 for FC 5 for FCoE

Switch supplier	Supported hop count
Cisco	7 for FC Up to 3 of the switches can be FCoE switches.

**Related information**

[NetApp Downloads: Brocade Scalability Matrix Documents](#)

[NetApp Downloads: Cisco Scalability Matrix Documents](#)

## Supported FC ports

The number of onboard FC ports and CNA/UTA2 ports configured for FC varies based on the model of the controller. FC ports are also available through supported FC target expansion adapters.

**Onboard FC ports**

- Onboard ports can be individually configured as either target or initiator FC ports.
- The number of onboard FC ports differs depending on controller model.  
A complete list of onboard FC ports on each controller model is available from the *Hardware Universe*.
- FC ports are only available on FAS2240 systems through the X1150A-R6 expansion adapter. FAS2220 and FAS2520 systems do not support FC.

**Target expansion adapter FC ports**

- Available target expansion adapters differ depending on controller model.  
A complete list of target expansion adapters for each controller model is available from the *Hardware Universe*.
- Except for the expansion adapter models listed in the table, the ports on FC expansion adapters are configured as initiators or targets at the factory and cannot be changed.  
The ports on the following expansion adapters can be individually configured as either target or initiator FC ports, just like the onboard FC ports.

Model	Type	Number of ports	Port speed	Supported slots
X2056-R6	FC	4 ports	8 Gb	Vertical I/O slot
X1132A-R6	FC	4 ports	8 Gb	Any regular I/O slot
X1143A-R6	CNA	2 ports	16 Gb	Any regular I/O slot

### How to prevent loss of connectivity when using the X1133A-R6 adapter

The X1133A-R6 HBA is a 4-port, 16-Gb, target-only FC adapter consisting of two, 2-port pairs. Each 2-port pair is supported by a single ASIC. If an error occurs with the ASIC supporting a pair, both ports in the pair will go offline.

To prevent loss of connectivity in the event of port failure, it is recommended that 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.

## Port configuration options for the X1143A-R6 adapter

By default the X1143A-R6 adapter is configured in FC target mode, but you can configure its ports as either 10 Gb Ethernet and FCoE (CNA) ports or as 16 Gb FC initiator or target ports.

Port pairs connected to the same ASIC must be configured in the same mode.

In FC mode, the X1143A-R6 device behaves just like any existing FC device with speeds up to 16 Gbps. In CNA mode, you can use the X1143A-R6 device for concurrent NIC and FCoE traffic sharing the same 10 GbE port. CNA mode only supports FC target mode for the FCoE function.

## FC supported speeds

FC target ports can be configured to run at different speeds. You should set the target port speed to match the speed of the device to which it connects. 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 instead of using autonegotiation. 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
<ul style="list-style-type: none"> <li>• 4 Gb</li> <li>• 2 Gb</li> <li>• 1 Gb</li> </ul>	<ul style="list-style-type: none"> <li>• 8 Gb</li> <li>• 4 Gb</li> <li>• 2 Gb</li> </ul>	<ul style="list-style-type: none"> <li>• 16 Gb</li> <li>• 8 Gb</li> <li>• 4 Gb</li> </ul>

## FC Target port configuration recommendations

For best performance and highest availability, you should use the recommended FC target port configuration.

The following table shows the preferred port usage order for onboard FC target ports. For expansion adapters, the FC ports should be spread so that they do not use the same ASIC for connectivity. The preferred slot order is listed in the *Hardware Universe* for the version of Data ONTAP software used by your controller.

**Note:** The FAS22xx and FAS2520 systems do not have onboard FC ports and do not support add-on HBAs.

Controller	Port pairs with shared ASIC	Number of target ports: Preferred ports
8080,8060 and 8040	0e+0f 0g+0h	1: 0e 2: 0e, 0g 3: 0e, 0g, 0h 4: 0e, 0g, 0f, 0h
8020	0c+0d	1: 0c 2: 0c, 0d

Controller	Port pairs with shared ASIC	Number of target ports: Preferred ports
62xx	0a+0b 0c+0d	1: 0a 2: 0a, 0c 3: 0a, 0c, 0b 4: 0a, 0c, 0b, 0d
6080 and 6040	0a+0b 0c+0d 0e+0f 0g+0h	1: 0h 2: 0h, 0d 3: 0h, 0d, 0f 4: 0h, 0d, 0f, 0b 5: 0h, 0d, 0f, 0b, 0g 6: 0h, 0d, 0f, 0b, 0g, 0c 7: 0h, 0d, 0f, 0b, 0g, 0c, 0e 8: 0h, 0d, 0f, 0b, 0g, 0c, 0e, 0a
32xx	0c+0d	1: 0c 2: 0c, 0d
31xx	0a+0b 0c+0d	1: 0d 2: 0d, 0b 3: 0d, 0b, 0c 4: 0d, 0b, 0c, 0a
FAS2554 and FAS2552	0c+0d 0e+0f	1: 0c 2: 0c, 0e 3: 0c, 0e, 0d 4: 0c, 0e, 0d, 0f

**Related information**

*[NetApp Hardware Universe](#)*

## Ways to Configure FCoE

---

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 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.

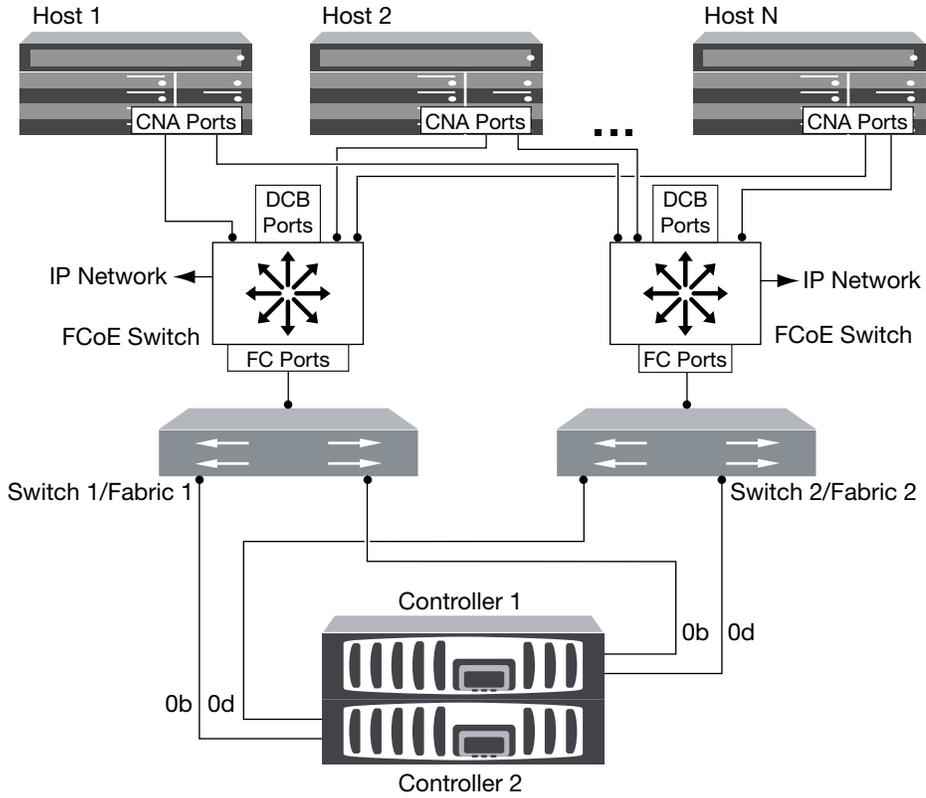
FCoE configurations require Ethernet switches that explicitly support FCoE features. FCoE configurations are validated through the same interoperability and quality assurance process as FC switches. Supported configurations are listed in the Interoperability Matrix. Some of the parameters included in 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

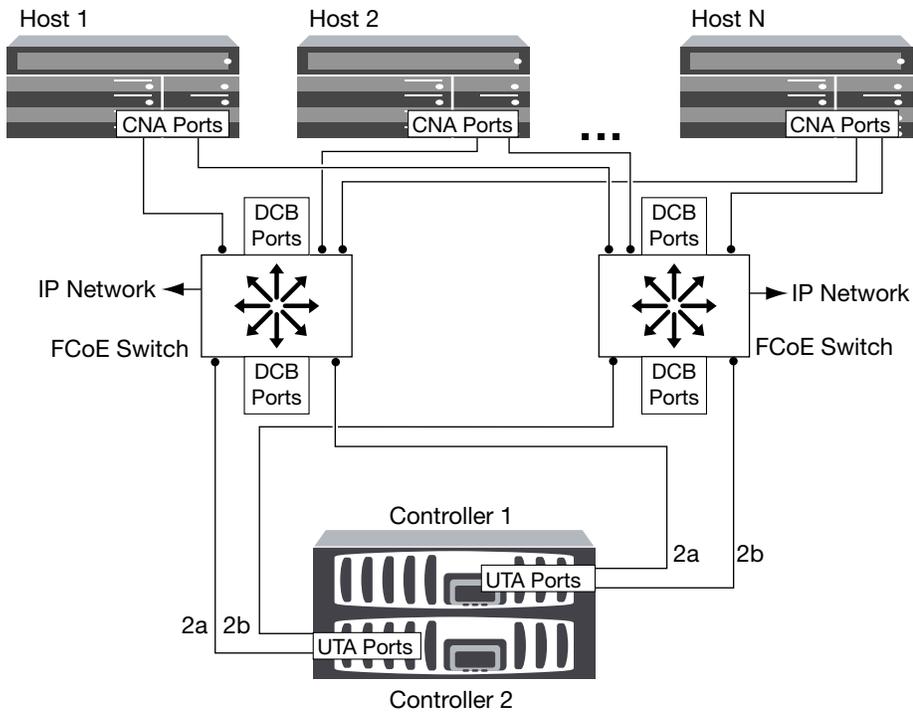
Using FCoE initiators (CNAs), you can connect hosts to both controllers in an HA pair through FCoE switches to FC target ports. The FCoE switch must also have FC ports. The host FCoE initiator always connects to the FCoE switch. The FCoE switch can connect directly to the FC target or can connect to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



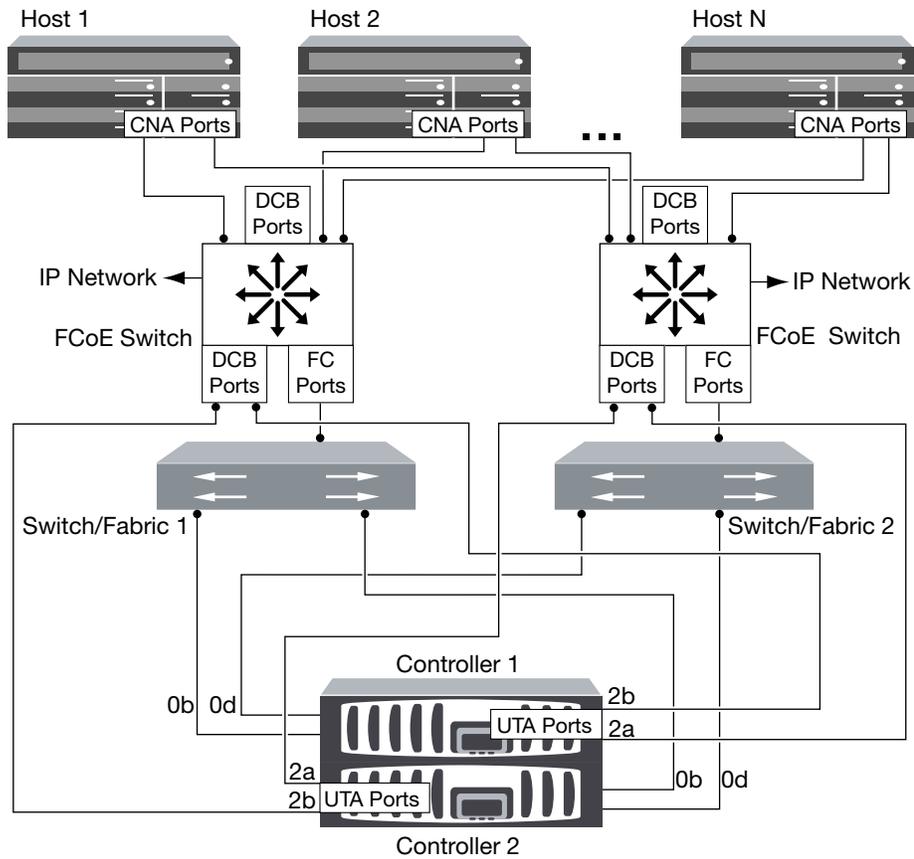
**FCoE initiator to FCoE target**

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.



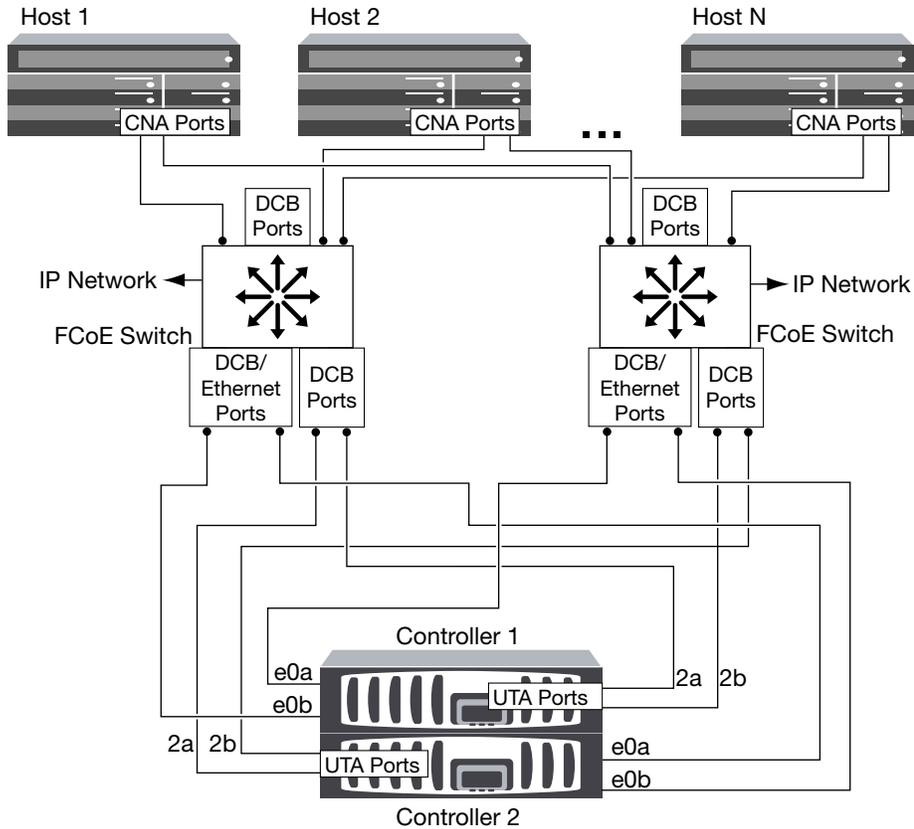
### FCoE initiator to FCoE and FC targets

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, CIFS, iSCSI, and other Ethernet traffic, can use regular Ethernet ports on the FCoE switches.



## FCoE initiator and target 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.

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

## FCoE supported hop count

The maximum supported Fibre Channel over Ethernet (FCoE) hop count between a host and storage system depends on the switch supplier and storage system support for FCoE configurations.

The hop count is defined as the number of switches in the path between the initiator (host) and target (storage system). 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).

The following table lists the maximum supported hop counts:

Switch supplier	Supported hop count
Brocade	7 for FC 5 for FCoE
Cisco	7 Up to 3 of the switches can be FCoE switches.

## Fibre Channel and FCoE zoning

---

An FC or FCoE zone is a logical grouping of one or more ports within a fabric. For devices to be able to see each other, connect, create sessions with one another, and communicate, both ports need to have a common zone membership. Single initiator zoning is recommended.

### Reasons for zoning

- Zoning reduces or eliminates *crosstalk* between initiator HBAs.  
This occurs even in small environments and is one of the best arguments for implementing zoning. The logical fabric subsets created by zoning eliminate crosstalk problems.
- Zoning reduces the number of available paths to a particular FC or FCoE port and reduces the number of paths between a host and a particular LUN that is visible.  
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 that an OS multipathing driver sees. If a host does not have a multipathing solution installed, you need to verify that only one path to a LUN is visible by using either zoning in the fabric or portsets in the SVM.
- Zoning increases security by limiting access and connectivity to end-points that share a common zone.  
Ports that have no zones in common cannot communicate with one another.
- Zoning improves SAN reliability by isolating problems that occur and helps to reduce problem resolution time by limiting the problem space.

### Recommendations for zoning

- You should implement zoning any time four or more hosts are connected to a SAN.
- Although World Wide Node Name zoning is possible with some switch vendors, World Wide Port Name zoning is required to properly define a specific port and to use NPIV effectively.
- You should limit the zone size while still maintaining manageability. Multiple zones can overlap to limit size. Ideally, a zone is defined for each host or host cluster.
- You should use single-initiator zoning to eliminate crosstalk between initiator HBAs.

## World Wide Name-based zoning

Zoning based on World Wide Name (WWN) specifies the WWN of the members to be included within the zone. When zoning in clustered Data ONTAP, you must use World Wide Port Name (WWPN) zoning.

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.

For Fibre Channel paths to storage controllers running clustered Data ONTAP, be sure the FC switches are zoned using the WWPNs of the target logical interfaces (LIFs), not the WWPNs of the physical ports on the node. For more information on LIFs, see the *Clustered Data ONTAP Network Management Guide*.

## Individual zones

In the recommended zoning configuration, there is one host initiator per zone. The zone consists of the host initiator port and one or more target LIFs on each storage node up to the desired number of paths per target. This means that hosts accessing the same nodes cannot see each other's ports, but each initiator can access any node.

For Fibre Channel paths to nodes running clustered Data ONTAP, be sure the FC switches are zoned using the WWPNs of the target 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."

## Single-fabric zoning

In a single-fabric configuration, you can still connect each host initiator to each storage node. Multipathing software is required on the host to manage multiple paths. Each host should have two initiators for multipathing to provide resiliency in the solution.

Each 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 every node 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, 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.

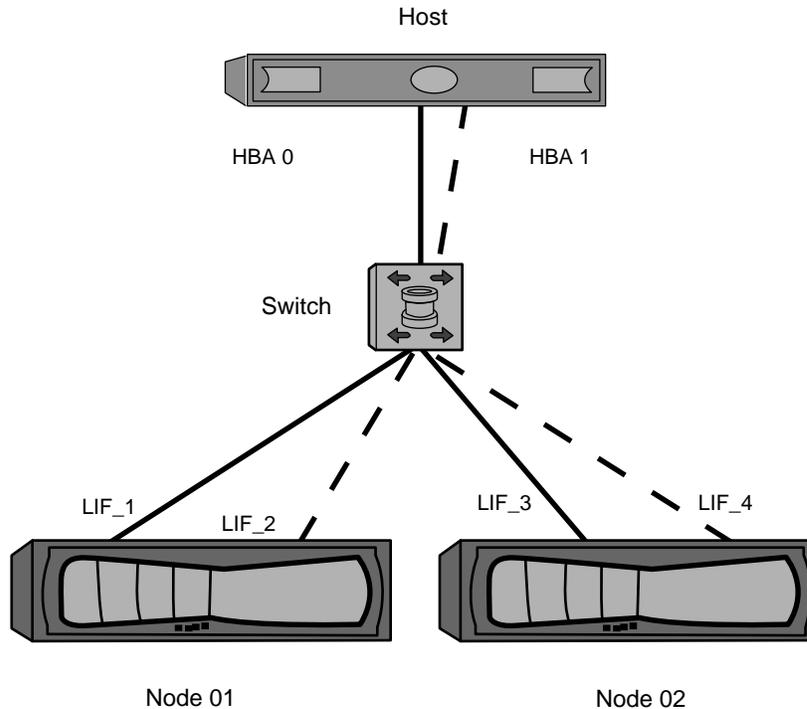
Single-fabric configurations are supported, but are not considered highly available. The failure of a single component can cause loss of access to data.

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

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

- Zone 1: HBA 0, LIF\_1, and LIF\_3
- Zone 2: HBA 1, LIF\_2, and LIF\_4

If the configuration included more nodes, the LIFs for the additional nodes would be included in these zones.



**Figure 1: Single-fabric zoning**

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

- 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

**Note:** 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 or storage controllers], see the configuration limits information elsewhere in this document.

## Dual-fabric HA pair zoning

In dual fabric configurations, you can connect each host initiator to each cluster node. Each host initiator uses a different switch to access the cluster nodes. Multipathing software is required on the host to manage multiple paths.

Dual fabric configurations are considered high availability because access to data is maintained in the event of a single component failure.

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

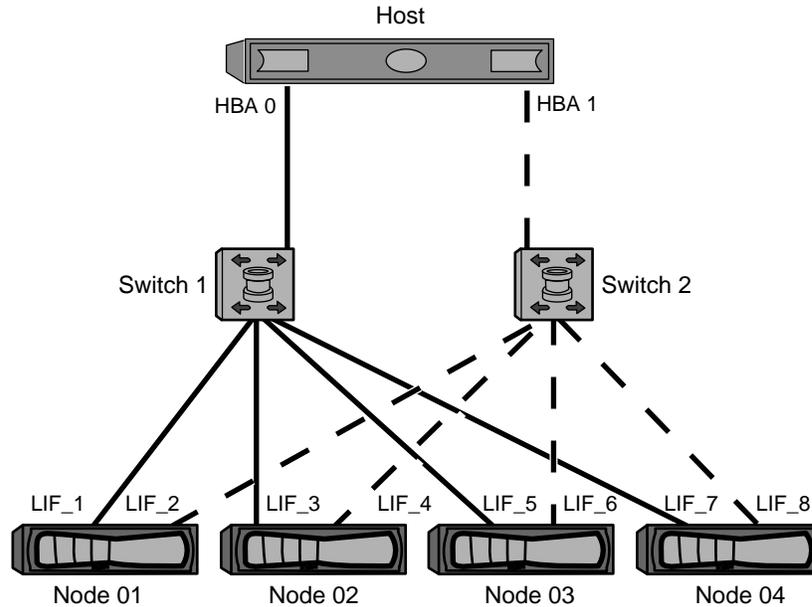
**Note:** The naming convention used in this figure is just a recommendation of one possible naming convention that you can choose to use for your Clustered Data ONTAP solution.

- 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 initiator can access a LIF on every node. This enables the host to still access its LUNs if a node fails. Storage Virtual Machines (SVMs) have access to all iSCSI and FCP LIFs on every node in a clustered solution. Portsets or FC switch zoning can be used to reduce the number of paths from a SVM to the host and the number of paths from a SVM to a LUN.

If the configuration included more nodes, the LIFs for the additional nodes would be included in these zones.



**Figure 2: Dual-fabric zoning**

**Note:** 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. Information on supported path and LUN limitations can be verified by using the configuration limits at the end of this document.

## FC and FCoE LIFs on the same port need to be in separate zones

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

Multiple LIFs for the FC and FCoE protocols can share physical ports on a node as long as they are in different zones. Cisco FC and FCoE switches require each LIF on a given port to be in a separate zone from the other LIFs on that port.

A single zone can have both FC and FCoE LIFs. A zone can contain a LIF from every target port in the cluster, but be careful to not exceed the host's path limits.

LIFs on different physical ports can be in the same zone.

While this is a requirement for Cisco switches, separating LIFs is a good idea for all switches.

## Requirements for shared SAN configurations

---

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

For all host operating systems, NetApp recommends using separate adapters to connect to each vendor's storage. Using separate adapters reduces the chances of conflicting drivers and settings. For connections to Data ONTAP storage, the adapter model, BIOS, firmware, and driver must be listed as supported in the NetApp Interoperability Matrix.

Set the required or recommended timeout values and other storage parameters for the host. Always install the NetApp software or apply the NetApp settings last.

- For AIX, apply the values from the AIX Host Utilities version listed in the Interoperability Matrix for your configuration.
- For ESX, apply host settings using Virtual Storage Console for VMware vSphere.
- For HP-UX, use the HP-UX default storage settings.
- For Linux, apply the values from the Linux Host Utilities version listed in the Interoperability Matrix for your configuration.
- For Solaris, apply the values from the AIX Host Utilities version listed in the Interoperability Matrix for your configuration.
- For Windows, install either the Data ONTAP DSM for Windows MPIO or the Windows Host Utilities version listed in the Interoperability Matrix for your configuration.

### Related information

[\*NetApp Interoperability Matrix Tool\*](#)

[\*NetApp Documentation: Host Utilities \(current releases\)\*](#)

## ALUA Configurations

---

Clustered Data ONTAP always uses asymmetric logical unit access (ALUA) for both FC and iSCSI paths. Be sure to use host configurations that support ALUA.

ALUA is an industry-standard protocol for identifying optimized paths between a storage system and a host computer. The administrator of the host computer does not need to manually select the paths to use.

You do not need to enable ALUA on storage nodes, and you cannot disable it.

For information about which specific host configurations support ALUA, see the Interoperability Matrix and the Host Utilities *Installation and Setup Guide* for your host operating system.

### Related information

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

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

## When host multipathing software is required

If there is more than one path from the Storage Virtual Machine (SVM) logical interfaces (LIFs) to the fabric, multipathing software is required. Multipathing software is required on the host any time the host can access a LUN through more than one path.

The multipathing software presents a single disk to the operating system for all paths to a LUN. Without multipathing software, the operating system could treat each path as a separate disk, which can lead 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

In single-fabric single-node configurations, multipathing software is not required if you only have a single path from the host to the node.

Multipathing 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 recommended.

Multipathing software is also known as MPIO (multipath I/O) software.

## Recommended number of paths from host to nodes in cluster

You should not exceed more than 8 paths from your host to each node in your cluster.

You should have a minimum of two paths per LUN connecting to each node 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 8 paths.

- Selective LUN Map (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 for all LUN maps created in Data ONTAP 8.3. SLM can be manually enabled on LUN maps created prior to Data ONTAP 8.3.
- Portsets for iSCSI
- FC igroup mappings from your host
- FC switch zoning

## Configuration limits for FC, FCoE, and iSCSI configurations

---

Configuration limits are available for FC, FCoE, and iSCSI configurations. In some cases, theoretical limits might be higher, but the published limits are tested and supported.

### SAN configuration requirement for FlexVol volumes

Volumes containing LUNs must be FlexVol volumes. SAN protocols can only be used with Storage Virtual Machines (SVMs) with FlexVol volumes. Infinite Volumes are not supported for SAN.

In this document, “volume” always means “FlexVol volume” and “SVM” always means “SVM” with FlexVol volumes.

### SAN configuration limit parameters and definitions

There are a number of parameters and definitions related to FC, FCoE, and iSCSI configuration limits.

Parameter	Definition
Visible target ports per host (iSCSI)	The maximum number of target iSCSI Ethernet ports that a host can see or access on iSCSI-attached controllers.
Visible target ports per host (FC)	The maximum number of FC adapters that a host can see or access on the attached FC controllers.
LUNs per host	The maximum number of LUNs that you can map from the controllers to a single host.
Maximum paths from host to LUN	The maximum number of paths from the host to a single LUN <b>Note:</b> Using the maximum number of paths is not recommended.
Maximum paths from host to storage solution	The maximum total number of paths from the host to the entire cluster.
Maximum LUN size	The maximum size of an individual LUN on the respective operating system.
Storage Virtual Machines (SVMs)	The total number of SVMs, including the default node SVM. There are limits on SVMs per node and SVMs per cluster.
Volumes per node	The total number of volumes supported on a single node.
LUNs per controller or node	The maximum number of LUNs that you can configure per controller, including cloned LUNs and LUNs contained within cloned volumes. LUNs contained in Snapshot copies do not count in this limit, and there is no limit on the number of LUNs that can be contained within Snapshot copies.
LUNs per volume	The maximum number of LUNs that you can configure within a single volume. LUNs contained in Snapshot copies do not count in this limit, and there is no limit on the number of LUNs that can be contained within Snapshot copies.

Parameter	Definition
FC port fan-in	The maximum number of initiator-target nexus (ITNs) that can connect to a single FC port on a controller. Connecting the maximum number of ITNs is generally not recommended, and you might need to tune the FC queue depths on the host to achieve this maximum value.
FC LIFs per port	The maximum number of FC logical interfaces (LIFs) that can be defined on a single physical FC port.
iSCSI sessions per controller or node	The recommended maximum number of iSCSI sessions that you can connect to a single controller. The general formula to calculate this is as follows: $\text{Maximum sessions} = 8 \times \text{System Memory} \div 512 \text{ MB}$ .
IP LIFs per port	The maximum combined number of iSCSI, NFS, CIFS, and logical interfaces (LIFs) that can be defined on a single physical Ethernet port.
LIFs per port set	The maximum number of logical interfaces that can be assigned to a single port set.
FC ITNs per controller	The maximum number of initiator-target nexus (ITNs) supported per port per controller.
igroups per controller	The maximum number of initiator groups that you can configure per controller.
Initiators per igroup	The maximum number of FC initiators (HBA WWNs) or iSCSI initiators (host iqn/eui node names) that you can include in a single igroup.
LUN mappings per controller	The maximum number of LUN mappings per controller. For example, a LUN mapped to two igroups counts as two mappings.
LUN path name length	The maximum number of characters in a full LUN name. For example, <code>/vol/abc/def</code> has 12 characters.
LUN size	The maximum capacity of an individual LUN on a controller.
FC queue depth available per port	The usable queue depth capacity of each FC target port. The number of LUNs is limited by available FC queue depth.
Ethernet ports per node	The maximum number of supported Ethernet ports per node.
FC target ports per controller or node	The maximum number of supported FC target ports per controller. FC initiator ports used for back-end disk connections, for example, connections to disk shelves, are not included in this number.
Port sets per node	The maximum number of port sets that can be created on a single node.

### Related tasks

*Calculating queue depth* on page 43

## Determining the number of supported nodes for SAN configurations

The number of nodes per cluster supported by clustered Data ONTAP varies depending on your version of Data ONTAP, the storage controller models in your cluster, and the protocol of your cluster nodes.

### About this task

If any node in the cluster is configured for FC, 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 the *Hardware Universe*.
2. Click **Controllers** in the upper left (next to the **Home** button).
3. Select the check box next to your version of Data ONTAP in the **Choose Your Data ONTAP Version** column.  
A new column is displayed for you to choose your platforms.
4. Select the check boxes next to the platforms used in your solution in the **Choose Your Platforms** column.
5. Unselect the **Select All** check box in the **Choose Your Specifications** column.
6. Select the **Max Nodes per Cluster - SAN** check box.
7. Click **Show Results**.

### Related references

[SAN configuration limits](#) on page 34

### Related information

[NetApp Hardware Universe](#)

## Determining the number of supported hosts per cluster in FC configurations

The maximum number of SAN hosts that can be connected to a cluster varies greatly based upon your specific combination of multiple cluster attributes, such as the number of hosts connected to each cluster node, initiators per host, sessions per host, and nodes in the cluster.

### About this task

For FC 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 configurations is 2,048. As long as 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:

```
fcip 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.

## Determining the supported number of hosts in iSCSI configurations

The maximum number of SAN hosts that can be connected in iSCSI configurations varies greatly based on your specific combination of multiple cluster attributes, such as the number of hosts connected to each cluster node, initiators per host, logins per host, and nodes in the cluster.

#### About this task

Single-node configurations are supported in iSCSI. The number of hosts that can be directly connected to 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 *Hardware Universe*.

For all multi-node cluster configurations, 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.

#### Related references

[SAN configuration limits](#) on page 34

## Host operating system limits for SAN configurations

Each host operating system has host-based configuration limits for FC, FCoE, and iSCSI.

The following table lists the maximum supported value for each parameter based on testing. All values are for FC, FCoE and iSCSI unless otherwise noted. The values listed are the maximum supported by NetApp. The operating system vendor might support a different value. For best performance, do not configure your system at the maximum values.

**Note:** AIX 6.0 Host Utilities do not support iSCSI.

Parameter	Windows	Linux	HP-UX	Solaris	AIX	ESX
Visible target ports per host	32	24	16	16	16	16
LUNs per host	255	2,048 devices max (where each path to a LUN is a device)	11iv2: 512 11iv3: 1,024	512	1,024	256 Local drives, CD-ROM, and so on count against this value.
Maximum paths from host to LUN	32	24 (max of 2,048 per host)	11iv2: 8 11iv3: 32	16	16	8 (max of 1,024 per host)
Maximum paths from host to storage solution	1,024	1,024	11iv2: 4,096 11iv3: 8,192	2,048	2,048	1,024 (including non-storage LUNs and devices)
Maximum LUN size	2 TB (MBR) 16 TB (GPT)	16 TB	11iv2: 2 TB 11iv3: 16 TB	16 TB	16 TB	16 TB (VMFS-5 and pass through RDM) 2 TB (VMFS-3 and non-pass-through RDM)

### Related references

[SAN configuration limit parameters and definitions](#) on page 29

## SAN configuration limits

You should use the tested and supported maximum configuration limits established for each storage controller in your SAN environment. For reliable operations, you should not exceed the tested and supported limits.

The following table lists the maximum supported value for each parameter based on per node testing, but for best performance, you should not configure your node with the maximum values. All values are for FC, FCoE, and iSCSI unless otherwise noted.

The maximum number of LUNs and the number of host HBAs that can connect to an FC port are limited by the available queue depth on the FC target ports.

Clusters with mixed controller types have cluster limits based on the node with the lowest limits. The cluster limits are the single node limits of the lowest limit node multiplied by the number of nodes in the cluster. The maximum number of nodes within a cluster is determined by the platform that supports the fewest number of nodes. If any node in the cluster is configured for FC, FCoE, or iSCSI, the cluster is limited to the SAN node limits. See the *Hardware Universe* for SAN node limits for your specific platforms.

Storage limits for Cloud ONTAP are documented in the *Cloud ONTAP Release Notes*.

Parameters	Maximum per node
Storage Virtual Machines (SVMs)	32 for the FAS2220 125 for all other models
Volumes	500 for the following models: <ul style="list-style-type: none"> <li>• FAS2220</li> <li>• FAS2240</li> <li>• FAS3220</li> <li>• FAS3240</li> </ul> 1,000 for all other models
Volume per node per SVM	500 for the following models: <ul style="list-style-type: none"> <li>• FAS2220</li> <li>• FAS2240</li> <li>• FAS3220</li> <li>• FAS3240</li> </ul> 1,000 for all other models

Parameters	Maximum per node
LUNs	<p data-bbox="678 254 922 279">1,024 for the FAS2220</p> <p data-bbox="678 296 1016 321">2,048 for the following models:</p> <ul data-bbox="678 338 821 548" style="list-style-type: none"><li data-bbox="678 338 821 363">• FAS2240</li><li data-bbox="678 373 821 399">• FAS2520</li><li data-bbox="678 409 821 434">• FAS2552</li><li data-bbox="678 445 821 470">• FAS2554</li><li data-bbox="678 480 821 506">• FAS3220</li><li data-bbox="678 516 821 541">• FAS3240</li></ul> <p data-bbox="678 569 1016 594">8,192 for the following models:</p> <ul data-bbox="678 611 821 926" style="list-style-type: none"><li data-bbox="678 611 821 636">• FAS3250</li><li data-bbox="678 646 773 672">• 3270</li><li data-bbox="678 682 773 707">• 6210</li><li data-bbox="678 718 773 743">• 6220</li><li data-bbox="678 753 773 779">• 6240</li><li data-bbox="678 789 773 814">• 8020</li><li data-bbox="678 825 773 850">• 8040</li><li data-bbox="678 861 821 886">• AFF8020</li><li data-bbox="678 896 821 921">• AFF8040</li></ul> <p data-bbox="678 949 1029 974">12,288 for the following models:</p> <ul data-bbox="678 991 821 1241" style="list-style-type: none"><li data-bbox="678 991 773 1016">• 6250</li><li data-bbox="678 1026 773 1052">• 6280</li><li data-bbox="678 1062 773 1087">• 6290</li><li data-bbox="678 1098 773 1123">• 8060</li><li data-bbox="678 1134 773 1159">• 8080</li><li data-bbox="678 1169 821 1194">• AFF8060</li><li data-bbox="678 1205 821 1230">• AFF8080</li></ul>

Parameters	Maximum per node
LUN mappings	1,024 for the FAS2220 2,048 for the following models: <ul style="list-style-type: none"> <li>• FAS2240</li> <li>• FAS2520</li> <li>• FAS2552</li> <li>• FAS2554</li> <li>• FAS3220</li> <li>• FAS3240</li> </ul> 8,192 for the following models: <ul style="list-style-type: none"> <li>• FAS3250</li> <li>• 3270</li> <li>• 6210</li> <li>• 6220</li> <li>• 6240</li> <li>• 8020</li> <li>• 8040</li> <li>• AFF8020</li> <li>• AFF8040</li> </ul> 12,288 for the following models: <ul style="list-style-type: none"> <li>• 6250</li> <li>• 6280</li> <li>• 6290</li> <li>• 8060</li> <li>• 8080</li> <li>• AFF8060</li> <li>• AFF8080</li> </ul>
LUNs per volume	512 for all models
LUN size	16 TB for all models

Parameters	Maximum per node
iSCSI sessions per node	1,024 for the FAS2220 2,048 for the following models: <ul style="list-style-type: none"> <li>• FAS2240</li> <li>• FAS2520</li> <li>• FAS2552</li> <li>• FAS2554</li> <li>• FAS3220</li> <li>• FAS3240</li> </ul> 4,096 for the following models: <ul style="list-style-type: none"> <li>• FAS3250</li> <li>• 3270</li> <li>• 6210</li> <li>• 6220</li> <li>• 6240</li> <li>• 8020</li> <li>• 8040</li> <li>• AFF8020</li> <li>• AFF8040</li> </ul> 8,192 for the following models: <ul style="list-style-type: none"> <li>• 6250</li> <li>• 6280</li> <li>• 6290</li> <li>• 8060</li> <li>• 8080</li> <li>• AFF8060</li> <li>• AFF8080</li> </ul>
iSCSI LIFs per node	256 for all models
iSCSI LIFs per port	32 for all models
iSCSI LIFs per port set	32 for all models

Parameters	Maximum per node
igroups	256 for the FAS2220 512 for the following models: <ul style="list-style-type: none"> <li>• FAS2240</li> <li>• FAS2520</li> <li>• FAS2552</li> <li>• FAS2554</li> <li>• FAS3220</li> <li>• FAS3240</li> </ul> 2,048 for the following models: <ul style="list-style-type: none"> <li>• FAS3250</li> <li>• 3270</li> <li>• 6210</li> <li>• 6220</li> <li>• 6240</li> <li>• 8020</li> <li>• 8040</li> <li>• AFF8020</li> <li>• AFF8040</li> </ul> 4,096 for the following models: <ul style="list-style-type: none"> <li>• 6250</li> <li>• 6280</li> <li>• 6290</li> <li>• 8060</li> <li>• 8080</li> <li>• AFF8060</li> <li>• AFF8080</li> </ul>
Initiators per igroup	128 for the FAS2220 256 for all other models

Parameters	Maximum per node
Initiators per node	<p>512 for the FAS2220</p> <p>1,024 for the following models:</p> <ul style="list-style-type: none"> <li>• FAS2240</li> <li>• FAS2520</li> <li>• FAS2552</li> <li>• FAS2554</li> <li>• FAS3220</li> <li>• FAS3240</li> </ul> <p>2,048 for the following models:</p> <ul style="list-style-type: none"> <li>• FAS3250</li> <li>• 3270</li> <li>• 6210</li> <li>• 6220</li> <li>• 6240</li> <li>• 8020</li> <li>• 8040</li> <li>• AFF8020</li> <li>• AFF8040</li> </ul> <p>4,096 for the following models:</p> <ul style="list-style-type: none"> <li>• 6250</li> <li>• 6280</li> <li>• 6290</li> <li>• 8060</li> <li>• 8080</li> <li>• AFF8060</li> <li>• AFF8080</li> </ul>

Parameters	Maximum per node
Portsets	256 for the FAS2220 512 for the following models: <ul style="list-style-type: none"> <li>• FAS2240</li> <li>• FAS2520</li> <li>• FAS2552</li> <li>• FAS2554</li> <li>• FAS3220</li> <li>• FAS3240</li> </ul> 2,048 for the following models: <ul style="list-style-type: none"> <li>• FAS3250</li> <li>• 3270</li> <li>• 6210</li> <li>• 6220</li> <li>• 6240</li> <li>• 8020</li> <li>• 8040</li> <li>• AFF8020</li> <li>• AFF8040</li> </ul> 4,096 for the following models: <ul style="list-style-type: none"> <li>• 6250</li> <li>• 6280</li> <li>• 6290</li> <li>• 8060</li> <li>• 8080</li> <li>• AFF8060</li> <li>• AFF8080</li> </ul>
Ethernet ports	See <i>Hardware Universe</i> for platform supported limits.
<p><b>Note:</b> The following FC limits do not apply to the FAS2220 or the FAS2520. The FAS2220 and FAS2520 do not support FC. These FC limits are the same for all other platforms.</p>	
FC ports	See <i>Hardware Universe</i> for platform supported limits.
FC queue depth available per port	2,048 for all models
FC ITNs per port	512 for all models

Parameters	Maximum per node
FC ITNs per node	8,192 for the following models: <ul style="list-style-type: none"> <li>• 6250</li> <li>• 6280</li> <li>• 6290</li> <li>• 8060</li> <li>• 8080</li> <li>• AFF8060</li> <li>• AFF8080</li> </ul> 4,096 for the following models: <ul style="list-style-type: none"> <li>• FAS3250</li> <li>• 3270</li> <li>• 6210</li> <li>• 6220</li> <li>• 6240</li> <li>• 8020</li> <li>• 8040</li> <li>• AFF8020</li> <li>• AFF8040</li> </ul> 2,048 for the following models: <ul style="list-style-type: none"> <li>• 6250</li> <li>• 6280</li> <li>• 6290</li> <li>• 8060</li> <li>• 8080</li> <li>• AFF8060</li> <li>• AFF8080</li> </ul>
FC (FC/FCoE) LIFs per port	32 for all models
FC LIFs portset	32 for all models
FC LIFs per node	512 for all models

**Related tasks**

[Determining the number of supported nodes for SAN configurations](#) on page 31

[Determining the number of supported hosts per cluster in FC configurations](#) on page 31

[Determining the supported number of hosts in iSCSI configurations](#) on page 32

**Related information**

[NetApp Hardware Universe](#)

## SAN configuration limits for Data ONTAP-v platforms

Data ONTAP-v platforms, such as the Data ONTAP Edge storage system, have configuration limits for reliable operation. For best performance, do not configure your system at the maximum values.

The following table lists the maximum supported value for each parameter based on testing for Data ONTAP Edge systems. Do not exceed the tested limits. All values are for iSCSI. Data ONTAP-v platforms do not support FC.

**Note:** Data ONTAP Edge systems are single-node clusters only.

Parameter	Data ONTAP Edge (FDvM200)
Nodes per cluster	1
Storage Virtual Machines (SVMs)	8
Volumes	200
Volumes per node per SVM	200
LUNs per node	512
LUN maps	1,024
LUN size	16 TB
LUNs per volume	512
LUN path name length	255
igroups	256
Initiators per igroup	128
iSCSI sessions per node	128
iSCSI LIFs per port	8
iSCSI LIFs per portset	8
Portsets	32
Ethernet ports	6

## FC switch configuration limits

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.

### Related information

*Cisco Configuration Limits* - [www.cisco.com/en/US/products/ps5989/products\\_installation\\_and\\_configuration\\_guides\\_list.html](http://www.cisco.com/en/US/products/ps5989/products_installation_and_configuration_guides_list.html)

## Calculating queue depth

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, 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. This prevents hosts with small queue depths from being starved by hosts with large queue depths.

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

- For small to mid-size systems, use a HBA queue depth of 32.
- For large systems, use a 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.
- Ensure that the storage controller target FC port queue depth is not exceeded to avoid performance penalties or errors.

### Steps

1. Count the total number of FC initiators in all 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.

### Example

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 your initiators to 128.

- If the result is greater than 2,048, go to step 3.

**Example**

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.

- Option 1:
  - a. Add more FC target ports.
  - b. Redistribute your FC initiators.
  - c. Repeat steps 1 and 2.

**Example**

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:
  - a. Designate each host as "large" or "small" based on its expected I/O need.
  - b. Multiply the number of large initiators by 128.
  - c. Multiply the number of small initiators by 32.
  - d. Add the two results together.
  - e. 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.
  - f. 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.

**Note:** 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) x (Response time)

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

**Example**

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  $(10 \times 128) + (20 \times 32) = 1,920$ .

**Example**

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

Your resulting queue depth per initiator is  $2,048/30 = 68$

## Setting queue depths on 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.

**Note:** 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.

## Setting queue depths on 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/rdisk/c2t2d0
scsictl -m queue_depth=16 /dev/rdisk/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`

## Setting queue depths on Solaris hosts

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

### About this task

- 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`  
**Note:** 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.

## Setting queue depths on VMware hosts

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

### To set maximum queue depth for a QLogic HBA

#### 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`  
**Note:** 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'`

### To change the queue depth of an Emulex HBA

#### Steps

1. Log on to the service console as the root user.
2. Use the `#vmkload_mod -l grep lpfcdd` 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`  
**Note:** 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`.

## Setting queue depths on Windows hosts

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

### To update Emulex HBA queue depths

#### 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**.

**Note:** 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\lpxnids\Parameters\Device\NumberOfRequests`

### To update Qlogic HBA queue depths

#### 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.

## Copyright information

---

Copyright © 1994–2015 NetApp, Inc. All rights reserved. Printed in the U.S.

No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

**RESTRICTED RIGHTS LEGEND:** Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

## Trademark information

---

NetApp, the NetApp logo, Go Further, Faster, AltaVault, ASUP, AutoSupport, Campaign Express, Cloud ONTAP, Clustered Data ONTAP, Customer Fitness, Data ONTAP, DataMotion, Fitness, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexArray, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexVol, FPolicy, GetSuccessful, LockVault, Manage ONTAP, Mars, MetroCluster, MultiStore, NetApp Insight, OnCommand, ONTAP, ONTAPI, RAID DP, RAID-TEC, SANtricity, SecureShare, Simplicity, Simulate ONTAP, Snap Creator, SnapCenter, SnapCopy, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, StorageGRID, Tech OnTap, Unbound Cloud, and WAFL and other names are trademarks or registered trademarks of NetApp, Inc., in the United States, and/or other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. A current list of NetApp trademarks is available on the web at <http://www.netapp.com/us/legal/netapptmlist.aspx>.

## How to send comments about documentation and receive update notifications

---

You can help us to improve the quality of our documentation by sending us your feedback. You can receive automatic notification when production-level (GA/FCS) documentation is initially released or important changes are made to existing production-level documents.

If you have suggestions for improving this document, send us your comments by email to [docomments@netapp.com](mailto:docomments@netapp.com). To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

If you want to be notified automatically when production-level documentation is released or important changes are made to existing production-level documents, follow Twitter account @NetAppDoc.

You can also contact us in the following ways:

- NetApp, Inc., 495 East Java Drive, Sunnyvale, CA 94089 U.S.
- Telephone: +1 (408) 822-6000
- Fax: +1 (408) 822-4501
- Support telephone: +1 (888) 463-8277

# Index

- 16-Gb FC port
  - supported speed [15](#)
- 4-Gb FC port
  - supported speed [15](#)
- 8-Gb FC port
  - supported speed [15](#)

## A

- adapters
  - how to prevent loss of connectivity when using X1133A-R6 [14](#)
  - port configuration options for X1143A-R6 [15](#)
- AIX hosts
  - setting queue depth for [45](#)
- ALUA configurations [27](#)
- asymmetric logical unit access (ALUA) configurations [27](#)

## B

- best practices
  - FC switch configuration [13](#)
- Brocade switches
  - FC configuration limits [42](#)

## C

- Cisco switches
  - FC and FCoE zoning requirement [25](#)
  - FC configuration limits [42](#)
- comments
  - how to send feedback about documentation [50](#)
- configuration limits
  - by SAN host operating system [33](#)
  - determining supported number of hosts in iSCSI configurations [32](#)
  - FC switch [42](#)
  - FDvM200 storage systems [42](#)
  - hosts per node [31](#)
  - SAN [34](#)
  - SAN parameter, defined [29](#)
- configuration options
  - X1143A-R6 adapter port [15](#)
- configurations
  - considerations for FC [10](#)
- configurations, SAN
  - requirements for shared [26](#)
- connectivity loss
  - how to prevent when using X1133A-R6 adapters [14](#)

## D

- Data ONTAP-v
  - configuration limits [42](#)
- documentation
  - how to receive automatic notification of changes to [50](#)

- how to send feedback about [50](#)
- dynamic VLANs [9](#)

## E

- ESX
  - SAN host configuration limits [33](#)
- expansion FC adapter
  - supported port speed [15](#)

## F

- FC
  - multifabric switch zoning [24](#)
  - number of supported hop counts [13](#)
  - port speed [15](#)
  - single-fabric switch zoning [23](#)
  - supported port speed [15](#)
  - switch configuration best practices [13](#)
  - switch WWN zoning [22](#)
  - switch zoning with individual zones [23](#)
  - target port configuration recommendations [15](#)
  - ways to configure with HA pairs [12](#)
  - ways to configure with single nodes [10](#)
  - zoning for [22](#)
- FC configurations
  - considerations for [10](#)
- FC LIF
  - zoning restrictions for Cisco switches [25](#)
- FC ports
  - support for onboard and expansion adapter [14](#)
- FC switches
  - configuration limits [42](#)
- FCoE
  - initiator and target combinations [20](#)
  - maximum supported switch hop count [21](#)
  - ways to configure [17](#)
  - zoning for [22](#)
- FCoE LIF
  - zoning restrictions for Cisco switches [25](#)
- FDvM200
  - configuration limits [42](#)
- feedback
  - how to send comments about documentation [50](#)
- FlexVol volumes
  - required for SAN [29](#)

## H

- HA pairs
  - ways to configure with FC [12](#)
  - ways to configure with iSCSI SANs [7](#)
- hop counts
  - maximum supported for FCoE switches [21](#)
  - number supported for FC [13](#)
- host multipathing software
  - when required [27](#)
- host utilities

- AIX queue depths [45](#)
- HP-UX queue depths [45](#)
- Solaris queue depths [45](#)
- VMware queue depths [46](#)
- Windows queue depths [47](#)

## hosts

- determining supported number in iSCSI configurations [32](#)

## hosts per node [31](#)

## hosts, iSCSI SAN

- ways to configure with HA pairs [7](#)
- ways to configure with single nodes [5](#)

## HP-UX

- SAN host configuration limits [33](#)

## HP-UX hosts

- setting queue depth for [45](#)

## I

### Infinite Volumes

- not supported for SAN [29](#)

### information

- how to send feedback about improving documentation [50](#)

### initiators

- FCoE and FC combinations [20](#)

### iSCSI

- dynamic VLANs [9](#)
- static VLANs [8](#)

### iSCSI configurations

- benefits of using VLANs [8](#)
- considerations for [5](#)
- determining supported number of hosts [32](#)

### iSCSI SAN hosts

- ways to configure with HA pairs [7](#)
- ways to configure with single nodes [5](#)

### ITNs

- per node limits [31](#)

## L

### LIFs

- FC and FCoE zoning restrictions for Cisco switches [25](#)

### Linux

- SAN host configuration limits [33](#)

### LUNs

- recommended number of paths to avoid single points of failure [27](#)

## M

### MPIO software

- when required [27](#)

### multipathing software

- when required [27](#)

## N

### nodes

- determining number of supported, for SAN configurations [31](#)

## O

### onboard FC port

- supported port speed [15](#)

## P

### parameters

- SAN configuration limit, defined [29](#)

### paths

- recommended number to avoid single points of failure [27](#)

### port configuration options

- X1143A-R6 adapter [15](#)

### port speed

- supported for FC [15](#)

### ports

- support for onboard and expansion adapter FC [14](#)

## Q

### queue depths

- calculating [43](#)
- setting for AIX hosts [45](#)
- setting for HP-UX hosts [45](#)
- setting for Solaris hosts [45](#)
- setting for VMware hosts [46](#)
- setting for Windows hosts [47](#)
- tuning [43](#)

## R

### recommended number of paths

- for avoiding single points of failure [27](#)

## S

### SAN

- configuration limits [34](#)

### SAN configuration limits

- by host operating system [33](#)
- FC switch [42](#)
- FDvM200 storage systems [42](#)
- parameters defined [29](#)

### SAN configurations

- determining supported number of nodes [31](#)
- requirements for shared [26](#)

### SAN hosts, iSCSI

- ways to configure with HA pairs [7](#)
- ways to configure with single nodes [5](#)

### shared SAN configurations

- requirements for [26](#)

### single points of failure

- recommended number of paths to avoid [27](#)

### single-node configurations

- ways to configure in FC [10](#)
- ways to configure iSCSI SAN hosts for [5](#)

### soft zoning

- FC switch [22](#)

### Solaris

- SAN host configuration limits [33](#)

### Solaris hosts

- setting queue depth for [45](#)
- static VLANs [8](#)
- suggestions
  - how to send feedback about documentation [50](#)
- supported nodes
  - determining number of, for SAN configurations [31](#)
- SVMs
  - with FlexVol volumes, required for SAN [29](#)
- switch
  - FC and FCoE zoning requirement [25](#)
  - FC multifabric zoning [24](#)
  - FC single-fabric zoning [23](#)
  - FC WWN zoning [22](#)
  - FC zoning with individual zones [23](#)
- switches
  - FC configuration best practices [13](#)
  - FC configuration limits [42](#)
  - maximum supported hop count for FCoE [21](#)

## T

- target port configurations
  - FC recommendations for [15](#)
- targets
  - FCoE and FC combinations [20](#)
- twitter
  - how to receive automatic notification of documentation changes [50](#)

## V

- virtual LANs
  - benefits of using in iSCSI configurations [8](#)
- VLANs

- benefits of using in iSCSI configurations [8](#)
  - dynamic [9](#)
  - static [8](#)
- VMware hosts
  - setting queue depth for [46](#)
- volumes
  - FlexVol, required for SAN [29](#)

## W

- Windows
  - SAN host configuration limits [33](#)
- Windows hosts
  - setting queue depth for [47](#)
- WWN zoning
  - FC switch [22](#)

## X

- X1133A-R6 adapters
  - how to prevent loss of connectivity when using [14](#)
- X1143A-R6 adapters
  - port configuration options [15](#)

## Z

- zoning
  - FC switch by WWN [22](#)
  - FC switch multifabric [24](#)
  - FC switch single-fabric [23](#)
  - FC switch with individual zones [23](#)
  - for FC and FCoE [22](#)
  - requirements for Cisco switches [25](#)