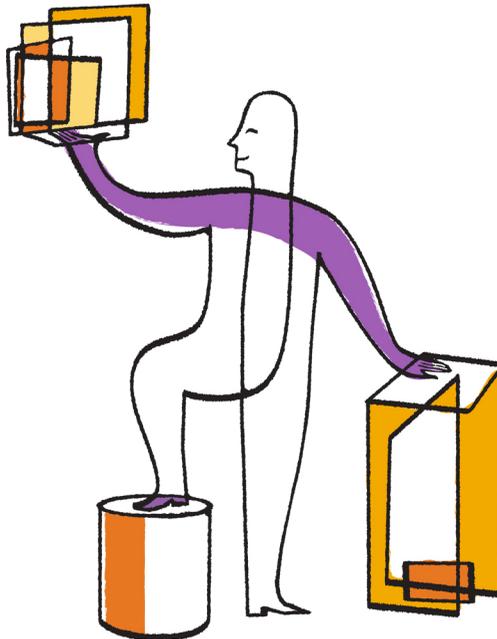




**NetApp®**

## Clustered Data ONTAP® 8.2

### 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: 215-07954\_A0  
May 2013



# Contents

<b>iSCSI configurations .....</b>	<b>6</b>
Recommended number of paths to avoid single points of failure .....	6
Single-network and multi-network single-node configuration in an iSCSI SAN .....	6
Single-network HA pair in an iSCSI SAN .....	9
Multi-network HA pair in an iSCSI SAN .....	10
Direct-attached single-node configurations in an iSCSI SAN .....	11
VLANs for iSCSI configurations .....	12
Static VLANs .....	12
Dynamic VLANs .....	12
<b>Fibre Channel configurations .....</b>	<b>13</b>
Recommended number of paths to avoid single points of failure .....	13
FC onboard and expansion port combinations .....	14
Fibre Channel supported hop counts .....	15
Fibre Channel supported speeds .....	15
Fibre Channel switch configuration best practices .....	16
Host multipathing software requirements .....	16
62xx supported fibre channel configurations .....	16
62xx target port configuration recommendations .....	17
62xx Single-node configurations .....	17
62xx: Single-fabric HA configurations .....	19
62xx: Multifabric HA configurations .....	20
6040 and 6080 supported fibre channel configurations .....	22
6040 and 6080 target port configuration recommendations .....	22
60xx Single-node configurations .....	22
6040 and 6080: Single-fabric HA configurations .....	24
6040 and 6080: Multifabric HA pair .....	26
32xx supported fibre channel configurations .....	27
32xx target port configuration recommendations .....	27
32xx Single-node configurations .....	28
32xx: Single-fabric HA configurations .....	29
32xx: Multifabric HA configurations .....	30
31xx supported fibre channel configurations .....	32

31xx target port configuration recommendations .....	32
31xx Single-node configurations .....	32
31xx: Single-fabric HA pair .....	34
31xx: Multifabric HA pair .....	35
FAS2240 Supported configurations .....	36
FAS2240 Single-fabric Single-node configurations .....	37
FAS2240 Multifabric single-node configurations .....	38
<b>Fibre Channel over Ethernet overview .....</b>	<b>39</b>
Recommended number of paths to avoid single points of failure .....	39
FCoE initiator and target combinations .....	40
Fibre Channel over Ethernet supported hop count .....	40
Fibre Channel over Ethernet supported configurations .....	41
FCoE initiator to FC target configuration .....	41
FCoE end-to-end configuration .....	43
FCoE mixed with FC .....	44
FCoE mixed with IP storage protocols .....	46
<b>Fibre Channel and FCoE zoning .....</b>	<b>48</b>
World Wide Name-based zoning .....	48
Individual zones .....	49
Single-fabric zoning .....	49
Dual-fabric HA pair zoning .....	50
FC and FCoE LIFs on the same port need to be in separate zones .....	52
<b>Shared SAN configurations .....</b>	<b>53</b>
<b>ALUA Configurations .....</b>	<b>54</b>
<b>Configuration limits for FC, FCoE, and iSCSI configurations .....</b>	<b>55</b>
SAN configurations always use FlexVol volumes .....	55
Supported number of nodes for SAN configurations .....	55
SAN Configuration limit parameters and definitions .....	55
Host operating system limits for SAN configurations .....	57
SAN configuration limits .....	59
SAN configuration limits specifically for the FAS2220 .....	61
FC switch configuration limits .....	62
Calculating queue depth .....	63
Setting queue depths on AIX hosts .....	65
Setting queue depths on HP-UX hosts .....	65
Setting queue depths on Solaris hosts .....	66

Setting queue depths on VMware hosts .....	67
Setting queue depths on Windows hosts .....	68
<b>Copyright information .....</b>	<b>69</b>
<b>Trademark information .....</b>	<b>70</b>
<b>How to send your comments .....</b>	<b>71</b>
<b>Index .....</b>	<b>72</b>

## iSCSI configurations

---

Network-attached iSCSI configurations that use HA pairs are supported for clustered Data ONTAP. Starting with clustered Data ONTAP 8.2, single-node configurations are also supported.

For clustered Data ONTAP, you need to create one or more iSCSI paths to each storage controller that can access a given LUN. This path setup differs from previous versions of Data ONTAP operating in 7-Mode. For clustered Data ONTAP, ports on a partner node do not assume the IP addresses of the failed partner. Instead, the MPIO software on the host is responsible for selecting the new paths. This behavior is very similar to Fibre Channel path failover.

In an iSCSI environment, you can connect Ethernet switches in any vendor-supported configuration. Ethernet switch only (No FCoE switches or FCoE traffic in the data path) counts are not a limitation in Ethernet iSCSI configurations. For specific recommendations and best practices, see the Ethernet switch vendor's documentation.

For Windows iSCSI multipathing options, see TR-4080 *Best Practices for Scalable SAN in Data ONTAP 8.2 Cluster-Mode*.

### Related information

[NetApp Interoperability Matrix - support.netapp.com/NOW/products/interoperability/TR-4080: Best Practices for Scalable SAN in Data ONTAP 8.2 Cluster-Mode.pdf](https://support.netapp.com/NOW/products/interoperability/TR-4080:BestPracticesforScalableSANinDataONTAP8.2Cluster-Mode.pdf)

## Recommended number of paths to avoid single points of failure

You should have a minimum of 2 paths per LUN connecting to each node in your cluster. This eliminates single points of failure and enables the system to survive component failures.

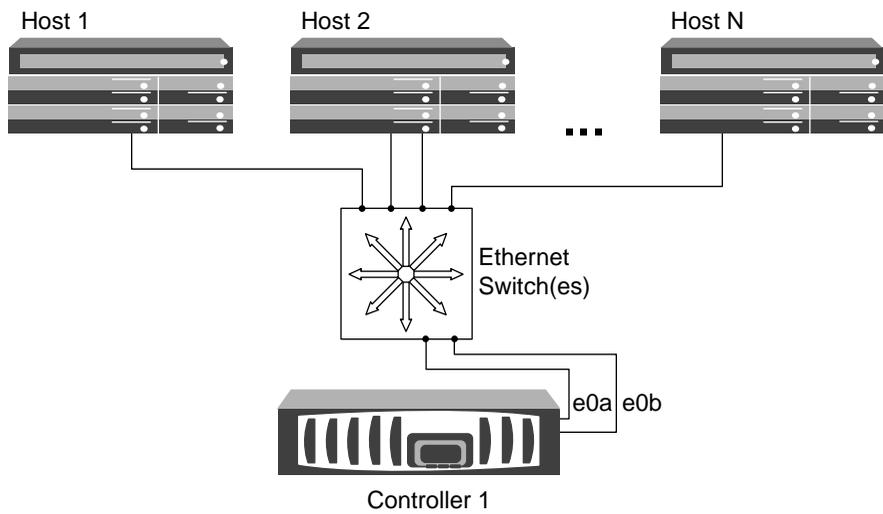
If you have more than four nodes in your cluster or more than four target ports in any of your nodes, portsets can be used to limit the number of paths that can be used to access LUNs on your controllers.

## Single-network and multi-network single-node configuration in an iSCSI SAN

You can connect hosts to a single-node that uses the iSCSI protocol over a single IP network or over a multiple IP network. The network can consist of one or more switches. Each node can have

multiple iSCSI connections to the network. The number of ports available depends on the storage controller model.

The following figure shows a single-network configuration. Multi-network configurations are also supported.

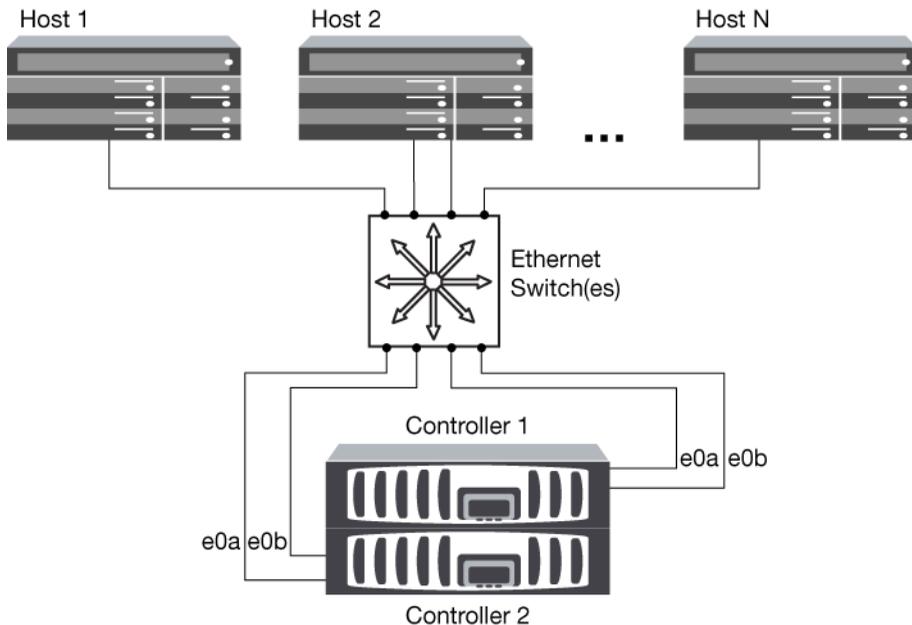


Attribute	Value
Fully redundant	No, due to the single network
Type of network	Single network
Different host operating systems	Yes, with multiple-host configurations
Multipathing required	Yes, if multiple connections per host are configured
Type of configuration	Single node

## Single-network HA pair in an iSCSI SAN

You can connect hosts to HA pair controllers that use the iSCSI protocol over a single IP network. The network can consist of one or more switches. Each controller can have multiple iSCSI connections to the network. The number of ports available depends on the storage controller model.

The following figure shows two Ethernet connections to the network per storage controller. Depending on the controller model, more connections are possible.



**Figure 2: iSCSI single-network HA pair**

Attribute	Value
Fully redundant	No, due to the single network

Attribute	Value
Type of network	Single network
Different host operating systems	Yes, with multiple-host configurations
Multipathing required	Yes, if multiple connections per host are configured
Type of configuration	HA pair

## Multi-network HA pair in an iSCSI SAN

You can connect hosts to HA pair controllers that use the iSCSI protocol over multiple IP networks. To be fully redundant, a minimum of two connections to separate networks per controller is necessary to protect against NIC, network, and cabling failure. The host will require multipathing software to be installed and configured.

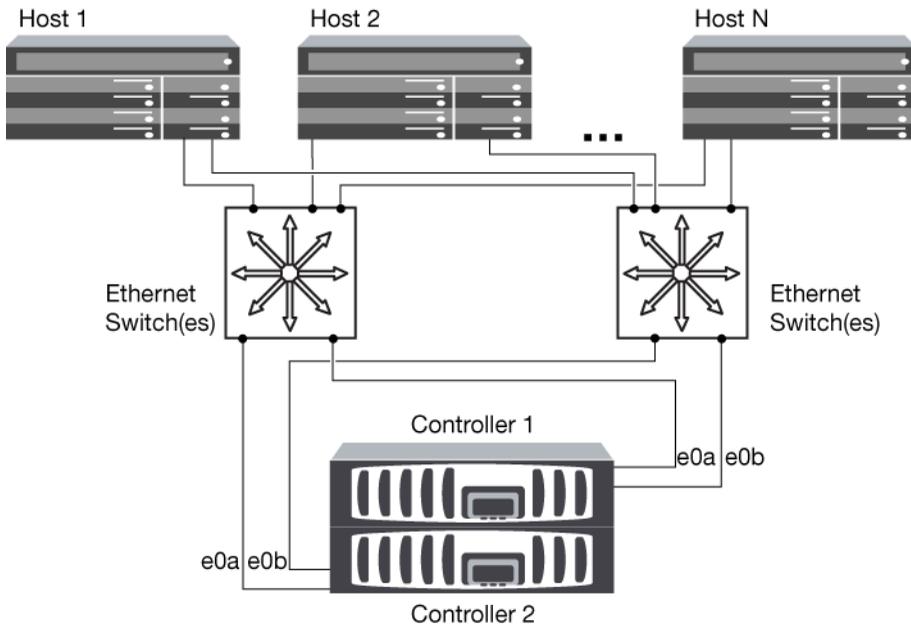


Figure 3: iSCSI multi-network HA pair

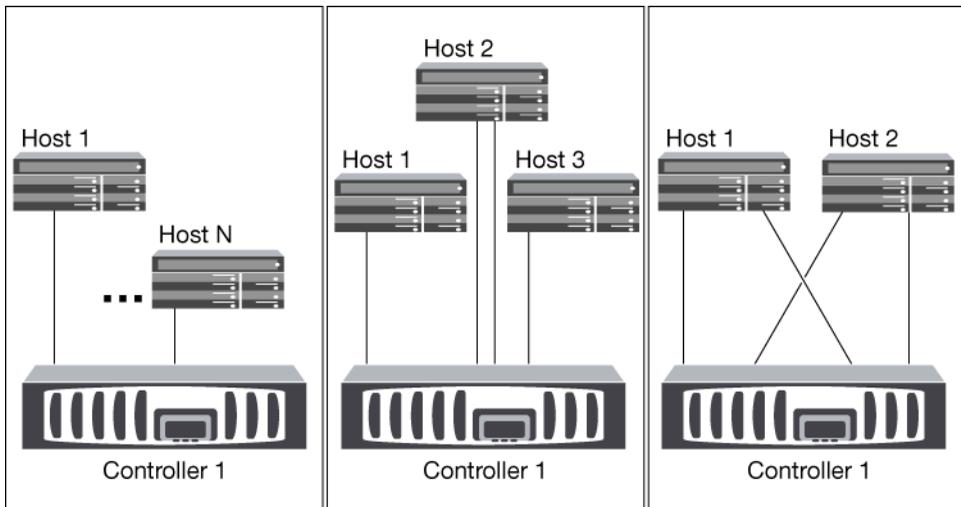
Attribute	Value
Fully redundant	Yes
Type of network	Multi-network
Different host operating systems	Yes, with multiple-host configurations

Attribute	Value
Multipathing required	Yes
Type of configuration	HA pair

## Direct-attached single-node configurations in an iSCSI SAN

You can connect hosts using iSCSI directly to controllers. The number of hosts that can be directly connected to a controller or pair of controllers depends on the number of available Ethernet ports.

**Note:** Direct-attached configurations are not supported in HA pairs.



**Figure 4: iSCSI direct-attached single-node configurations**

Attribute	Value
Fully redundant	No, due to the single node
Type of network	None, direct-attached
Different host operating systems	Yes, with multiple-host configurations
Multipathing required	Yes, if multiple connections per host are configured
Type of configuration	Single node

## VLANs for 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.

### Reasons for implementing VLANs

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.

## Fibre Channel configurations

---

Fabric-attached FC configurations that use HA pairs are supported for clustered Data ONTAP. Starting with clustered Data ONTAP 8.2, single-node fabric configurations are also supported.

For multiple-host configurations, hosts can use different operating systems, such as Windows or UNIX. Hosts require that a supported multipathing solution be installed and configured.

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. For specific exceptions, see the Interoperability Matrix on the NetApp Support Site.

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.

**Note:** The following sections show detailed SAN configuration diagrams for each type of storage system. For simplicity, the diagrams show only a single fabric or, in the case of the dual-fabric configurations, two fabrics. However, it is possible to have multiple fabrics connected to a single storage system. In the case of dual-fabric configurations, even multiples of fabrics are supported.

### Related information

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

## Recommended number of paths to avoid single points of failure

You should have a minimum of 2 paths per LUN connecting to each node in your cluster. This eliminates single points of failure and enables the system to survive component failures.

If you have more than four nodes in your cluster or more than four target ports in any of your nodes, portsets can be used to limit the number of paths that can be used to access LUNs on your controllers.

## FC onboard and expansion port combinations

You can use storage controller onboard FC ports as both initiators and targets. You can also add storage controller FC ports on expansion adapters and use them as initiators or targets.

**Note:** The X1132A-R6 and X2056-R6 FC HBA for the 62xx series systems are the only cards that can be configured with initiators and targets on the same card.

The following table lists supported combinations:

Onboard ports	Expansion ports	Supported?
Initiator + Target	None	Yes
Initiator + Target	Target only	Yes
Initiator + Target	Initiator only	Yes
Initiator + Target	Initiator + Target	Yes
Initiator only	Target only	Yes
Initiator only	Initiator + Target	Yes
Initiator only	Initiator only	Yes, but no FC SAN support
Initiator only	None	Yes, but no FC SAN support
Target only	Initiator only	Yes
Target only	Initiator + Target	Yes
Target only	Target only	Yes, but no FC disk shelf, V-Series configurations, or FC tape support
Target only	None	Yes, but no FC disk shelf, V-Series configurations, or FC tape support

### Related concepts

[Configuration limits for FC, FCoE, and iSCSI configurations](#) on page 55

### Related references

[FCoE initiator and target combinations](#) on page 40

## Fibre Channel supported hop counts

The maximum supported FC hop count between a host and storage system depends on 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*.

The following table lists 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 supported speeds

Fibre Channel 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.

You should set the port speed to match the speed of the device connected to the port. This is instead of using autonegotiation; a port set to autonegotiation can take longer to reconnect after a takeover/giveback or other interruption.

You should also set all ports used by a given host to the same speed.

### 4-Gb target ports

You can configure 4-Gb onboard ports and 4-Gb expansion adapters to run at the following speeds. Each controller and expansion adapter port can be configured individually. This allows configuration with a different speed from the other ports as needed.

- 4 Gb
- 2 Gb
- 1 Gb

### 8-Gb target ports

You can configure 8-Gb onboard ports and 8-Gb expansion adapters to run at the following speeds. Each controller and expansion adapter port can be configured individually. This allows configuration with a different speed from the other ports as needed.

- 8 Gb

- 4 Gb
- 2 Gb

## Fibre Channel switch configuration best practices

A fixed link speed setting works best, especially for large fabrics, because it provides the best performance for fabric rebuilds. In large fabrics, fixed link speeds can create significant time savings.

Although autonegotiation provides the greatest flexibility, it does not always perform as expected. Also, it adds time to the overall fabric-build sequence because the FC port has to autonegotiate. All switches connected into the fabric have to support NPIV and have it enabled. For FC and iSCSI best practices, refer to 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).

Clustered Data ONTAP uses NPIV to present Fibre Channel targets to a fabric. Make sure to use an NPIV-capable switch and to enable NPIV functionality. For details on which environments are supported, see the online NetApp Interoperability Matrix.

### Related information

[NetApp Interoperability Matrix - support.netapp.com/NOW/products/interoperability/TR-4080: Best Practices for Scalable SAN in Data ONTAP 8.2 Cluster-Mode.pdf](https://support.netapp.com/NOW/products/interoperability/TR-4080:BestPracticesforScalableSANinDataONTAP8.2Cluster-Mode.pdf)

## Host multipathing software requirements

Multipathing software is required on a host any time it 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.

Multipathing software is also known as MPIO (multipath I/O) software. For operating system support for multipathing software, see the Interoperability Matrix.

## 62xx supported fibre channel configurations

62xx controllers are supported in clusters of HA configurations.

The 62xx systems have four onboard 8-Gb FC ports per controller and each one can be configured as either a target or initiator FC port.

The 62xx systems also have vertical I/O slots (slots 1, 11, and 12) that can use a special 4-port 8-Gb FC adapter (Model X2056-R6). Each port on these adapters can be individually configured as either a target or initiator FC port, just like the onboard FC ports.

The 62xx systems also support a 4-port 8-Gb FC adapter (Model X1132A-R6) that can go into any regular I/O slot. Each port on this adapter can be individually configured as either a target or initiator FC port, just like the onboard FC ports.

Each 62xx controller supports 4-Gb and 8-Gb FC target expansion adapters.

## 62xx target port configuration recommendations

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

The port pairs on a 62xx controller that share an ASIC are 0a+0b and 0c+0d.

The following table shows the preferred port usage order for onboard FC target ports. For target expansion adapters, the preferred slot order is listed in the *System Configuration Guide* for the version of Data ONTAP software that the controllers use.

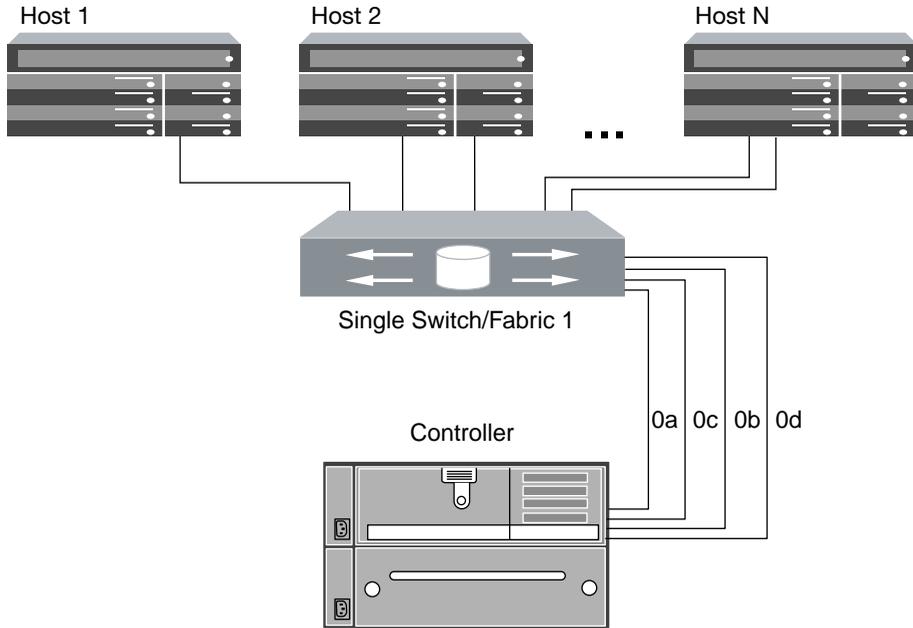
Number of target ports	Ports
1	0a
2	0a, 0c
3	0a, 0c, 0b
4	0a, 0c, 0b, 0d

## 62xx Single-node configurations

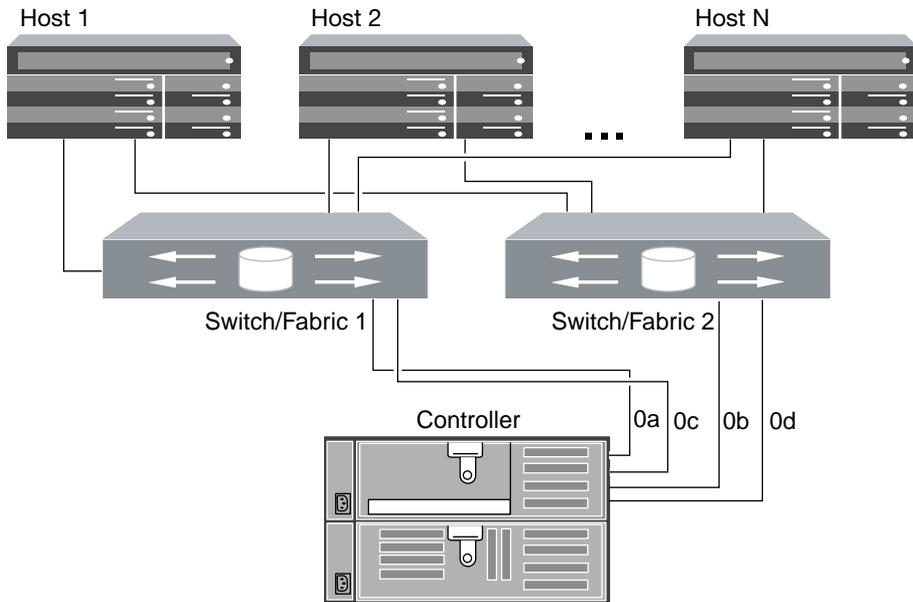
You can connect hosts to single nodes using a single FC fabric or using multiple FC fabrics. If you use multiple paths, multipathing software is required on the host. Direct-attached FC configurations are not supported.

FC switch zoning or portsets are recommended to limit the number of paths between hosts and LUNs in configurations with multiple target ports connected to the same fabric.

**Note:** The FC target port numbers in the following figure are examples. The actual port numbers might vary depending on whether you are using onboard ports or FC target expansion adapters. If you are using FC target expansion adapters, the target port numbers also depend on the expansion slots into which your target expansion adapters are installed.



**Figure 5: 62xx single-fabric single-node configuration**



**Figure 6: 62xx multifabric single-node configuration**

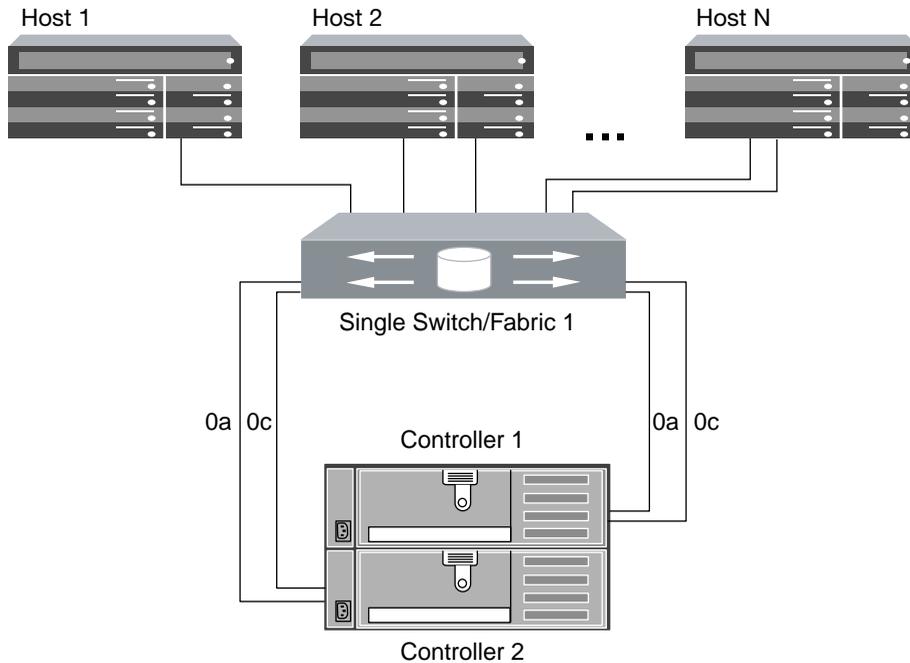
Attribute	Value
Fully redundant	No, due to the single node
Type of fabric	Single fabric or multifabric
Different host operating systems	Yes, with multiple-host configurations
FC ports or adapters	One to the maximum number of supported onboard FC ports per controller  One to the maximum number of supported 4-Gb or 8-Gb FC target expansion adapters
Type of configuration	Single-node configuration

## 62xx: Single-fabric HA configurations

You can connect hosts to both controllers in an HA configuration through a single FC switch.

If you use multiple paths to a LUN, multipathing software is required on the host. FC switch zoning or portsets are recommended to limit the number of paths between hosts and LUNs in configurations with multiple target ports connected to the same fabric. Portsets should be used to help limit the paths that can be used to access LUNs on the controllers.

**Note:** The FC target port numbers in the following figure (0a, 0c) are examples. The actual port numbers might vary depending on whether you are using onboard ports or FC target expansion adapters. If you are using FC target expansion adapters, the target port numbers also depend on the expansion slots into which your target expansion adapters are installed.



**Figure 7: 62xx single-fabric HA configuration**

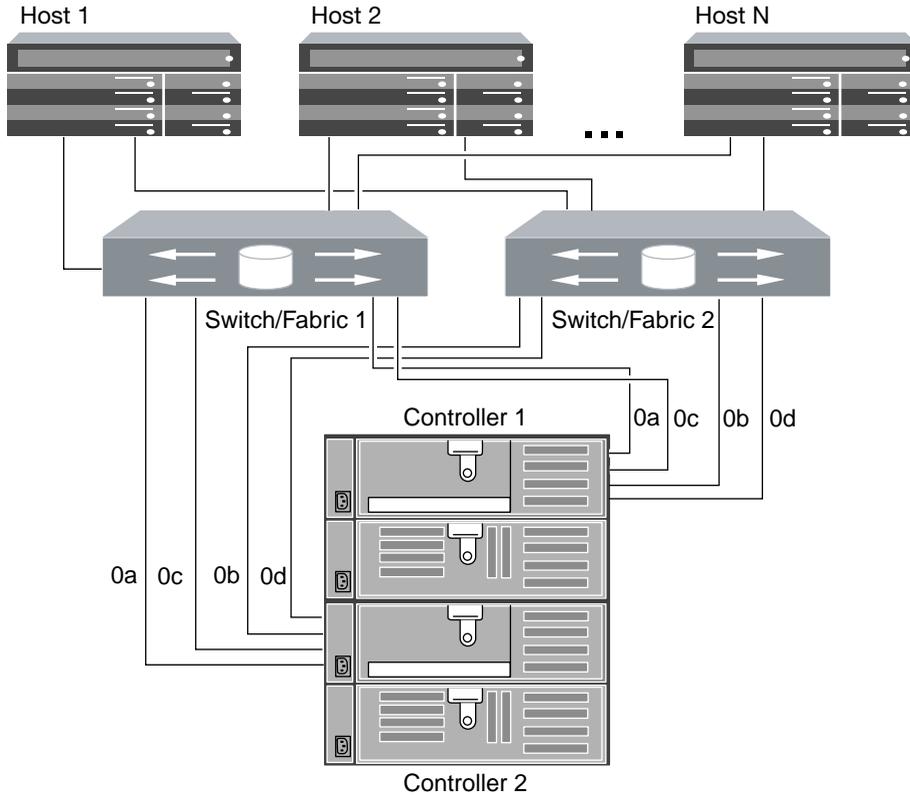
Attribute	Value
Fully redundant	No, due to the single fabric
Type of fabric	Single fabric
Different host operating systems	Yes, with multiple-host configurations
FC ports or adapters	One to the maximum number of supported onboard FC ports per controller One to the maximum number of supported 4-Gb or 8-Gb FC ports using target expansion adapters per controller
Type of configuration	HA configuration

## 62xx: Multifabric HA configurations

You can connect hosts to both controllers in an HA configuration through two or more FC switch fabrics for redundancy.

If you use multiple paths to a LUN, multipathing software is required on the host. FC switch zoning or portsets are recommended to limit the number of paths between hosts and LUNs in configurations with multiple target ports connected to the same fabric. Portsets should be used to help limit the paths that can be used to access LUNs on the controllers.

**Note:** The FC target port numbers in the following figure (0a, 0b, 0c, and 0d) are examples. The actual port numbers might vary depending on whether you are using onboard ports or FC target expansion adapters. If you are using FC target expansion adapters, the target port numbers also depend on the expansion slots into which your target expansion adapters are installed.



**Figure 8: 62xx multifabric HA configurations**

Attribute	Value
Fully redundant	Yes
Type of fabric	Multifabric
Different host operating systems	Yes, with multiple-host configurations
FC ports or adapters	One to the maximum number of supported onboard FC ports per controller One to the maximum number of supported 4-Gb or 8-Gb FC ports using target expansion adapters per controller
Type of configuration	HA configuration

## 6040 and 6080 supported fibre channel configurations

6040 and 6080 controllers are supported in clusters of HA configurations.

The 6040 and 6080 systems have eight onboard 4-Gb FC ports per controller, and each one can be configured as either a target or initiator FC port.

Additional target connections can be supported using 4-Gb and 8-Gb target expansion adapters.

### 6040 and 6080 target port configuration recommendations

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

The port pairs on a 6040 or 6080 controller that share an ASIC are 0a+0b, 0c+0d, 0e+0f, and 0g+0h.

The following table shows the preferred port usage order for onboard FC target ports. For target expansion adapters, the preferred slot order is listed in the *System Configuration Guide* for the version of Data ONTAP software that the controllers use.

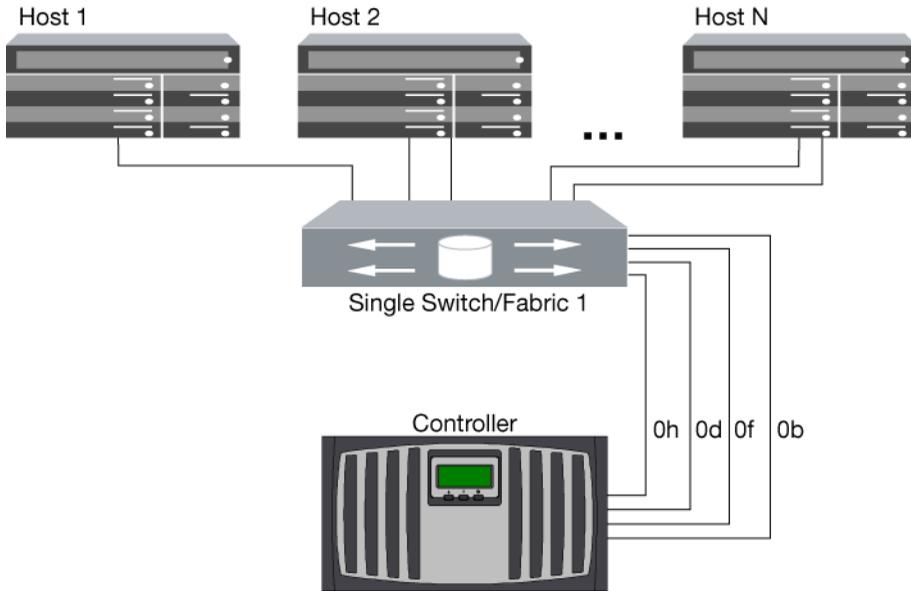
Number of target ports	Ports
1	0h
2	0h, 0d
3	0h, 0d, 0f
4	0h, 0d, 0f, 0b
5	0h, 0d, 0f, 0b, 0g
6	0h, 0d, 0f, 0b, 0g, 0c
7	0h, 0d, 0f, 0b, 0g, 0c, 0e
8	0h, 0d, 0f, 0b, 0g, 0c, 0e, 0a

### 60xx Single-node configurations

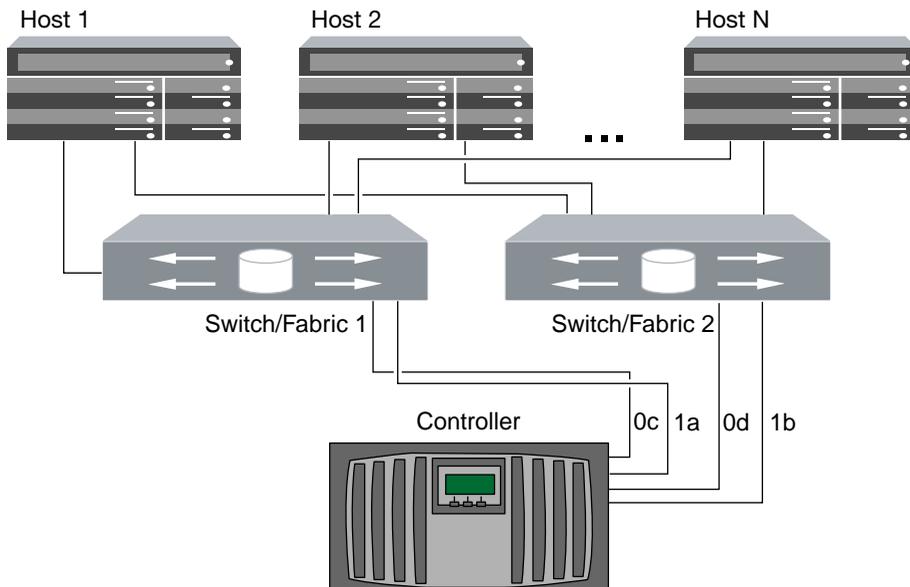
You can connect hosts to a single node using a single FC fabric or using multiple FC fabrics. If you use multiple paths, multipathing software is required on the host. Direct-attached FC configurations are not supported.

FC switch zoning or portsets are recommended to limit the number of paths between hosts and LUNs in configurations with multiple target ports connected to the same fabric.

**Note:** The FC target port numbers in the following figures are examples. The actual port numbers might vary depending on whether you are using onboard ports or FC target expansion adapters. If you are using FC target expansion adapters, the target port numbers also depend on the expansion slots into which your target expansion adapters are installed.



**Figure 9: 60xx single-fabric single-node configuration**



**Figure 10: 60xx multifabric single-node configuration**

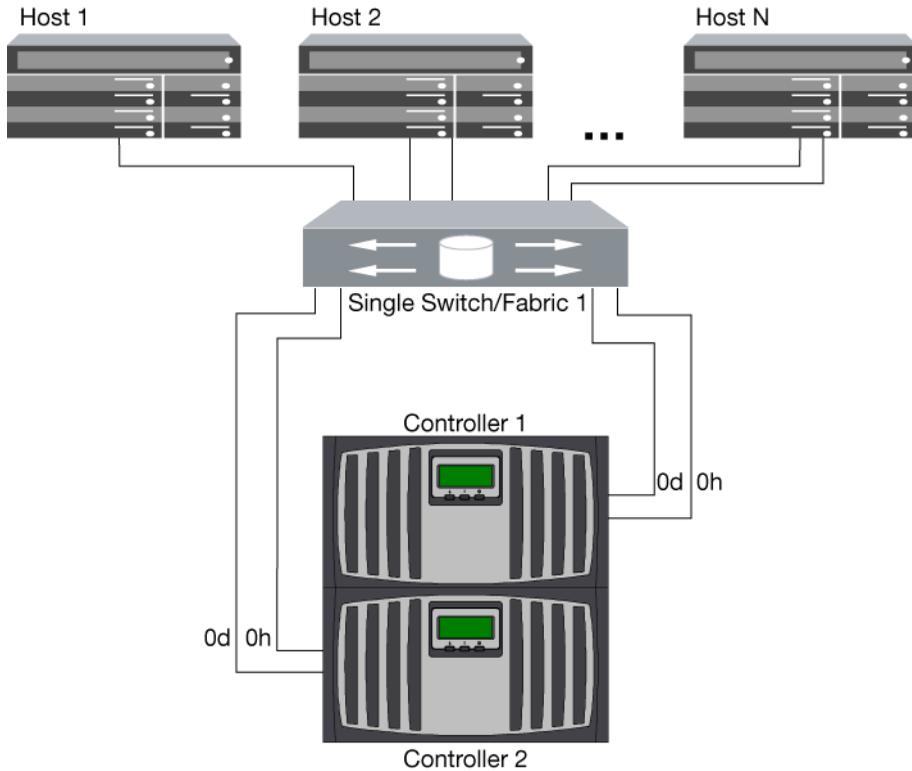
Attribute	Value
Fully redundant	No, due to the single node
Type of fabric	Single fabric or multifabric
Different host operating systems	Yes, with multiple-host configurations
FC ports or adapters	One to the maximum number of supported onboard FC ports per controller  One to the maximum number of supported 4-Gb or 8-Gb FC target expansion adapters
Type of configuration	Single-node configuration

## 6040 and 6080: Single-fabric HA configurations

You can connect hosts to both controllers in an HA pair using a single FC switch.

If you use multiple paths to a LUN, multipathing software is required on the host. FC switch zoning is recommended to limit the number of paths between hosts and LUNs in configurations with multiple target ports connected to the same fabric. Portsets should be used to help limit the paths that can be used to access LUNs on the controllers.

**Note:** The FC target port numbers in the following figure (0d and 0h) are examples. The actual port numbers might vary depending on whether you are using onboard ports or FC target expansion adapters. If you are using FC target expansion adapters, the target port numbers also depend on the expansion slots into which your target expansion adapters are installed.



**Figure 11: 6040 and 6080 single-fabric HA pair**

Attribute	Value
Fully redundant	No, due to the single fabric
Type of fabric	Single fabric
Different host operating systems	Yes, with multiple-host configurations
FC ports or adapters	One to the maximum number of supported onboard FC ports per controller One to the maximum number of supported 4-Gb or 8-Gb FC ports using target expansion adapters per controller
Type of configuration	HA pair

**Related references**

[6040 and 6080 target port configuration recommendations](#) on page 22

### 6040 and 6080: Multifabric HA pair

You can connect hosts to both controllers in an HA pair through two or more FC switch fabrics for redundancy.

If you use multiple paths to a LUN, multipathing software is required on the host. FC switch zoning is recommended to limit the number of paths between hosts and LUNs in configurations with multiple target ports connected to the same fabric. Portsets should be used to help limit the paths that can be used to access LUNs on the controllers.

**Note:** The FC target port numbers in the following figure (5a, 5b, 6a, and 6b) are examples. The actual port numbers might vary depending on whether you are using onboard ports or FC target expansion adapters. If you are using FC target expansion adapters, the target port numbers also depend on the expansion slots into which your target expansion adapters are installed.

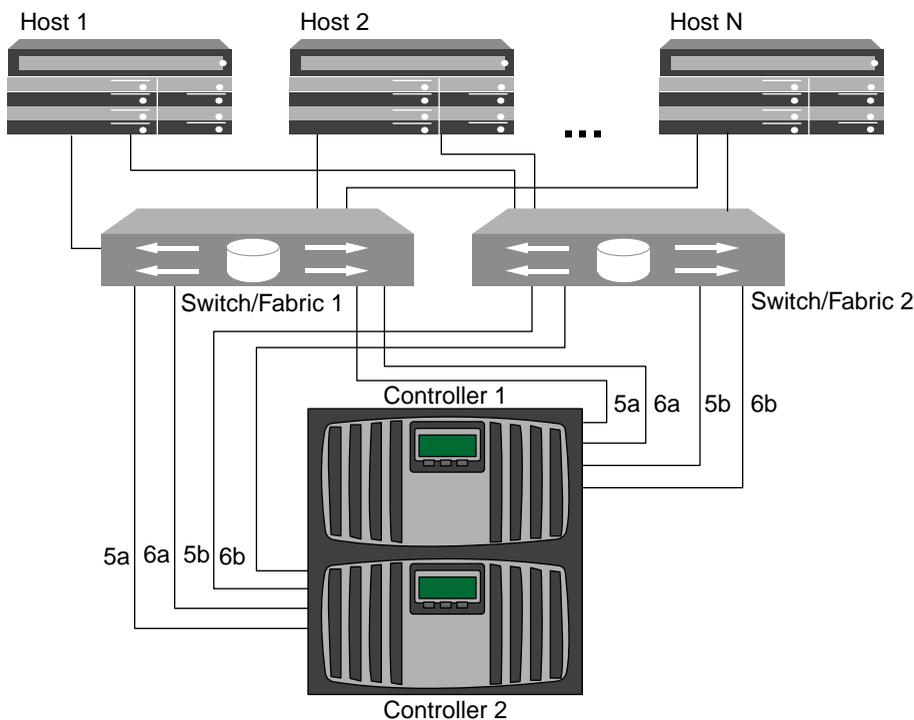


Figure 12: 6040 and 6080 multifabric HA pair

Attribute	Value
Fully redundant	Yes
Type of fabric	Multifabric
Different host operating systems	Yes, with multiple-host configurations

Attribute	Value
FC ports or adapters	One to the maximum number of supported onboard FC ports per controller  One to the maximum number of supported 4-Gb or 8-Gb FC ports using target expansion adapters per controller
Type of configuration	HA pair

### Related references

[6040 and 6080 target port configuration recommendations](#) on page 22

## 32xx supported fibre channel configurations

32xx controllers are supported in clusters of HA configurations.

The 32xx systems have two onboard 4-Gb FC ports per controller that can be configured as FC target or initiator ports. There are also two SAS ports for connecting disk shelves.

Each 32xx controller supports 4-Gb and 8-Gb FC target expansion adapters.

The 32xx systems support a 4-port 8-Gb FC adapter (Model X1132A-R6) that can go into any regular I/O slot. Each port on this adapter can be individually configured as either a target or initiator FC port, just like the onboard FC ports.

### 32xx target port configuration recommendations

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

The following table shows the preferred port usage order for 32xx onboard FC target ports. For target expansion adapters, the preferred slot order is given in the *System Configuration Guide* for the version of Data ONTAP software being used by the controllers.

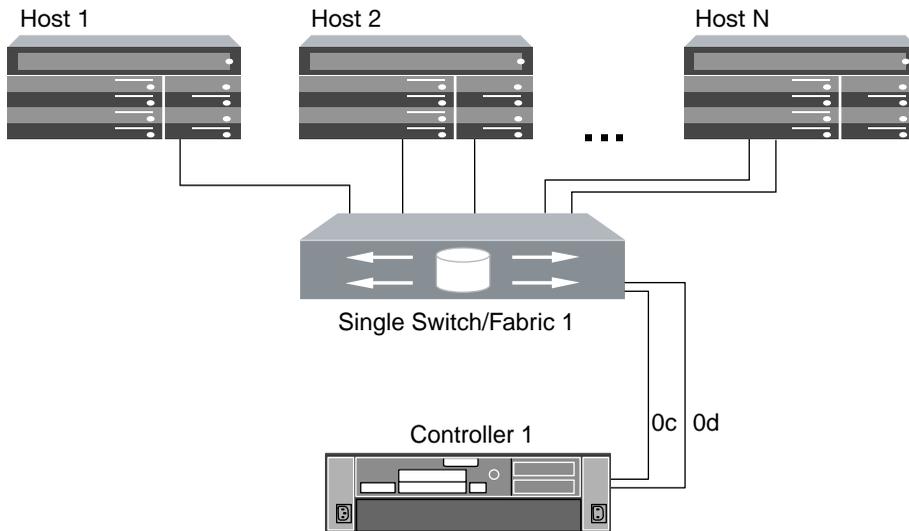
Number of target ports	Ports
1	0c
2	0c, 0d

### 32xx Single-node configurations

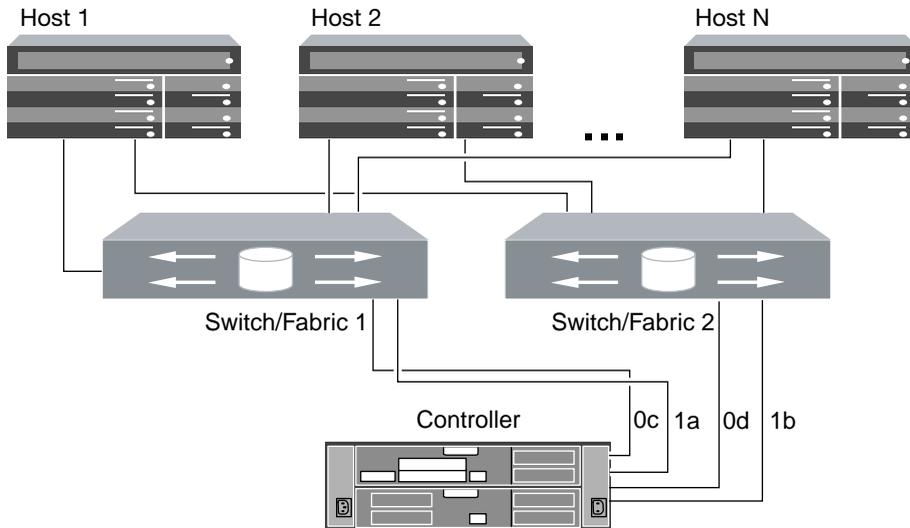
You can connect hosts to a single node using a single FC fabric or using multiple FC fabrics. Direct-attached FC configurations are not supported.

If you use multiple paths to a LUN, multipathing software is required on the host. FC switch zoning or portsets are recommended to limit the number of paths between hosts and LUNs in configurations with multiple target ports connected to the same fabric.

**Note:** The FC target port numbers in the following figures are examples. The actual port numbers might vary depending on whether you are using onboard ports or FC target expansion adapters. If you are using FC target expansion adapters, the target port numbers also depend on the expansion slots into which your target expansion adapters are installed.



**Figure 13: 32xx single-fabric single-node configuration**



**Figure 14: 32xx multifabric single-node configuration**

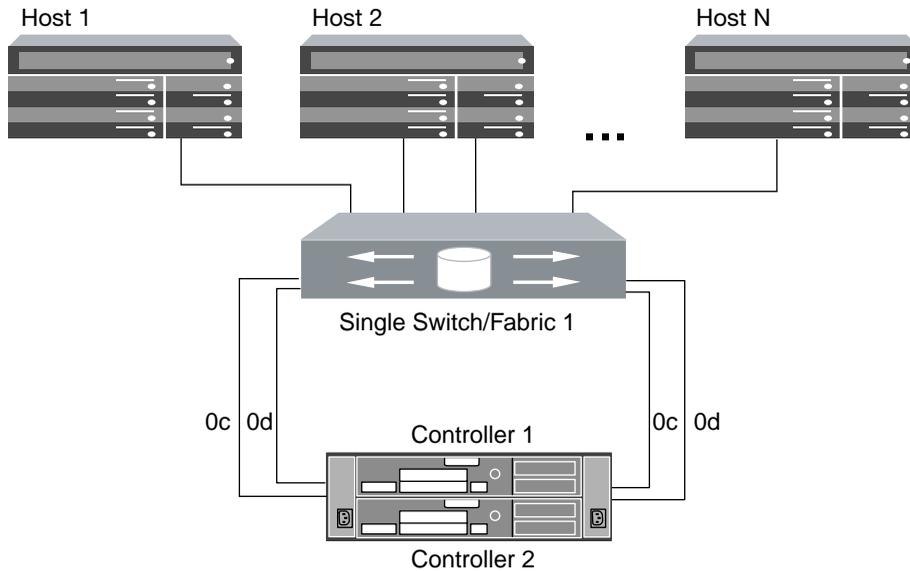
Attribute	Value
Fully redundant	No, due to the single node
Type of fabric	Single fabric or multifabric
Different host operating systems	Yes, with multiple-host configurations
FC ports or adapters	One to the maximum number of supported onboard FC ports per controller  One to the maximum number of supported 4-Gb or 8-Gb FC target expansion adapters
Type of configuration	Single-node configuration

### 32xx: Single-fabric HA configurations

You can connect hosts to both controllers in an HA configuration using a single FC switch.

If you use multiple paths to a LUN, multipathing software is required on the host. FC switch zoning or portsets are recommended to limit the number of paths between hosts and LUNs in configurations with multiple target ports connected to the same fabric. Portsets should be used to help limit the paths that can be used to access LUNs on the controllers

**Note:** The FC target port numbers in the following figure (0c, 0d) are examples. The actual port numbers might vary depending on whether you are using onboard ports or FC target expansion adapters. If you are using FC target expansion adapters, the target port numbers also depend on the expansion slots into which your target expansion adapters are installed.



**Figure 15: 32xx single-fabric HA configuration**

Attribute	Value
Fully redundant	No, due to the single fabric
Type of fabric	Single fabric
Different host operating systems	Yes, with multiple-host configurations
FC ports or adapters	One to the maximum number of supported onboard FC ports per controller One to the maximum number of supported 4-Gb or 8-Gb FC target expansion adapters
Type of configuration	HA configuration

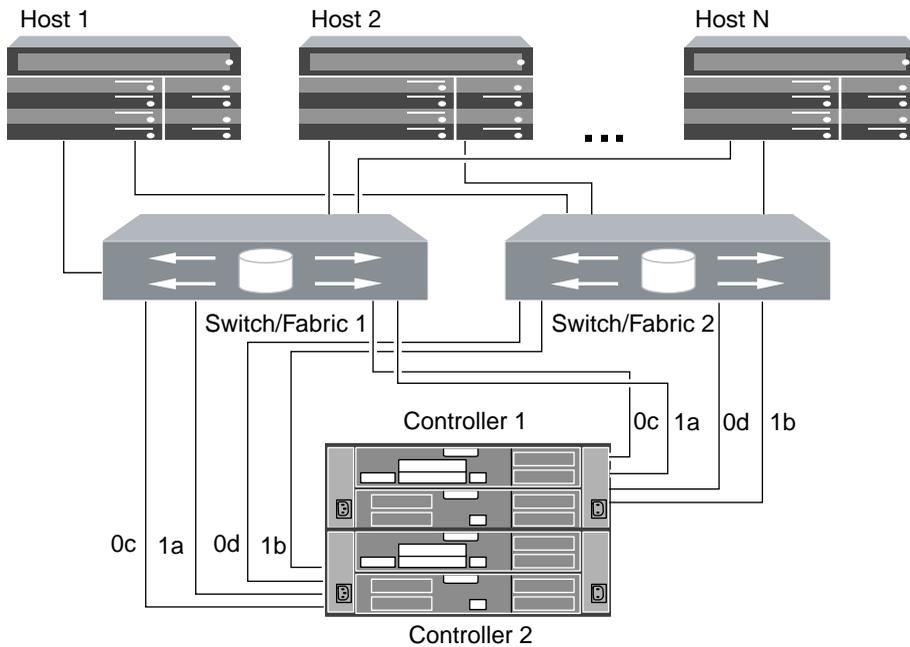
### 32xx: Multifabric HA configurations

You can connect hosts to both controllers in an HA configuration using two or more FC switch fabrics for redundancy.

If you use multiple paths to a LUN, multipathing software is required on the host. FC switch zoning or portsets are recommended to limit the number of paths between hosts and LUNs in configurations with multiple target ports connected to the same fabric. Portsets should be used to help limit the paths that can be used to access LUNs on the controllers

**Note:** The FC target port numbers in the following figure (0c, 0d, 1a, 1b) are examples. The actual port numbers might vary depending on whether you are using onboard ports or FC target

expansion adapters. If you are using FC target expansion adapters, the target port numbers also depend on the expansion slots into which your target expansion adapters are installed.



**Figure 16: 32xx multifabric HA configuration**

Attribute	Value
Fully redundant	Yes
Type of fabric	Multifabric
Different host operating systems	Yes, with multiple-host configurations
FC ports or adapters	One to the maximum number of supported onboard FC ports per controller  One to the maximum number of supported 4-Gb or 8-Gb FC target expansion adapters
Type of configuration	HA configuration

## 31xx supported fibre channel configurations

31xx systems are available in HA pairs. Clustered Data ONTAP operating does not support single controller systems on 31xx systems.

The 31xx systems have four onboard 4-Gb FC ports per controller and each port can be configured as either an FC target port or an initiator port. For example, you can configure two ports as SAN targets and two ports as initiators for disk shelves.

Each 31xx controller supports 4-Gb and 8-Gb FC target expansion adapters.

### 31xx target port configuration recommendations

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

The port pairs on a 31xx controller that share an ASIC are 0a+0b and 0c+0d.

The following table shows the preferred port usage order for onboard FC target ports. For target expansion adapters, the preferred slot order is given in the *System Configuration Guide* for the version of Data ONTAP software being used by the controllers.

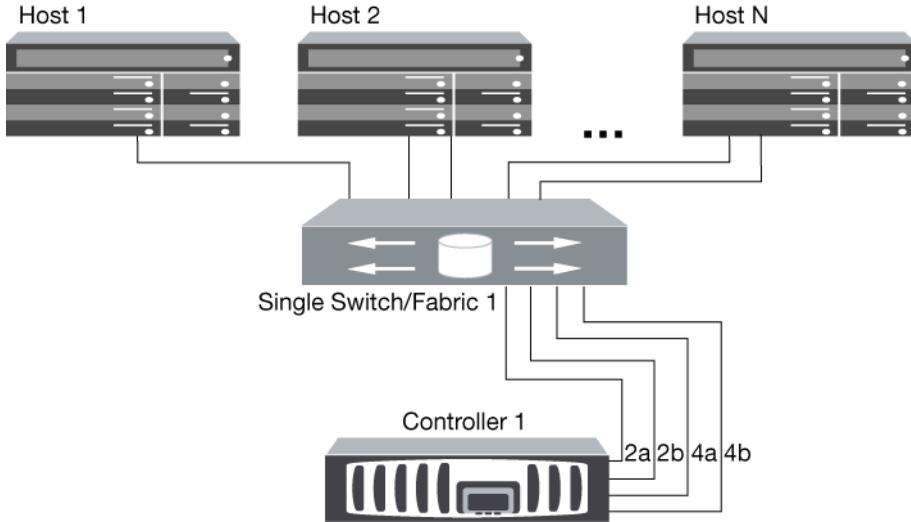
Number of target ports	Ports
1	0d
2	0d, 0b
3	0d, 0b, 0c
4	0d, 0b, 0c, 0a

### 31xx Single-node configurations

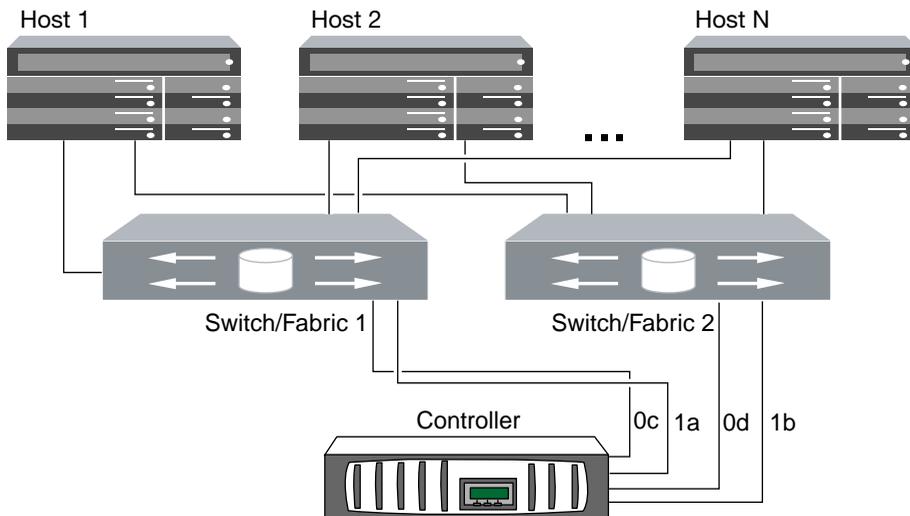
You can connect hosts to a single node using a single FC fabric or using multiple FC fabrics. If you use multiple paths, multipathing software is required on the host. Direct-attached FC configurations are not supported.

FC switch zoning is recommended to limit the number of paths between hosts and LUNs in configurations with multiple target ports connected to the same fabric.

**Note:** The FC target port numbers in the following figures are examples. The actual port numbers might vary depending on whether you are using onboard ports or FC target expansion adapters. If you are using FC target expansion adapters, the target port numbers also depend on the expansion slots into which your target expansion adapters are installed.



**Figure 17: 31xx single-fabric single-node configuration**



**Figure 18: 31xx multifabric single-node configuration**

Attribute	Value
Fully redundant	No, due to the single node
Type of fabric	Single fabric or multifabric
Different host operating systems	Yes, with multiple-host configurations

Attribute	Value
FC ports or adapters	One to the maximum number of supported onboard FC ports per controller  One to the maximum number of supported 4-Gb or 8-Gb FC target expansion adapters
Type of configuration	Single-node configuration

### 31xx: Single-fabric HA pair

You can connect hosts to both controllers in an HA pair using a single FC switch.

If you use multiple paths to a LUN, multipathing software is required on the host. FC switch zoning is recommended to limit the number of paths between hosts and LUNs in configurations with multiple target ports connected to the same fabric. Portsets should be used to help limit the paths that can be used to access LUNs on the controllers.

**Note:** The FC target port numbers in the following figure (0b and 0d) are examples. The actual port numbers might vary depending on whether you are using onboard ports or FC target expansion adapters. If you are using FC target expansion adapters, the target port numbers also depend on the expansion slots into which your target expansion adapters are installed.

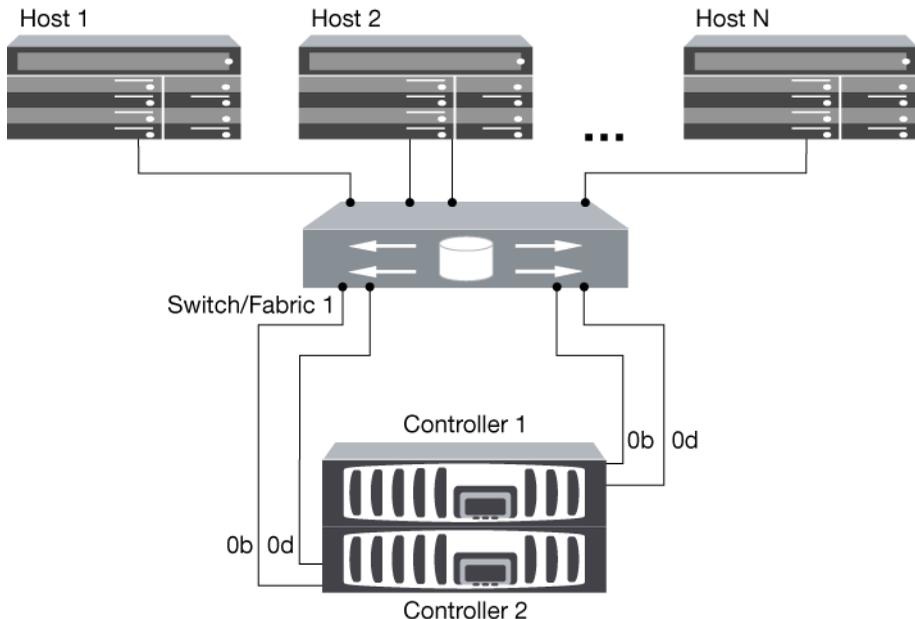


Figure 19: 31xx single-fabric HA pair

Attribute	Value
Fully redundant	No, due to the single fabric
Type of fabric	Single fabric
Different host operating systems	Yes, with multiple-host configurations
FC ports or adapters	One to the maximum number of supported onboard FC ports per controller  One to the maximum number of supported 4-Gb or 8-Gb FC target expansion adapters
Type of configuration	HA pair

### Related references

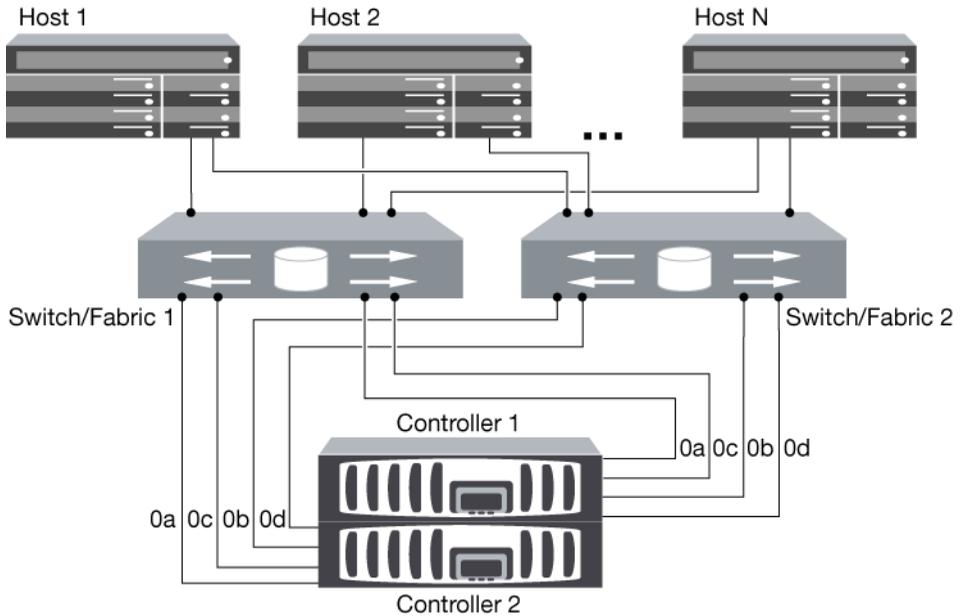
[31xx target port configuration recommendations](#) on page 32

## 31xx: Multifabric HA pair

You can connect hosts to both controllers in an HA pair using two or more FC switch fabrics for redundancy.

If you use multiple paths to a LUN, multipathing software is required on the host. FC switch zoning is recommended to limit the number of paths between hosts and LUNs in configurations with multiple target ports connected to the same fabric. Portsets should be used to help limit the paths that can be used to access LUNs on the controllers.

**Note:** The FC target port numbers in the following figure (0a, 0b, 0c, and 0d) are examples. The actual port numbers might vary depending on whether you are using onboard ports or FC target expansion adapters. If you are using FC target expansion adapters, the target port numbers also depend on the expansion slots into which your target expansion adapters are installed.



**Figure 20: 31xx multifabric HA pair**

Attribute	Value
Fully redundant	Yes
Type of fabric	Multifabric
Different host operating systems	Yes, with multiple-host configurations
FC ports or adapters	One to the maximum number of supported onboard FC ports per controller  One to the maximum number of supported 4-Gb or 8-Gb FC target expansion adapters
Type of configuration	HA pair

**Related references**

[31xx target port configuration recommendations](#) on page 32

## FAS2240 Supported configurations

The FAS2240 is supported for FC only in single-node configurations using the X1150A-R6 FC adapter. FAS2240 single-node FC configurations cannot be joined to other nodes in a cluster. To join the single node to a cluster, you must remove the FC adapter and replace it with the 10-Gb Ethernet

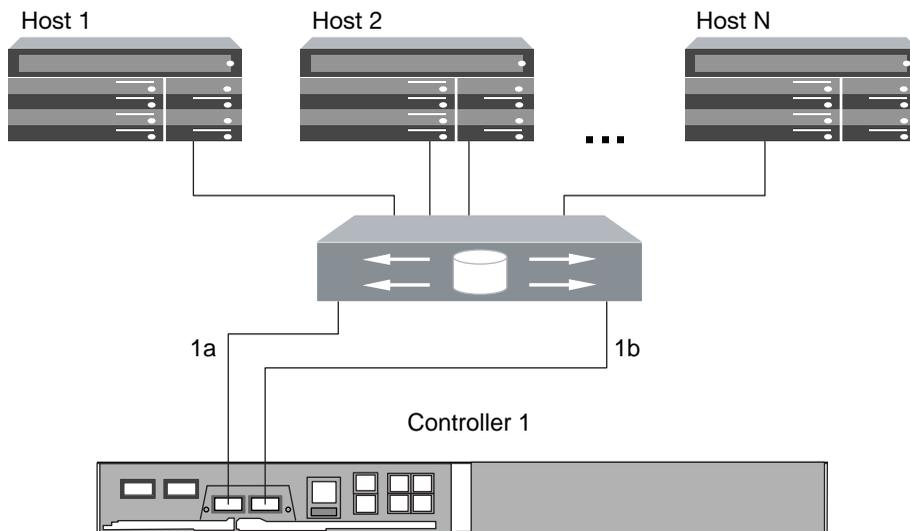
adapter required for the cluster interconnect. iSCSI is supported for both single-node and multi-node clusters.

## FAS2240 Single-fabric Single-node configurations

You can connect hosts to a single node using a single FC switch. If you use multiple paths, multipathing software is required on the host. Direct-attached FC configurations are not supported.

FC switch zoning or portsets are recommended to limit the number of paths between hosts and LUNs in configurations with multiple target ports connected to the same fabric.

**Note:** The following illustration 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 between the SAN configuration for the two models.



**Figure 21: FAS2240 single-fabric single-node configuration**

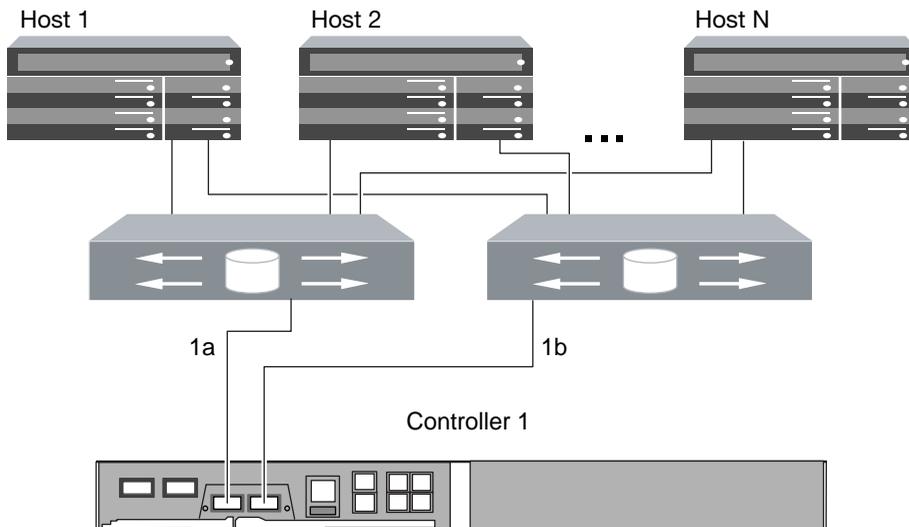
Attribute	Value
Fully redundant	No, due to the single fabric and single node
Type of fabric	Single fabric
Different host operating systems	Yes, with multiple-host configurations
FC ports or adapters	One or two ports on the 2-port 8-Gb FC expansion adapter configured as targets
Type of configuration	Single-node configuration

## FAS2240 Multifabric single-node configurations

You can connect hosts to one node using two or more FC switch fabrics for redundancy. Direct-attached FC configurations are not supported.

If you use multiple paths to a LUN, multipathing software is required on the host. FC switch zoning or portsets are recommended to limit the number of paths between hosts and LUNs in configurations with multiple target ports connected to the same fabric.

**Note:** The following illustration 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 between the two models.



**Figure 22: FAS2240 multifabric single-node configuration**

Attribute	Value
Fully redundant	No, due to the single node
Type of fabric	Multifabric
Different host operating systems	Yes, with multiple-host configurations
FC ports or adapters	Both ports on the supported 2-port 8-Gb FC expansion adapter configured as targets
Type of configuration	Single-node configuration

## Fibre Channel over Ethernet overview

---

Fibre Channel over Ethernet (FCoE) is a model for connecting hosts to storage systems. As with Fibre Channel (FC), FCoE maintains existing FC management and controls. However, the hardware transport is a lossless 10-Gb Ethernet network.

Setting up an FCoE connection on the host or storage requires one or more supported converged network adapters (CNAs) connected to a supported FCoE switch. The CNA is a consolidation point and effectively serves as both an FC HBA and an Ethernet adapter

The CNA is presented to the host and target as both an FCoE Initiator HBA and a 10-Gb Ethernet adapter. The FCoE Initiator HBA portion of the CNA handles the FCoE traffic when traffic is sent and received as FC frames mapped into Ethernet packets (FC over Ethernet). The Ethernet adapter portion of the CNA handles the standard Ethernet IP traffic, such as iSCSI, CIFS, NFS, and HTTP, for the host. Both the FCoE and standard Ethernet portions of the CNA communicate over the same Ethernet port, which connects to the FCoE switch.

The FCoE target adapter is also sometimes called a "unified target adapter" or UTA. Like the CNA, the UTA supports both FCoE and regular Ethernet traffic.

You should configure jumbo frames (MTU = 9000) for the Ethernet adapter portion of the CNA. You cannot change the MTU for the FCoE portion of the adapter.

**Note:** Unified target adapters (UTAs) are 10-Gb converged network adapters (CNAs) that you install in your storage systems.

In general, you configure and use FCoE connections just like traditional FC connections. You can use UTAs for non-FCoE IP traffic such as NFS, CIFS, or iSCSI.

**Note:** For detailed information about how to set up and configure your host to run FCoE, see your host documentation.

## Recommended number of paths to avoid single points of failure

You should have a minimum of 2 paths per LUN connecting to each node in your cluster. This eliminates single points of failure and enables the system to survive component failures.

If you have more than four nodes in your cluster or more than four target ports in any of your nodes, portsets can be used to limit the number of paths that can be used to access LUNs on your controllers.

## 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 or 8-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.

### Related references

[Rules for combining onboard and expansion FC ports](#) on page 14

## Fibre Channel over Ethernet supported hop count

The maximum supported 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 over Ethernet supported configurations

Supported FCoE native configurations include single-fabric and multi-fabric configurations.

Supported storage systems with native FCoE target expansion adapters (called unified target adapters or UTAs) are the 62xx, 60xx, 32xx, and 31xx.

The host FCoE initiator with FC target configuration is also supported on all storage systems using an FCoE/DCB switch. Direct-attached configurations are not supported in FCoE using UTAs and are not recommended in iSCSI using UTAs.

**Note:** While iSCSI configurations are supported on standard Ethernet switches and follow the same supportability requirements as other IP-based protocols, FCoE configurations require Ethernet switches that explicitly support FCoE features. If a CNA is used as a 10-Gb Ethernet card and only IP protocols are going to be used then a switch that supports FCoE for connectivity is not required.

FCoE configurations are validated through the same interoperability and quality assurance process as fibre channel switches. Supported configurations are listed in the Interoperability Matrix Tool (IMT). 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.

### Related information

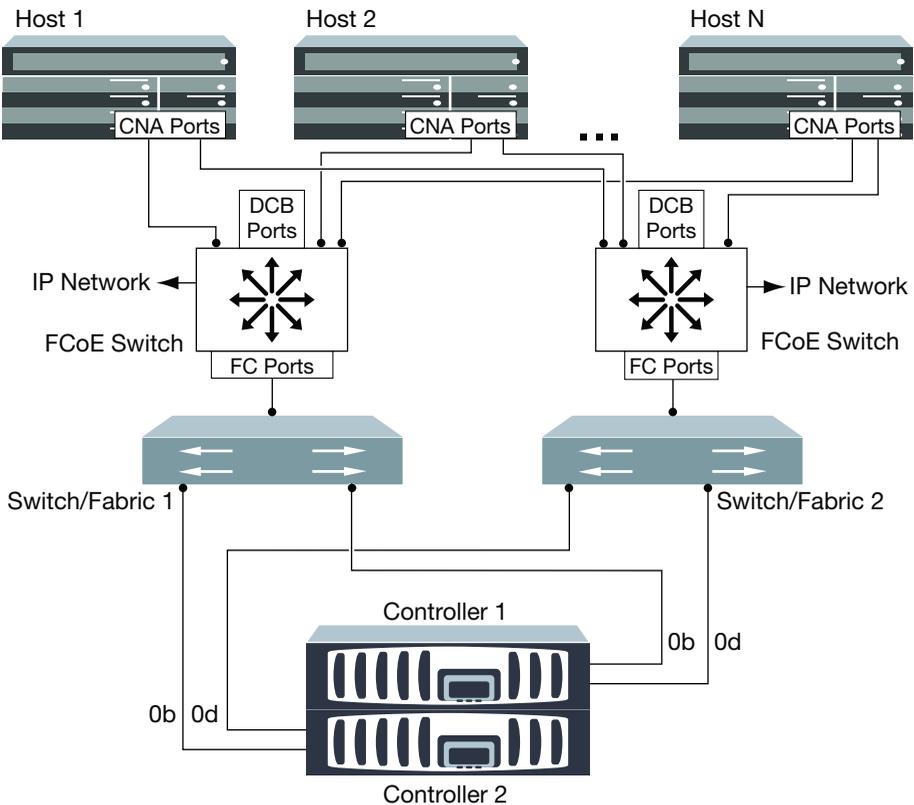
*[NetApp Interoperability Matrix - support.netapp.com/NOW/products/interoperability/](http://support.netapp.com/NOW/products/interoperability/)*

## FCoE initiator to FC target configuration

You can connect hosts to both controllers in an HA pair using host FCoE initiators (CNAs) through FCoE switches to FC target ports. This requires an FCoE switch that also has FC ports.

The host FCoE initiator always connects to a supported FCoE switch. The FCoE switch can connect directly to an FC target, or can connect through FC switches to the FC target.

**Note:** The FC target expansion adapter port numbers (0b and 0d) in the following figure are examples. The actual port numbers might vary, depending on the expansion slot in which the FC target expansion adapter is installed.



**Figure 23: FCoE initiator to FC dual-fabric HA pair**

Attribute	Value
Fully redundant	Yes
Type of fabric	Dual fabric
Different host operating systems	Yes, with multiple-host configurations
FC ports or adapters	One to the maximum number of supported onboard FC ports per controller One to the maximum number of supported 4-Gb or 8-Gb FC ports per controller using FC target expansion adapters
Multipathing required	Yes
Type of configuration	HA pair

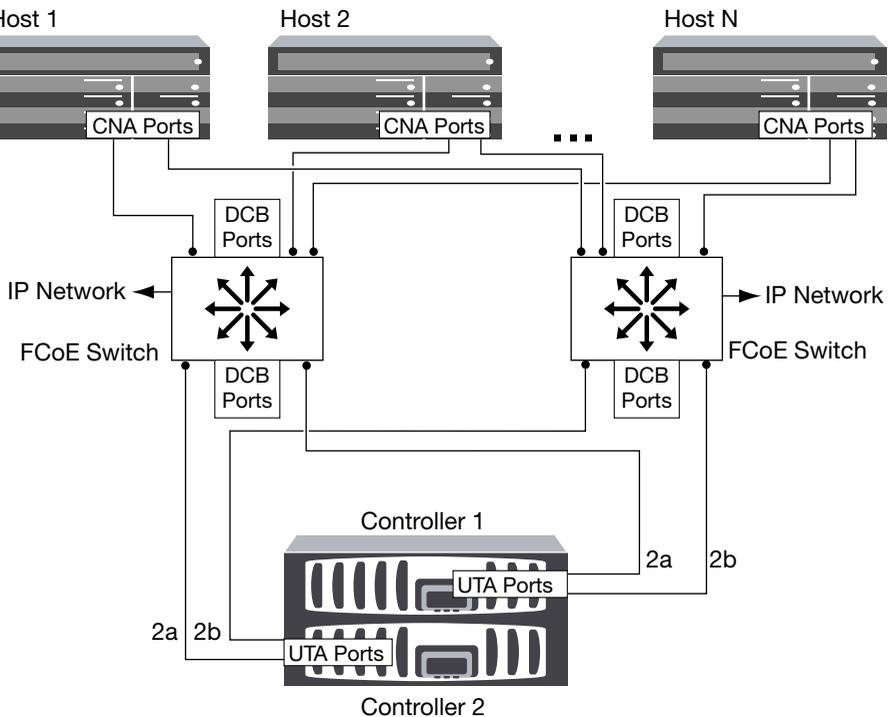
## FCoE end-to-end configuration

You can connect hosts to both controllers in an HA pair. This connection is achieved host using FCoE initiators (CNAs) through DCB switches to FCoE target ports.

The FCoE target adapter is also called a unified target adapter or UTA. Like the CNA in a host, the UTA supports both FCoE and regular Ethernet traffic.

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 FCoE switches to each other, the switches must run a firmware version that supports Ethernet ISLs.

**Note:** The FCoE target expansion adapter port numbers (2a and 2b) in the following figure are examples. The actual port numbers might vary, depending on the expansion slot in which the FCoE target expansion adapter is installed.



**Figure 24: FCoE end-to-end**

Attribute	Value
Fully redundant	Yes
Type of fabric	Dual fabric
Different host operating systems	Yes, with multiple host-configurations

Attribute	Value
FCoE ports or adapters	One or more FCoE target expansion adapters per controller
Multipathing required	Yes
Type of configuration	HA pair

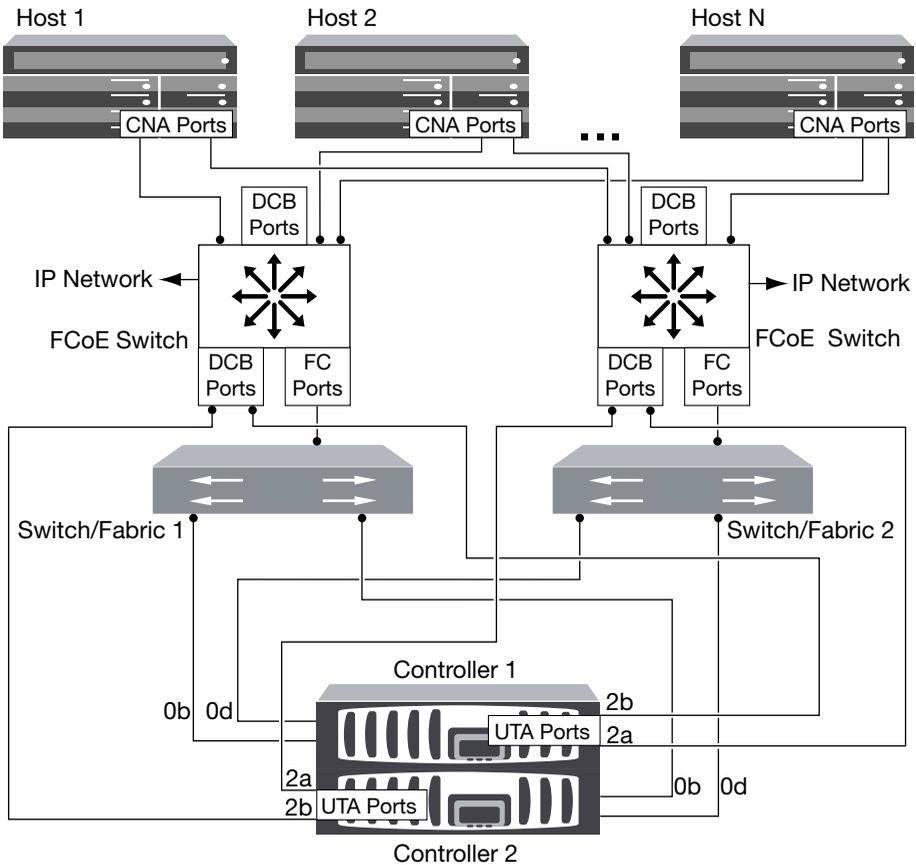
## FCoE mixed with FC

You can connect hosts to both controllers in an HA pair. This connection is achieved using host FCoE initiators (CNAs) through FCoE switches to FCoE and FC mixed target ports.

The FCoE target adapter is also called a unified target adapter or UTA. Like the CNA in a host, the UTA supports both FCoE and regular Ethernet traffic.

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 FCoE switches to each other, the switches must run a firmware version that supports Ethernet ISLs.

**Note:** The FCoE target expansion adapter port numbers (2a and 2b) and FC target port numbers (0b and 0d) are examples. The actual port numbers might vary, depending on the expansion slots in which the FCoE target expansion adapter and FC target expansion adapter are installed.



**Figure 25: FCoE mixed with FC**

Attribute	Value
Fully redundant	Yes
Type of fabric	Dual fabric
Different host operating systems	Yes, with multiple-host configurations
FC/FCoE ports or adapters	One to the maximum number of supported onboard FC ports per controller One or more FCoE target expansion adapters per controller At least one 4-Gb or 8-Gb FC target expansion adapter per controller
Multipathing required	Yes
Type of configuration	HA pair

## FCoE mixed with IP storage protocols

You can connect hosts to both controllers in an HA pair using host FCoE initiators (CNAs) through FCoE switches to FCoE target ports. You can also run non-FCoE Ethernet traffic through the same switches.

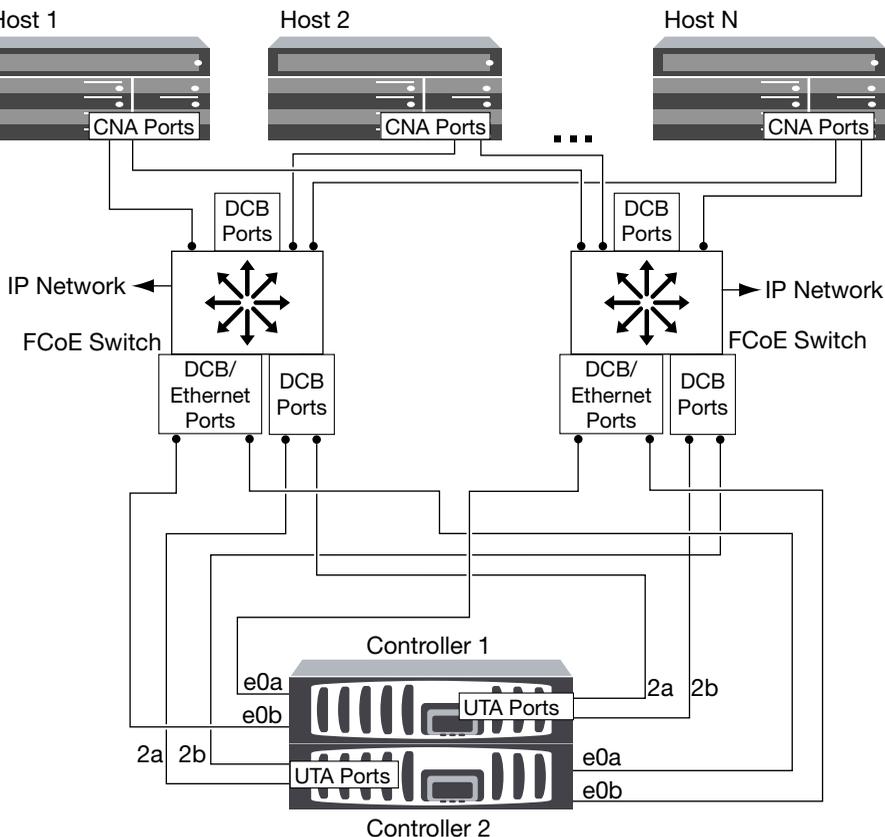
The FCoE target adapter is also called a unified target adapter or UTA. Like the CNA in a host, the UTA supports both FCoE and regular Ethernet traffic.

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 FCoE switches to each other over Ethernet, the switches must run a firmware version that supports Ethernet ISLs. You can also connect switches using FC ISLs.

The FCoE ports are connected to DCB ports on the FCoE switches. FCoE ports cannot use traditional link aggregation to a single switch. Cisco switches support a special type of link aggregation called *Virtual Port Channel* that does support FCoE. A Virtual Port Channel aggregates individual links to two switches. You can also use the Virtual Port Channel for other Ethernet traffic. For more information about configuring Virtual Port Channels, see Technical Report TR-3800: *Fibre Channel over Ethernet (FCoE) End-to-End Deployment Guide* or your Cisco switch documentation.

Ports used for traffic other than FCoE, including NFS, CIFS, iSCSI, and other Ethernet traffic, can use regular Ethernet ports on the FCoE switches.

**Note:** The FCoE target expansion adapter port numbers (2a and 2b) and the Ethernet port numbers (e0a and e0b) in the following figure are examples. The actual port numbers might vary, depending on the expansion slots in which the FCoE target expansion adapters are installed.



**Figure 26: FCoE mixed with IP storage protocols**

Attribute	Value
Fully redundant	Yes
Type of fabric	Dual fabric
Different host operating systems	Yes, with multiple-host configurations
FCoE ports or adapters	One or more FCoE target expansion adapters per controller
Multipathing required	Yes
Type of configuration	HA pair

**Related information**

*TR-3800: Fibre Channel over Ethernet (FCoE) End-to-End Deployment Guide - [media.netapp.com/documents/TR-3800.pdf](http://media.netapp.com/documents/TR-3800.pdf)*

## Fibre Channel and FCoE zoning

---

An FC or FCoE zone is a subset of the fabric that consists of a group of FC or FCoE ports or nodes that can communicate with each other. You must contain the nodes within the same zone to allow communication.

### Reasons for zoning

- Zoning reduces or eliminates *cross talk* 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.
- Zoning increases security because there is limited access between different nodes of a SAN.
- 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 anytime 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 recommended.
- 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. Depending on the switch vendor, either World Wide Node Names or World Wide Port Names can be used. You should always use World Wide Port Name 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 can access a LIF on every node, but should only be zoned to the nodes that are providing the paths used for 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 if a Vserver is moved to a different node.

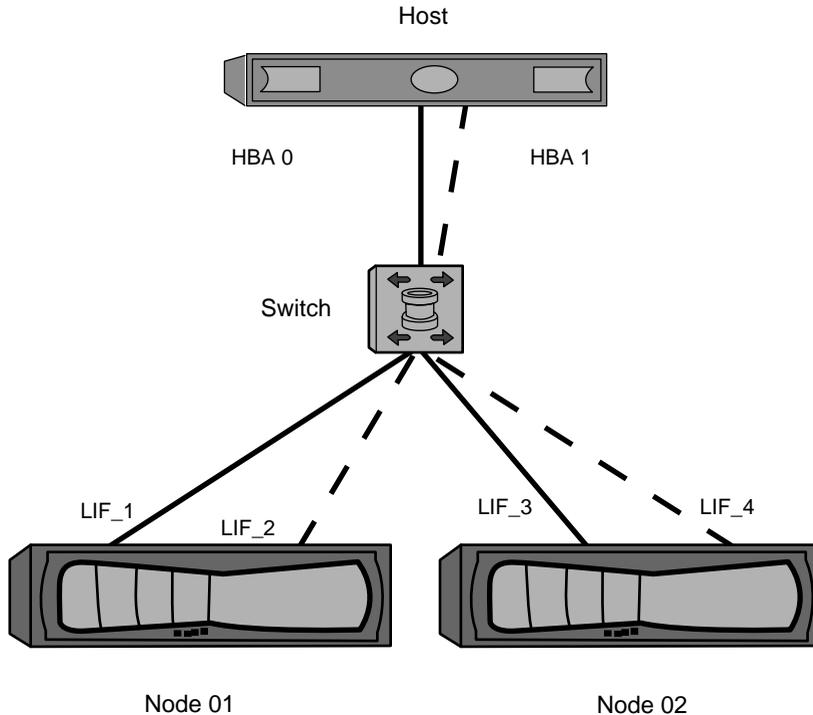
Single-fabric configurations are supported, but are not considered highly available. The failure of a single component can cause loss of access to data.

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\_1, and LIF\_4

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



**Figure 27: 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. The number of paths can be verified by using configuration limits at the end of this document.

## Dual-fabric HA pair zoning

In dual fabric configurations, you can connect each host initiator to each storage node. Each host initiator uses a different switch to access the storage 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.

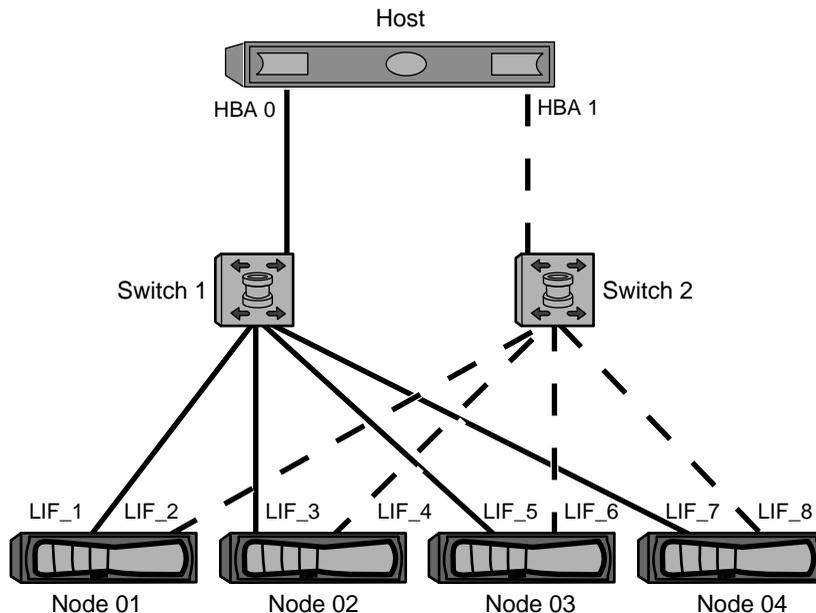
- Zone 1: HBA 0, LIF\_1, LIF\_3, LIF\_5, and LIF\_7

- 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. Vservers have access to all iSCSI and FCP LIFs on every node on a controller. To reduce the paths from a Vserver to the host, the use of portsets is recommended.

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



**Figure 28: 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.

## Shared SAN configurations

---

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

To access NetApp storage systems and other vendors' storage systems from a single host requires:

- Native Host OS multipathing or VERITAS DMP is used for multipathing (see exception for EMC PowerPath co-existence below)
- NetApp configuration requirements (such as timeout settings) as specified in the appropriate NetApp Host Utilities documents have been met

Native Host OS multipathing in combination with EMC PowerPath is supported for the following configurations. For configurations that do meet these requirements, a PVR is required to determine supportability.

Host	Supported configuration
Windows	EMC CX3, CX4, and VNX arrays with PowerPath 4.5 or later connected to a NetApp storage system using Data ONTAP DSM for Windows MPIO. This applies equally to Windows 2003, Windows 2008, and Windows 2012.
Solaris	EMC CX3, CX4, and VNX arrays PowerPath 5 or later connected to a NetApp storage system using MPxIO
AIX	EMC CX3, CX4, and VNX arrays PowerPath 5 or later connected to a NetApp storage system using AIX MPIO

## ALUA Configurations

---

Clustered Data ONTAP always uses asymmetric logical unit access (ALUA). 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.

Clustered Data ONTAP always uses ALUA for both FC and iSCSI paths. 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: [support.netapp.com](http://support.netapp.com)*

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

# 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 configurations always use FlexVol volumes

Volumes containing LUNs must be FlexVol volumes. SAN protocols can only be used with Vservers with FlexVol volumes. Infinite Volumes are not supported for SAN.

In this document, "volume" always means FlexVol volume. "Vserver" always means Vserver with FlexVol volumes.

## Supported number of nodes for SAN configurations

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

In Data ONTAP 8.2, the FAS2220, FAS2240, 3160, 3170, and 3240 platforms support a maximum of four nodes per cluster. All other platforms support a maximum of eight nodes per cluster.

If any node in the cluster is configured for FC, FCoE, or iSCSI, the cluster is limited to the SAN node limits.

### Related references

[SAN configuration limits](#) on page 59

[SAN configuration limits specifically for the FAS2220](#) on page 61

## 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 Fibre Channel controllers.

Parameter	Definition
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. If you have more than 4 nodes, the best practice is to use portsets to limit paths to a subset of the nodes in the cluster. You should not exceed more than 8 paths.
Maximum LUN size	The maximum size of an individual LUN on the respective operating system.
Vservers	The total number of Vservers, including the default node Vserver. There are limits on Vservers per node and Vservers 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.
FC port fan-in	The maximum number of hosts that can connect to a single FC port on a controller. Connecting the maximum number of hosts 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 Fibre Channel logical interfaces (LIFs) that can be defined on a single physical FC port.
FC port fan-out	The maximum number of LUNs mapped to a host through a FC target port on a controller.
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: Maximum sessions = 8 * System Memory divided by 512 MB.
IP LIFs per port	The maximum combined number of iSCSI, NFS, CIFS, and FlexCache logical interfaces (LIFs) that can be defined on a single physical Ethernet port.

Parameter	Definition
iSCSI LIFs per portset	The maximum number of iSCSI logical interfaces that can be assigned to a single portset.
Hosts per controller (FC)	The maximum number of hosts that can connect to a controller. Connecting the maximum number of hosts is generally not recommended, and you might need to tune the FC queue depths on the host to achieve this maximum value. This value assumes two initiators per host.  <b>Note:</b> Virtualized hosts, using NPIV to access the fabric with virtual WWPNs, count as additional hosts for the purpose of this limitation.
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, /vol/abc/def 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.
Portsets per node	The maximum number of portsets that can be created on a single node.

### Related tasks

[Calculating queue depth](#) on page 63

## 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	16	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	4.x: 256 Local drives, CD-ROM, and so on count against this value.
Maximum paths from host to LUN	8	24 (max of 2,048 per host)	11iv2: 8 11iv3: 32	16	16	4.x: 8 (max of 1,024 per host)
Maximum paths from host to storage solution	1,024	1,024	11iv2: 8 11iv3: 32	16	16	4.x: 8 (max of 1,024 per host)
Max 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 55

## SAN configuration limits

Each system model has configuration limits for reliable operation. The values listed are the maximum that can be supported. Do not exceed the tested limits.

**Note:** The following limits do not apply to the FAS2220. See SAN configuration limits specific to the FAS2220 for limits applicable to that platform.

The following table lists the maximum supported value for each parameter based on testing. For best performance, do not configure your system with the maximum values. All values are for FC, FCoE, and iSCSI unless otherwise noted. If any node in the cluster is configured for FC, FCoE, or iSCSI, the cluster is limited to the SAN node limits.

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.

**Note:** The following system characteristics impact your cluster limits. Therefore, observe the per-node and per-cluster limits.

- The cluster limits are not simply the node limits multiplied by the number of nodes.
- Clusters with mixed controller types have cluster limits based on the lowest node limits. Be sure to observe both per-node and per-cluster limits.

Parameters	Max per node	Max per cluster
Nodes per cluster	n/a	<p><b>Note:</b> Maximum number of nodes within a cluster is determined by platform which supports the fewest number of nodes.</p> <p>4 for the following models:</p> <ul style="list-style-type: none"> <li>• FAS2240</li> <li>• 3160</li> <li>• 3170</li> <li>• 3210</li> <li>• 3240</li> </ul> <p>8 for all other models.</p>
Vservers	125	750

Parameters	Max per node	Max per cluster
Volumes	500 for the following models: <ul style="list-style-type: none"> <li>• FAS2240</li> <li>• 3140</li> <li>• 3160</li> <li>• 3210</li> <li>• 3220</li> <li>• 3240</li> </ul> 1,000 for all other models.	3,000 for the following models: <ul style="list-style-type: none"> <li>• FAS2240</li> <li>• 3140</li> <li>• 3160</li> <li>• 3210</li> <li>• 3220</li> <li>• 3240</li> </ul> 6,000 for all other models.
Volumes per node per Vserver	500	n/a
LUNs	2,048 for the following models: <ul style="list-style-type: none"> <li>• FAS2240</li> <li>• 3140</li> <li>• 3160</li> <li>• 3210</li> <li>• 3220</li> <li>• 3240</li> </ul> 8,192 for all other models.	8,192 for the following models: <ul style="list-style-type: none"> <li>• FAS2240</li> <li>• 3140</li> <li>• 3160</li> <li>• 3210</li> <li>• 3220</li> <li>• 3240</li> </ul> 49,152 for all other models.
LUN mappings	2,048 for the following models: <ul style="list-style-type: none"> <li>• FAS2240</li> <li>• 3140</li> <li>• 3160</li> <li>• 3210</li> <li>• 3220</li> <li>• 3240</li> </ul> 8,192 for all other models.	8,192 for the following models: <ul style="list-style-type: none"> <li>• FAS2240</li> <li>• 3140</li> <li>• 3160</li> <li>• 3210</li> <li>• 3220</li> <li>• 3240</li> </ul> 49,152 for all other models.
LUNs per volume	512	n/a
LUN size	16 TB	n/a
FC queue depth available per port	2,048	n/a
FC port fan-in	256	n/a
FCP (FC/FCoE) LIFs per port	16	n/a

Parameters	Max per node	Max per cluster
Connected hosts per node (FC) - assumes two initiators per host	1,024	n/a
iSCSI sessions per node	1,024	n/a
iSCSI LIFs per port	16	n/a
iSCSI LIFs per node	256	n/a
FCP LIFs per node	512	n/a
FCP/iSCSI LIFs per portset	32	n/a
igroups	512	3,072
Initiators per igroup	256	n/a
Portsets	512	3,072
FCP ports <b>Note:</b> FAS2240: 2 ports, single-node configurations only	32	256
Ethernet ports	Platform supported limits	Platform supported limits

### Related concepts

[Supported number of nodes for SAN configurations](#) on page 55

## SAN configuration limits specifically for the FAS2220

The FAS2220 has configuration limits for reliable operation. The values listed are the maximum that can be supported. For best performance, do not configure your system at the maximum values. Do not exceed the tested limits.

The following table lists the maximum supported value for each parameter based on testing for the FAS2220 platform. All values are for iSCSI. The FAS2220 does not support FC.

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.

**Note:** The following will impact your cluster limits. Therefore, please observe the per-node and per-cluster limits in the following table.

- The cluster limits are not simply the node limits multiplied by the number of nodes.

- Clusters with mixed controller types will have cluster limits based on the lowest node limits. Be sure to observe both per-node and per-cluster limits.

<b>FAS2220 Parameters</b>	<b>FAS2220 Max per node</b>	<b>FAS2220 Max per cluster</b>
Nodes per cluster	n/a	4
Vservers	8	16
Volumes	200	400
LUNs	200	400
LUN maps	200	400
igroups	32	64
Initiators per igroup	32	n/a
iSCSI sessions per node	64	n/a
iSCSI LIFs per port	8	n/a
iSCSI LIFs per portset	8	n/a
Portsets	32	64
LUN size	16 TB	n/a
LUNs per volume	200	n/a
Volumes per node per Vserver	200	n/a
Ethernet ports	Platform supported limit	Platform supported limit

### Related tasks

[Calculating queue depth](#) on page 63

## 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.

For clustered Data ONTAP, each FC logical interface (LIF) logs into an FC switch port. The total number of logins from a single initiator port equals the number of LIFs plus one login for the underlying physical port. You should not exceed the switch vendor's configuration limits for logins or other configuration values.

## 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

*Brocade Scalability Guidelines* - [www.brocade.com/downloads/documents/matrices/scalability-matrix-fos-v7.0a.pdf](http://www.brocade.com/downloads/documents/matrices/scalability-matrix-fos-v7.0a.pdf)

*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 hosts per controller 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.  $15 \times 128 = 1,920$ . 1,920 is less than the total queue depth limit of 2,048, so 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.  $30 \times 128 = 3,840$ . 3,840 is greater than the total queue depth limit of 2,048, so 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" host to 128 and the queue depth for "small" host 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 x (.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/rdsd/c2t2d0
```

```
scsictl -m queue_depth=16 /dev/rdsd/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.

## HBA Queue Depth

To update the queue depth for the Emulex HBA, execute the following commands:

### 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.

### About this task

## 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:
5. Reboot the server using the following command:

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

```
#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, run 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. Run 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 pull down 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\lpxnnds\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 drop down list box.
4. Update the Execution Throttle parameter.

## Copyright information

---

Copyright © 1994–2013 NetApp, Inc. All rights reserved. Printed in the U.S.

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

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

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

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

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

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

## Trademark information

---

NetApp, the NetApp logo, Network Appliance, the Network Appliance logo, Akorri, ApplianceWatch, ASUP, AutoSupport, BalancePoint, BalancePoint Predictor, Bypass, Campaign Express, ComplianceClock, Cryptainer, CryptoShred, CyberSnap, Data Center Fitness, Data ONTAP, DataFabric, DataFort, Decru, Decru DataFort, DenseStak, Engenio, Engenio logo, E-Stack, ExpressPod, FAServer, FastStak, FilerView, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexSuite, FlexVol, FPolicy, GetSuccessful, gFiler, Go further, faster, Imagine Virtually Anything, Lifetime Key Management, LockVault, Mars, Manage ONTAP, MetroCluster, MultiStore, NearStore, NetCache, NOW (NetApp on the Web), Onaro, OnCommand, ONTAPI, OpenKey, PerformanceStak, RAID-DP, ReplicatorX, SANscreen, SANshare, SANtricity, SecureAdmin, SecureShare, Select, Service Builder, Shadow Tape, Simplicity, Simulate ONTAP, SnapCopy, Snap Creator, SnapDirector, SnapDrive, SnapFilter, SnapIntegrator, SnapLock, SnapManager, SnapMigrator, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapSuite, SnapValidator, SnapVault, StorageGRID, StoreVault, the StoreVault logo, SyncMirror, Tech OnTap, The evolution of storage, Topio, VelocityStak, vFiler, VFM, Virtual File Manager, VPolicy, WAFL, Web Filer, and XBB are trademarks or registered trademarks of NetApp, Inc. in the United States, other countries, or both.

IBM, the IBM logo, and [ibm.com](http://ibm.com) are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. A complete and current list of other IBM trademarks is available on the web at [www.ibm.com/legal/copytrade.shtml](http://www.ibm.com/legal/copytrade.shtml).

Apple is a registered trademark and QuickTime is a trademark of Apple, Inc. in the United States and/or other countries. Microsoft is a registered trademark and Windows Media is a trademark of Microsoft Corporation in the United States and/or other countries. RealAudio, RealNetworks, RealPlayer, RealSystem, RealText, and RealVideo are registered trademarks and RealMedia, RealProxy, and SureStream are trademarks of RealNetworks, Inc. in the United States and/or other countries.

All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

NetApp, Inc. is a licensee of the CompactFlash and CF Logo trademarks.

NetApp, Inc. NetCache is certified RealSystem compatible.

## How to send your comments

---

You can help us to improve the quality of our documentation by sending us your feedback.

Your feedback is important in helping us to provide the most accurate and high-quality information. If you have suggestions for improving this document, send us your comments by email to [doccomments@netapp.com](mailto:doccomments@netapp.com). To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

You can also contact us in the following ways:

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

# Index

10-Gb [39](#)

22xx

- multifabric single-node FC configurations [38](#)
- single-fabric single-node FC configurations [37](#)
- single-node FC configurations [36](#)

31xx

- configuration limits [59](#)
- FC configurations [32](#)
- multifabric HA pair FC configurations [35](#)
- single-fabric HA pair FC configurations [34](#)
- single-node FC configurations [32](#)
- target port configuration [32](#)

32xx

- configuration limits [59](#)
- FC configurations [27](#)
- multifabric HA configuration [30](#)
- single-fabric HA configuration [29](#)
- single-node FC configurations [28](#)
- target port configuration [27](#)

4-Gb FC port

- supported speed [15](#)

6040 or 6080

- FC configurations [22](#)
- single-fabric HA pair FC configurations [24](#)
- target port configuration [22](#)

6040 or 6080 multifabric HA pair FC configurations [26](#)

60xx

- configuration limits [59](#)
- single-node FC configurations [22](#)

62xx

- configuration limits [59](#)
- FC configurations [16](#)
- multifabric HA pair FC configurations [20](#)
- single-fabric HA pair FC configurations [19](#)
- single-node FC configurations [17](#)
- target port configuration [17](#)

8-Gb FC port

- supported speed [15](#)

## A

AIX hosts

- setting queue depth for [65](#)

ALUA configurations [54](#)

asymmetric logical unit access (ALUA) configurations [54](#)

## B

Brocade switch

- configuration limits [62](#)

## C

Cisco switches

- configuration limits [62](#)
- FC and FCoE zoning requirement [52](#)

configuration limits

- by host operating system [57](#)
- FAS2220 storage systems [61](#)
- FC switch [62](#)
- parameters defined [55](#)
- volume and Vserver requirements [55](#)

configurations

- 31xx FC configurations [32](#)
- 32xx FC configurations [27](#)
- 6040 or 6080 FC configurations [22](#)
- 62xx FC configurations [16](#)
- FC [13](#)
- FCoE [41](#)
- FCoE initiator to FC target [41](#)
- FCoE initiator to FCoE and FC mixed target [44](#)
- FCoE initiator to FCoE target [43](#)
- FCoE initiator to FCoE target mixed with IP traffic [46](#)
- iSCSI [6](#)

configurations, 22xx

- multifabric single-node FC configurations [38](#)
- single-fabric single-node FC configurations [37](#)
- single-node FC configurations [36](#)

configurations, 31xx

- multifabric HA pair FC configuration [35](#)
- single-fabric HA pair FC configuration [34](#)
- single-node FC configurations [32](#)

configurations, 32xx

- multifabric HA configuration [30](#)
- single-fabric HA configuration [29](#)
- single-node FC configurations [28](#)

configurations, 6040 or 6080

- single-fabric HA pair FC configuration [24](#)

configurations, 6040 or 6080 multifabric HA pair FC

configuration [26](#)

configurations, 60xx

- single-node FC configurations [22](#)
- configurations, 62xx
  - multifabric HA pair FC configuration [20](#)
  - single-fabric HA pair FC configuration [19](#)
  - single-node FC configurations [17](#)

## D

- DCB (data center bridging) switch
  - for FCoE [39](#)
- direct-attached configuration
  - iSCSI [11](#)
- dynamic VLANs [12](#)

## E

- EMC CLARiON
  - shared configurations [53](#)
- ESX
  - host configuration limits [57](#)
- expansion FC adapter
  - supported port speed [15](#)
- expansion FC ports
  - usage rules [14](#)

## F

- FAS2220
  - configuration limits [61](#)
- FAS2240
  - configuration limits [59](#)
- FC
  - 31xx configurations [32](#)
  - 31xx target port configuration [32](#)
  - 32xx configurations [27](#)
  - 32xx target port configuration [27](#)
  - 6040 or 6080 configurations [22](#)
  - 6040 or 6080 target port configuration [22](#)
  - 62xx configurations [16](#)
  - 62xx target port configuration [17](#)
  - configurations overview [13](#)
  - multifabric switch zoning [50](#)
  - onboard and expansion port usage rules [14](#)
  - port speed [15](#)
  - single-fabric switch zoning [49](#)
  - supported port speed [15](#)
  - switch configuration [16](#)
  - switch hop count [15](#)
  - switch WWN zoning [48](#)
  - switch zoning [48](#)

- switch zoning with individual zones [49](#)
- FC LIF
  - zoning restrictions for Cisco switches [52](#)
- FC switch
  - configuration limits [62](#)
- FCoE
  - initiator and target combinations [40](#)
  - supported configurations [41](#)
  - switch hop count [40](#)
  - switch zoning [48](#)
- FCoE configurations
  - FCoE initiator to FC target [41](#)
  - FCoE initiator to FCoE and FC mixed target [44](#)
  - FCoE initiator to FCoE target [43](#)
  - FCoE initiator to FCoE target mixed with IP traffic [46](#)
- FCoE LIF
  - zoning restrictions for Cisco switches [52](#)
- Fibre Channel over Ethernet (FCoE)
  - overview [39](#)
- FlexVol volumes
  - required for SAN [55](#)

## H

- HA pair
  - iSCSI direct-attached configuration [11](#)
  - iSCSI multi-network configuration [6, 10](#)
  - iSCSI single-network configuration [6, 9](#)
- HBA [39](#)
- heterogeneous SAN
  - using VSAN [13](#)
- hop count
  - for FC switches [15](#)
  - for FCoE switches [40](#)
- host multipathing software
  - when required [16](#)
- host utilities
  - AIX queue depths [65](#)
  - HP-UX queue depths [65](#)
  - Solaris queue depths [66](#)
  - VMWare queue depths [67](#)
  - Windows queue depths [68](#)
- HP-UX
  - host configuration limits [57](#)
- HP-UX hosts
  - setting queue depth for [65](#)

## I

- Infinite Volumes
  - not supported for SAN [55](#)
- initiator FC ports
  - onboard and expansion usage rules [14](#)
- initiators
  - FCoE and FC combinations [40](#)
- IP traffic
  - in FCoE configurations [46](#)
- iSCSI
  - configurations [6](#)
  - direct-attached configuration [11](#)
  - dynamic VLANs [12](#)
  - multi-network configuration [6, 10](#)
  - single-network configuration [6, 9](#)
  - static VLANs [12](#)
  - using VLANs [12](#)

## L

- LIFs
  - FC and FCoE zoning restrictions for Cisco switches [52](#)
- Linux
  - host configuration limits [57](#)

## M

- MPIO software
  - when required [16](#)
- multifabric HA pair FC configurations
  - 31xx [35](#)
  - 32xx [30](#)
  - 6040 or 6080 [26](#)
  - 62xx [20](#)
- multifabric single-node FC configurations
  - 22xx [38](#)
  - 31xx [32](#)
  - 32xx [28](#)
  - 60xx [22](#)
  - 62xx [17](#)
- multipathing software
  - when required [16](#)

## O

- onboard FC port
  - supported port speed [15](#)
- onboard FC ports

usage rules [14](#)

## P

- parameters
  - configuration limit definitions [55](#)
- point-to-point
  - FC switch port topology [16](#)
- port speed
  - supported for FC [15](#)
- port topology
  - FC switch [16](#)
- PowerPath
  - with shared configurations [53](#)

## Q

- queue depth
  - calculating [63](#)
  - tuning [63](#)
- queue depths
  - setting for AIX hosts [65](#)
  - setting for HP-UX hosts [65](#)
  - setting for Solaris hosts [66](#)
  - setting for VMware hosts [67](#)
  - setting for Windows hosts [68](#)

## R

recommended number of paths [6, 13, 39](#)

## S

- SAN configuration limits
  - by host operating system [57](#)
  - FAS2220 storage systems [61](#)
  - FC switch [62](#)
  - parameters defined [55](#)
  - volume and Vserver requirements [55](#)
- shared SAN configurations [53](#)
- single points of failure
  - how to avoid [6, 13, 39](#)
- single-fabric HA configurations
  - 32xx [29](#)
- single-fabric HA pair FC configurations
  - 31xx [34](#)
  - 6040 or 6080 [24](#)
  - 62xx [19](#)
- single-fabric single-node FC configurations

- 22xx [37](#)
- 31xx [32](#)
- 32xx [28](#)
- 60xx [22](#)
- 62xx [17](#)
- soft zoning
  - FC switch [48](#)
- Solaris
  - host configuration limits [57](#)
- Solaris hosts
  - setting queue depth for [66](#)
- static VLANs [12](#)
- supported configurations
  - FCoE [41](#)
- switch
  - configuration limits [62](#)
  - FC and FCoE zoning requirement [52](#)
  - FC configuration [16](#)
  - FC hop count [15](#)
  - FC multifabric zoning [50](#)
  - FC single-fabric zoning [49](#)
  - FC WWN zoning [48](#)
  - FC zoning [48](#)
  - FC zoning with individual zones [49](#)
  - FCoE hop count [40](#)
  - FCoE zoning [48](#)

## T

- target FC ports
  - onboard and expansion usage rules [14](#)
- target port configurations
  - 31xx [32](#)
  - 32xx [27](#)
  - 6040 or 6080 [22](#)
  - 62xx [17](#)
- targets
  - FCoE and FC combinations [40](#)

## V

- virtual LANs
  - reasons for using [12](#)
- VLANs
  - dynamic [12](#)
  - reasons for using [12](#)
  - static [12](#)
- VMware hosts
  - setting queue depth for [67](#)
- volumes
  - restrictions for SAN [55](#)
- VSAN
  - for heterogeneous SAN [13](#)
- Vservers
  - restrictions for SAN [55](#)

## W

- Windows
  - host configuration limits [57](#)
- Windows hosts
  - setting queue depth for [68](#)
- WWN zoning
  - FC switch [48](#)

## Z

- zoning
  - FC switch [48](#)
  - FC switch by WWN [48](#)
  - FC switch multifabric [50](#)
  - FC switch single-fabric [49](#)
  - FC switch with individual zones [49](#)
  - FCoE switch [48](#)
  - requirements for Cisco switches [52](#)