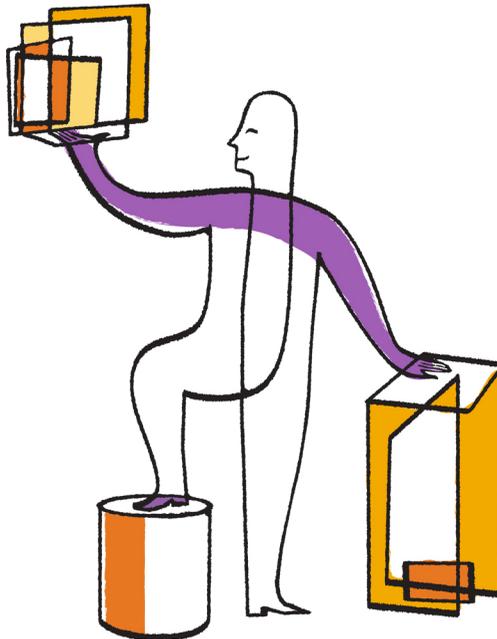




NetApp®

Clustered Data ONTAP® 8.3

SAN Configuration Guide



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

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

Part number: 215-09166_C0
March 2015

Contents

iSCSI configurations	5
Recommended number of paths to avoid single points of failure	5
Single-network and multi-network single-node configuration in an iSCSI SAN	6
Single-network dual-node HA pair in an iSCSI SAN	7
Multi-network dual-node HA pair in an iSCSI SAN	8
Direct-attached single-node configurations in an iSCSI SAN	9
VLANs for iSCSI configurations	10
Static VLANs	10
Dynamic VLANs	10
Supported FC configurations	11
Recommended number of paths to avoid single points of failure	11
When host multipathing software is required	12
FC supported hop counts	12
Supported FC ports	13
How to prevent loss of connectivity when using the X1133A-R6 adapter ...	14
Unified Target Adapter 2 (UTA2) in NetApp collateral is referred to as Converged Network Adapter (CNA) in product documentation	14
Function of the X1143A-R6 adapter	14
Default configuration for the X1143A-R6	15
FC supported speeds	15
FC Target port configuration recommendations	16
FC switch configuration best practices	17
Single-node configurations	18
Single-fabric configurations	20
FAS2240 single-fabric single-node configurations	20
Single-fabric HA configurations	21
Multifabric configurations	23
FAS2240 multifabric single-node configurations	23
Multifabric HA configurations	24
Fibre Channel over Ethernet overview	26
Recommended number of paths to avoid single points of failure	26
FCoE initiator and target combinations	27

FCoE supported hop count	27
FCoE supported configurations	28
FCoE initiator to FC target configuration	28
FCoE end-to-end configuration	30
FCoE mixed with FC	31
FCoE mixed with IP storage protocols	33
Fibre Channel and FCoE zoning	35
World Wide Name-based zoning	36
Individual zones	36
Single-fabric zoning	36
Dual-fabric HA pair zoning	38
FC and FCoE LIFs on the same port need to be in separate zones	39
Support for shared SAN configurations	40
ALUA Configurations	41
Configuration limits for FC, FCoE, and iSCSI configurations	42
SAN configurations always use FlexVol volumes	42
SAN configuration limit parameters and definitions	42
Determining the number of supported nodes for SAN configurations	44
Determining the number of supported hosts per cluster in FC configurations	45
Determining the supported number of hosts per cluster in iSCSI configurations	46
Host operating system limits for SAN configurations	46
SAN configuration limits	48
SAN configuration limits for Data ONTAP-v platforms	56
FC switch configuration limits	57
Calculating queue depth	57
Setting queue depths on AIX hosts	60
Setting queue depths on HP-UX hosts	60
Setting queue depths on Solaris hosts	61
Setting queue depths on VMware hosts	61
Setting queue depths on Windows hosts	63
Copyright information	64
Trademark information	65
How to send comments about documentation and receive update notification	66
Index	67

iSCSI configurations

Clustered Data ONTAP supports iSCSI configurations for both single-node and multi-node configurations. Multi-node configurations are built with HA pairs up to the supported number of nodes for SAN configurations.

You need to create one or more iSCSI paths from each storage controller, using logical interfaces (LIFs) that can access a given LUN. LIFs do not migrate or assume the IP addresses of the failed partner node. Instead, the MPIO software, using ALUA on the host, is responsible for selecting the appropriate paths for LUN access through LIFs. This behavior is just like the FC path access method.

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. 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* and TR-3441 *Windows Multipathing Options with Data ONTAP: Fibre Channel and iSCSI*.

Related tasks

[Determining the number of supported nodes for SAN configurations](#) on page 44

[Determining the supported number of hosts per cluster in iSCSI configurations](#) on page 46

Related references

[FCoE supported hop count](#) on page 27

Related information

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

[NetApp Technical Report 3441: Windows Multipathing Options with Data ONTAP: Fibre Channel and iSCSI](#)

Recommended number of paths to avoid single points of failure

You should have a minimum of two paths per LUN connecting to each node being used by the Storage Virtual Machine (SVM) in your cluster. This eliminates single points of failure and enables the system to survive component failures.

If you have four or more nodes in your cluster or more than four target ports being used by the SVMs in any of your nodes, you can use the following methods to limit the number of paths that can be used to access LUNs on your controllers:

- Selective LUN Map (SLM)
SLM is enabled by default for all LUN maps created in Data ONTAP 8.3.
- Portsets for iSCSI
- FC igroup mappings from your host
- FC switch zoning

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 multiple IP networks. The network can consist of one or more switches. Each node can have multiple iSCSI connections to the network. The number of available iSCSI target ports depends on the model of the storage controller, the number of onboard Ethernet ports on the controller, and the number of ports available through installed Ethernet network cards on the controller.

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

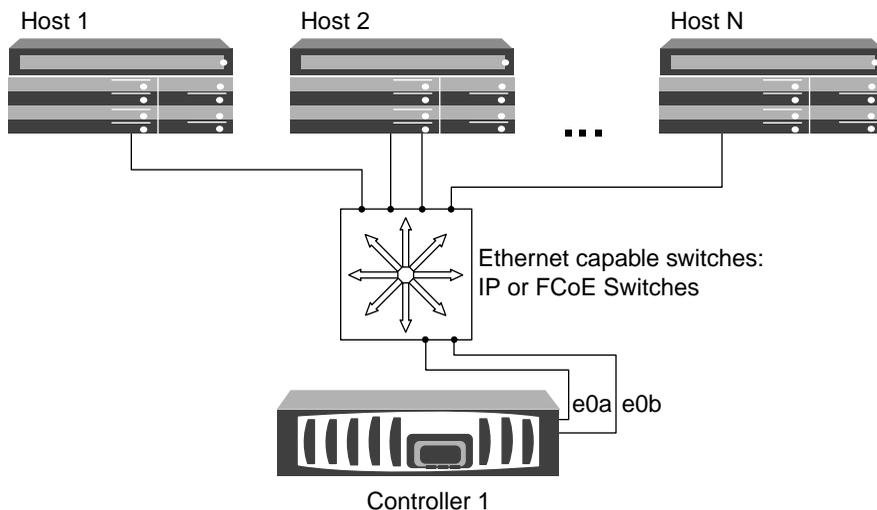


Figure 1: iSCSI single-network single-node configuration

Attribute	Value
Fully redundant	No, due to the single network
Type of network	Single network

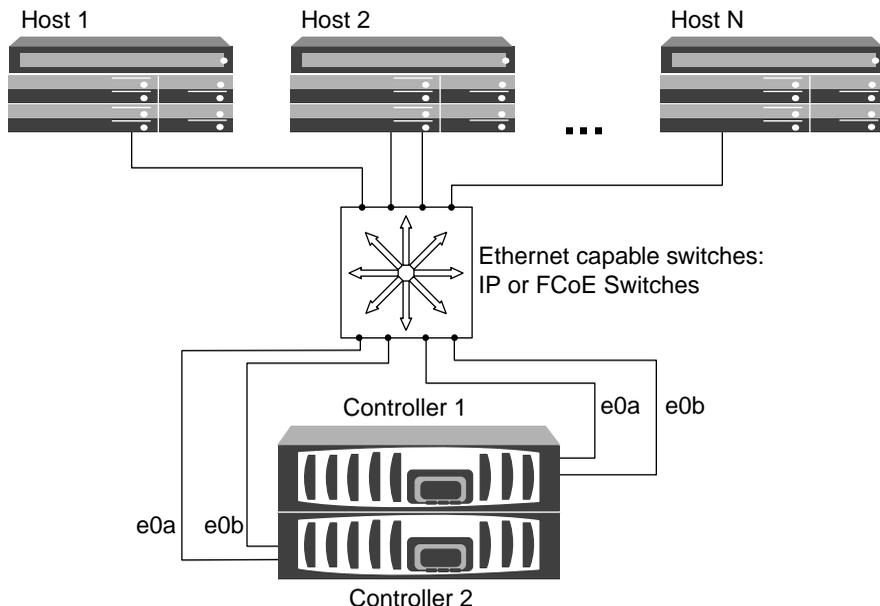
Attribute	Value
Different host operating systems	Yes, with multiple-host configurations
Multipathing required	Yes, if multiple connections per host are configured
Type of configuration	Single node

Single-network dual-node HA pair in an iSCSI SAN

You can connect hosts to dual-node HA pairs that use iSCSI 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 IP ports available depends on the model of the storage controller and any additional IP-capable PCIe cards installed in the controller, including network cards or UTAs (CNAs installed in the controllers).

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

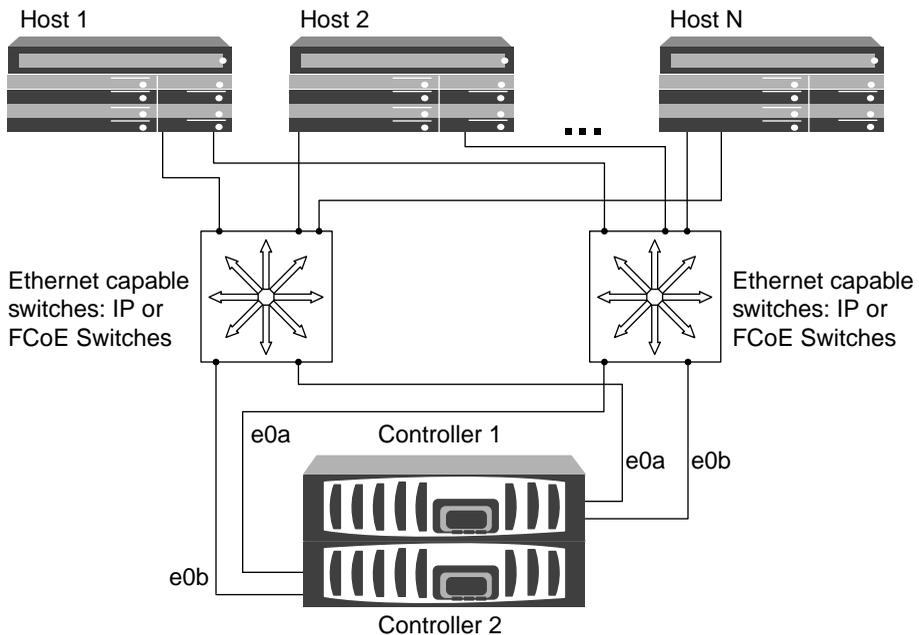


Attribute	Value
Fully redundant	No, due to the single network

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

Multi-network dual-node HA pair in an iSCSI SAN

You can connect hosts to dual-node HA pair controllers that use iSCSI 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 requires multipathing software to be installed and configured.



Attribute	Value
Fully redundant	Yes
Type of network	Multi-network

Attribute	Value
Different host operating systems	Yes, with multiple-host configurations
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.

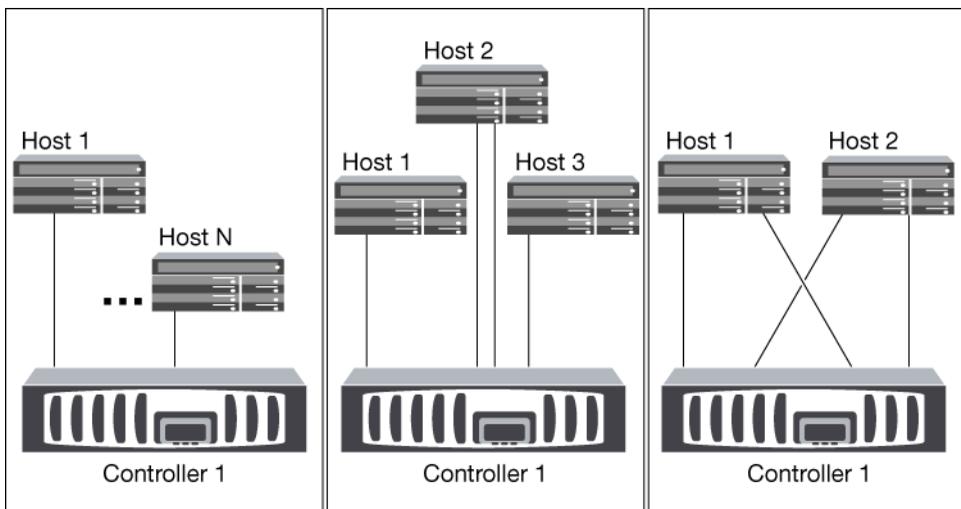


Figure 2: iSCSI direct-attached single-node configurations

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.

Supported FC configurations

Fibre Channel (FC) configurations are supported by various host OS types. Supported FC configurations include single-fabric and multifabric configurations. Both single-controller and HA pairs are supported. Direct-attached configurations are not supported.

For multiple-host configurations, hosts with different operating systems, such as Windows, Linux, or UNIX, can access the storage solution at the same time. Hosts require that a supported multipathing solution be installed and configured. Supported operating systems and multipathing solutions can be verified on the Interoperability Matrix.

HA pairs with multiple, physically independent storage fabrics (minimum of two) are recommended for SAN solutions. This provides redundancy at the fabric and storage system layers. Redundancy is particularly important because these layers typically support many hosts.

The use of heterogeneous FC switch fabrics is not supported, except in the case of embedded blade switches. Specific exceptions are listed on the Interoperability Matrix.

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

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

Related information

[*NetApp Interoperability*](#)

Recommended number of paths to avoid single points of failure

You should have a minimum of two paths per LUN connecting to each node being used by the Storage Virtual Machine (SVM) in your cluster. This eliminates single points of failure and enables the system to survive component failures.

If you have four or more nodes in your cluster or more than four target ports being used by the SVMs in any of your nodes, you can use the following methods to limit the number of paths that can be used to access LUNs on your controllers:

- Selective LUN Map (SLM)
SLM is enabled by default for all LUN maps created in Data ONTAP 8.3.
- Portsets for iSCSI
- FC igroup mappings from your host

- FC switch zoning

When host multipathing software is required

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

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

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

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

In single-fabric single-node configurations, multipathing software is not required if you only have a single path from the host to the node. You can use zoning or port sets to limit paths.

In HA configurations, multipathing software is required unless you use zoning or port sets to limit the host to a single path.

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

Related information

[NetApp Interoperability](#)

FC supported hop counts

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

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

The following table lists supported hop counts:

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

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

Supported FC ports

The number of onboard FC ports varies based on the model of the controller. FC ports are also available through supported target expansion adapters.

Onboard FC ports:

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

Target expansion adapter FC ports:

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

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

Related information

[NetApp Hardware Universe](#)

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

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

To prevent loss of connectivity in the event of port failure, it is recommended that you configure your system with redundant paths to separate X1133A-R6 HBAs, or with redundant paths to ports supported by different ASICs on the HBA.

Unified Target Adapter 2 (UTA2) in NetApp collateral is referred to as Converged Network Adapter (CNA) in product documentation

NetApp uses the marketing term “UTA2” to refer to a new generation of network adapters that can be configured as either 10-Gb Ethernet and FCoE ports or as 16-Gb FC initiator or target ports. The Data ONTAP CLI and NetApp hardware and software documentation refer to these as “CNA ports”.

These ports are available as the X1143A-R6 adapter and as onboard ports on the FAS80xx systems.

Function of the X1143A-R6 adapter

The X1143A-R6 network adapter is a new generation adapter with ports that can be configured as either 10 Gb Ethernet and FCoE ports or as 16Gb FC initiator or target ports.

Data ONTAP has supported FC and converged network adapter (CNA) modes in previous releases, but used different add-on host bus adapters (HBAs) and onboard devices for each mode. The X1143A-R6 adapter supports both FC and CNA modes on the same device using dual 16-Gb FC and 10-GbE/FCoE ports. The X1143A-R6 is the first device that can be configured into either one of the personality modes within the same hardware.

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

Data ONTAP discovers different number of PCI functions per chip depending on its mode. Each PCI function is an independent PCI device with its own PCI configuration space. For example, under FC mode, only one PCI function of FC will be available for each port, and the physical port allows the use of either 16G FC SFP+, 8G FC SFP+, or 4G FC SFP+. While under CNA mode, two PCI functions will be available for each port. One is the NIC function while the other is the FCoE function and the physical port only allows the use of 10G FC SFP+. Data ONTAP drivers identify and bind to PCI functions based on PCI Vendor_ID and PCI Device_ID combinations from their PCI configuration space. The X1143A-R6 device will assign unique but different PCI Device_IDs for different types of PCI functions, such as FC, NIC and FCoE.

In addition, from a Data ONTAP device driver perspective, once a PCI function is discovered, a device driver will be attached to use the X1143A-R6. For the NIC function, the NIC driver will be attached. In the case of the FC or FCoE function, a different driver will be used depending on its type. The X1143A-R6 supports the following types in Data ONTAP:

- FCP Initiator
- FCP target

Note: CNA mode only supports FCP target for the FCoE function.

Data ONTAP can attach the proper driver to the X1143A-R6 device during boot-up and initialize the device for FCP target. The Unified Connect Management (UCM) framework provides the infrastructure to manage these personality modes and FC-4 types, enabling you to change configurations. However UCM itself does not require any one particular configuration default for personality mode and/or FC-4 type. Instead it is capable of working with any default value that is set.

Default configuration for the X1143A-R6

You must connect to the X1143A-R6 ports using LC fiber optic cables with supported SFP+ optical modules, or by using copper Twinax cables (in 10GbE mode only). You can monitor the activity on the ports through the LEDs on the bezel.

The X1143A-R6 device can be configured as Fibre Channel (FC) or Converged Network Adapter (CNA). By default, the X1143A-R6 ships configured in FC target mode. You must use the `system hardware unified-connect modify` command to change the operational mode (FC or Ethernet) of the card. See the man page for details.

Note: Both ports that are connected to the same chip must operate in the same mode (FC or CNA).

Data ONTAP provides a mechanism to manage these configurations, which enables you to reconfigure the device into different modes. For the FC type, there is an existing command in Data ONTAP, `system hardware unified-connect modify`, which can be used to manage FC types of ports between FCP initiator and FCP target.

FC supported speeds

FC target ports can be configured to run at different speeds. You should set the target port speed to match the speed of the device to which it connects. All target ports used by a given host should be set to the same speed.

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

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

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

FC Target port configuration recommendations

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

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

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

Controller	Port pairs with shared ASIC	Number of target ports: Preferred ports
FAS8080, FAS8060 and FAS8040	0e+0f 0g+0h	1: 0e 2: 0e, 0g 3: 0e, 0g, 0h 4: 0e, 0g, 0f, 0h
FAS8020	0c+0d	1: 0c 2: 0c, 0d
62xx	0a+0b 0c+0d	1: 0a 2: 0a, 0c 3: 0a, 0c, 0b 4: 0a, 0c, 0b, 0d

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

Related information

[NetApp Hardware Universe](#)

FC switch configuration best practices

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

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

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

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

Related information

[NetApp Interoperability](#)

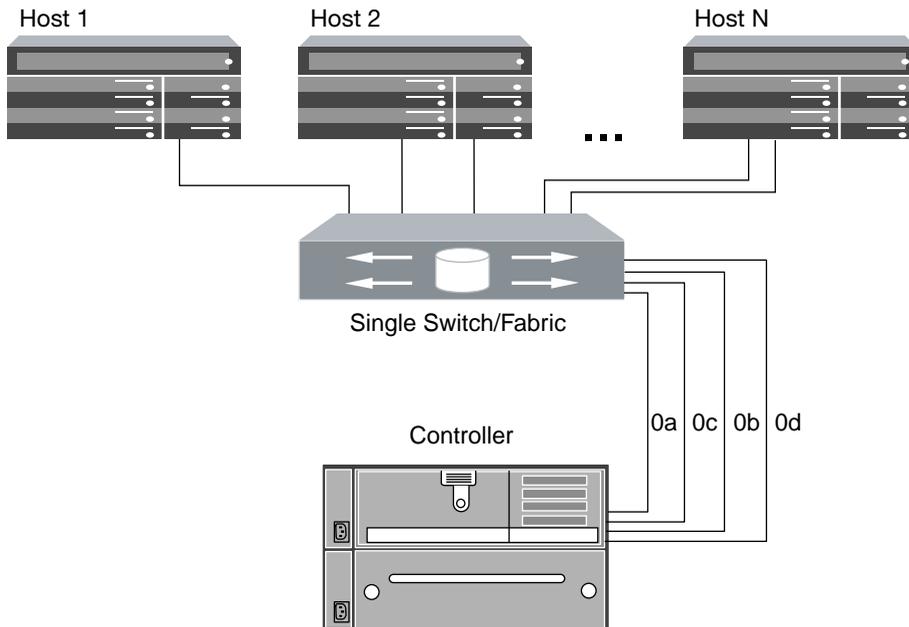
[NetApp Technical Report 4080: Best Practices for Scalable SAN in Clustered Data ONTAP 8.2](#)

Single-node configurations

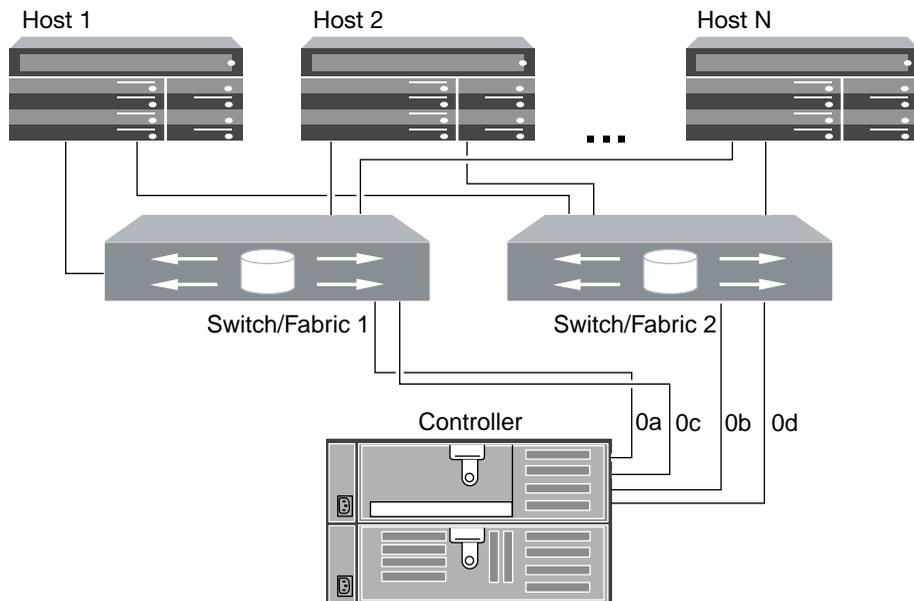
You can connect hosts to a single node using a single FC fabric or using multiple FC fabrics. All hardware platforms that support FC also support single-node configurations. Single-node configurations are not fully redundant.

In single-node configurations, you can use the maximum number of supported onboard FC ports per controller and the maximum number of supported target expansion adapters. For more information about the number of onboard ports and target expansion adapters supported by your specific platform, see the *Hardware Universe*.

The following illustration shows a single-fabric single-node configuration. The FC target port numbers (0a, 0c, 0b, 0d) are examples. The actual port numbers vary depending on the model of your storage node and the type of port you are using. If you are using FC target expansion adapters (FC HBAs) or the CNA expansion adapters, the target port numbers depend on the expansion slots into which the expansion adapters are installed.



The following illustration shows a multifabric single-node configuration. The FC target port numbers (0a, 0c, 0b, 0d) are examples. The actual port numbers vary depending on the model of your storage node and the type of port you are using. If you are using FC target expansion adapters (FC HBAs) or the CNA expansion adapters, the target port numbers depend on the expansion slots into which the expansion adapters are installed.



Related concepts

[When host multipathing software is required](#) on page 12

Related references

[Single-fabric zoning](#) on page 36

[Dual-fabric HA pair zoning](#) on page 38

Related information

[Hardware Universe: hwi.netapp.com](http://hwi.netapp.com)

Single-fabric configurations

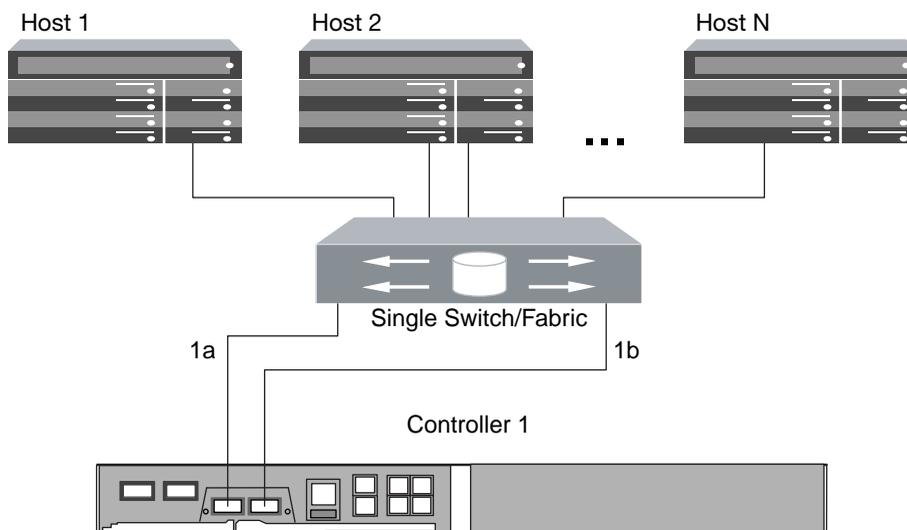
You can connect hosts to a single node or to HA pairs using a single FC fabric. Direct-attached FC configurations are not supported with multiple node clusters.

FAS2240 single-fabric single-node configurations

The FAS2240 only supports single-fabric single-node FC configurations through the X1150A-R6 expansion adapter. The FAS2240 does not support single-fabric HA configurations. Connecting through a single switch means that this configuration is not fully redundant. Direct-attached FC configurations are not supported with multiple node clusters.

If you use multiple paths to a LUN, multipathing software is required on the host. You can use FC switch zoning, portsets, and Selective LUN Map (SLM) to limit the number of paths between hosts and LUNs in configurations with multiple target ports connected to the same fabric. You can use SLM to specify which nodes will advertise paths to a given LUN. You can use portsets to help limit the paths that can be used to connect to a given LUN. Zoning can be used to limit connectivity between various ports, thus limiting effective paths to only those sharing a common zone.

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



Single-fabric HA configurations

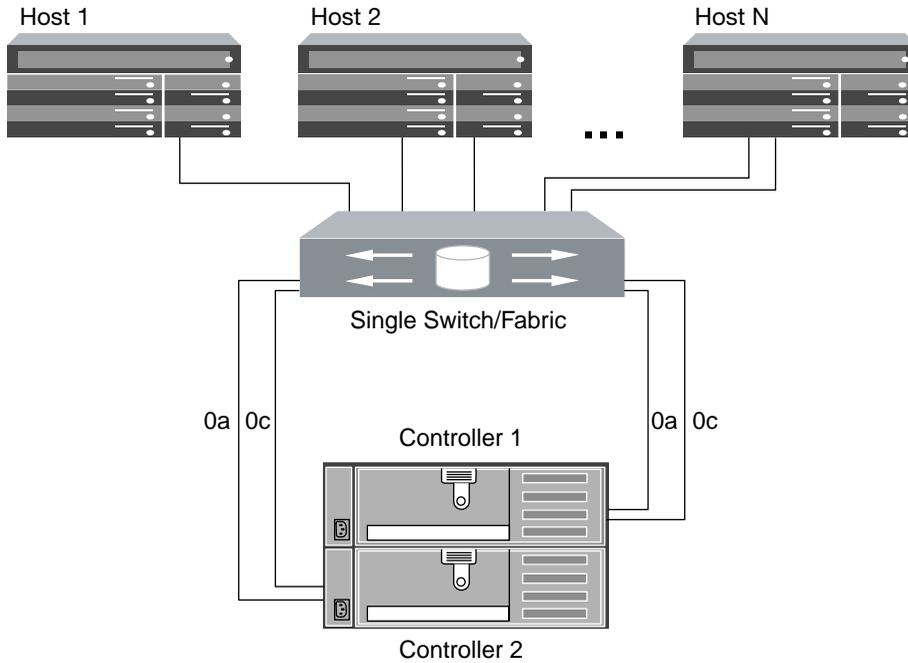
You can connect hosts to both controllers in an HA configuration through a single FC switch. Because the hosts and controllers are connected through a single switch, single-fabric HA configurations are not fully redundant. Direct-attached FC configurations are not supported with multiple node clusters.

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

In single-fabric HA configurations, you can use the maximum number of supported onboard FC ports per controller and the maximum number of supported target expansion adapters. For more information about the number of onboard ports and target expansion adapters supported by your specific platform, see the *Hardware Universe*.

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. You can use portsets to help limit the paths that can be used to access LUNs on the controller.

The following illustration shows a single-fabric HA configuration. The FC target port numbers (0a, 0c) are examples. The actual port numbers vary depending on the model of your storage node and the type of port you are using. If you are using FC target expansion adapters (FC HBAs) or the CNA expansion adapters, the target port numbers depend on the expansion slots into which the expansion adapters are installed.



Related concepts

[FC Target port configuration recommendations](#) on page 16

[When host multipathing software is required](#) on page 12

Related references

[Single-fabric zoning](#) on page 36

Related information

[NetApp Hardware Universe](#)

Multifabric configurations

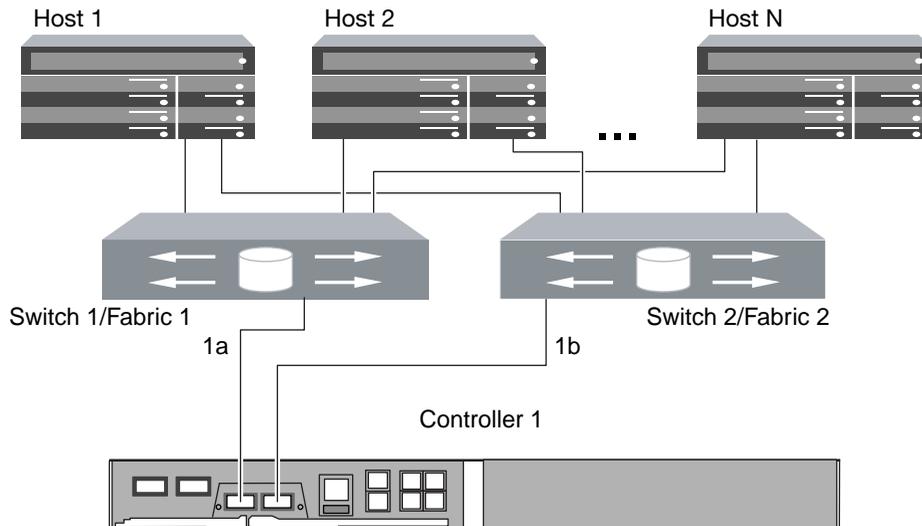
You can connect hosts to a single node or to HA pairs through two or more FC fabrics for redundancy. The FAS2240 can only support single-node configurations. All other platforms support HA configurations. Direct attached FC configurations are not supported.

FAS2240 multifabric single-node configurations

You can connect hosts to a single FAS2240 node using two or more FC fabrics. However, using a single node means that this configuration is not fully redundant. The FAS2240 does not support multifabric HA configurations or direct-attached FC configurations.

If you use multiple paths to a LUN, multipathing software is required on the host. Using FC switch zoning or portsets with Selective LUN Map (SLM) in the Storage Virtual Machine (SVM) is recommended, to limit the number of paths between hosts initiators and LUNs in configurations with multiple target LIFs in the SVM connected to the same fabric.

For simplicity, the following illustration shows a FAS2240 multifabric single-node configuration with only two fabrics. However, you can have two or more fabrics in any multifabric configuration. In this illustration the storage controllers are 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.



Related concepts

When host multipathing software is required on page 12

Related references

Dual-fabric HA pair zoning on page 38

Multifabric HA configurations

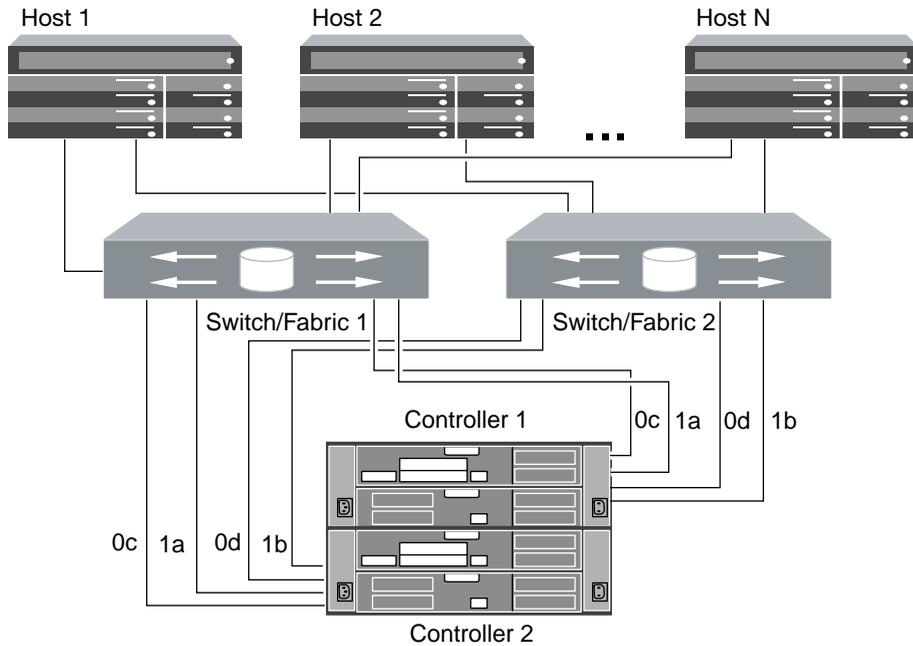
You can connect hosts to both controllers in an HA configuration through two or more FC switch fabrics. Multifabric HA configurations are fully redundant and support different host operating systems with multiple-host configurations.

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

In multifabric HA configurations, you can use the maximum number of supported onboard FC ports per controller and the maximum number of supported target expansion adapters. For more information about the number of onboard ports and target expansion adapters supported by your specific platform, see the *Hardware Universe*.

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. You can use portsets to help limit the paths that can be used to access LUNs on the controller

For simplicity, the following multifabric HA illustration shows only two fabrics. However, you can have two or more fabrics in any multifabric configuration. The FC target port numbers (0c, 1a, 0d, 1bd) are examples. The actual port numbers vary depending on the model of your storage node and the type of port you are using. If you are using FC target expansion adapters (FC HBAs) or the CNA expansion adapters, the target port numbers depend on the expansion slots into which the expansion adapters are installed.



Related concepts

[FC Target port configuration recommendations](#) on page 16

[When host multipathing software is required](#) on page 12

Related references

[Dual-fabric HA pair zoning](#) on page 38

Related information

[NetApp Hardware Universe](#)

Fibre Channel over Ethernet overview

Fibre Channel over Ethernet (FCoE) encapsulates a Fibre Channel frame inside an Ethernet frame and uses a lossless 10-Gb Ethernet network instead of an FC network. FCoE maintains existing FC management and controls.

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 FCoE traffic that 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 (UTA) or converged network adapter (CNA). Generally, a UTA defines a FCoE target port and a CNA is used to specify a host-based FCoE initiator. Like the CNA, the UTA supports both FCoE and regular Ethernet traffic and can be used for non-FCoE IP traffic such as NFS, CIFS, or iSCSI.

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

Related information

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

Recommended number of paths to avoid single points of failure

You should have a minimum of two paths per LUN connecting to each node being used by the Storage Virtual Machine (SVM) in your cluster. This eliminates single points of failure and enables the system to survive component failures.

If you have four or more nodes in your cluster or more than four target ports being used by the SVMs in any of your nodes, you can use the following methods to limit the number of paths that can be used to access LUNs on your controllers:

- Selective LUN Map (SLM)
SLM is enabled by default for all LUN maps created in Data ONTAP 8.3.
- Portsets for iSCSI

- FC igroup mappings from your host
- FC switch zoning

FCoE initiator and target combinations

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

FCoE initiators

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

The following table lists the supported combinations:

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

FCoE targets

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

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

FCoE supported hop count

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

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

For FCoE, you can have FCoE switches connected to FC switches.

For end-to-end FCoE connections, the FCoE switches must be running a firmware version that supports Ethernet inter-switch links (ISLs).

The following table lists the maximum supported hop counts:

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

FCoE supported configurations

All storage systems are supported with native Fibre Channel over Ethernet (FCoE) target expansion adapters (called unified target adapters or UTAs), in single-fabric and multi-fabric configurations, except the FAS22xx and FAS2520. The FAS22xx and FAS2520 does not support FCoE.

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 for FC or iSCSI.

Note: Although iSCSI configurations are supported on standard Ethernet switches and have 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 FC switches. Supported configurations are listed in the Interoperability Matrix. Some of the parameters included in these supported configurations are the switch model, the number of switches that can be deployed in a single fabric, and the supported switch firmware version.

Related information

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

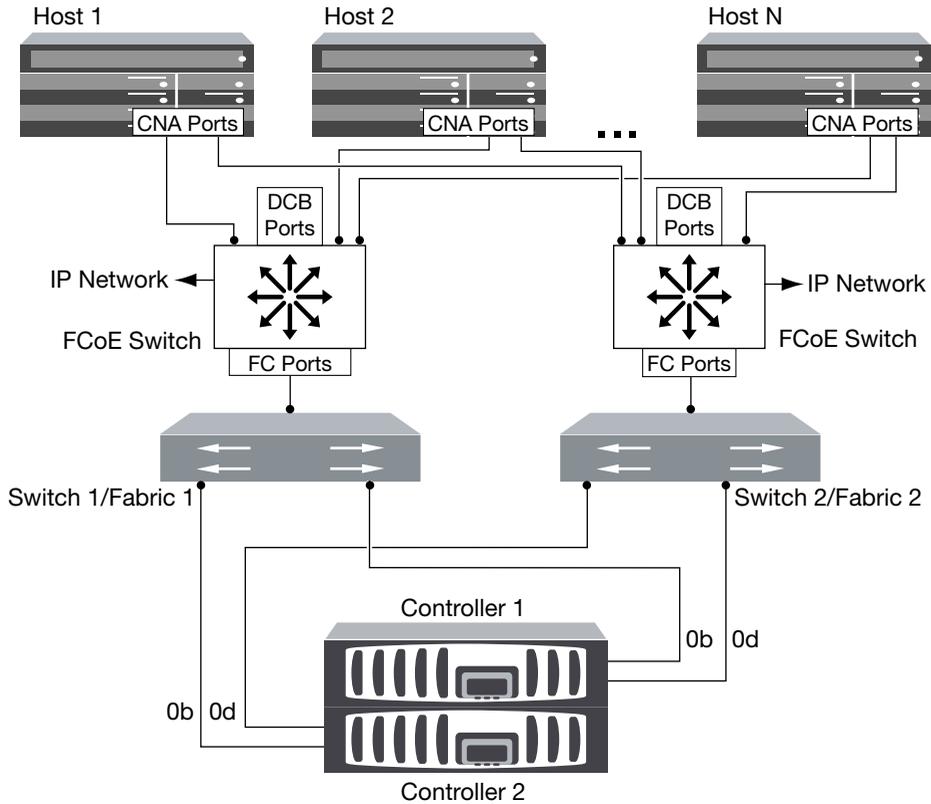
FCoE initiator to FC target configuration

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 following illustration shows an FCoE initiator to FC dual-fabric HA configuration. The FC target expansion adapter port numbers (0b and 0d) are examples. The actual port numbers

might vary, depending on the expansion slots in which the FCoE target expansion adapters are installed.



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, 8-Gb, or 16-Gb FC ports per controller using FC target expansion adapters
Multipathing required	Yes
Type of configuration	HA pair

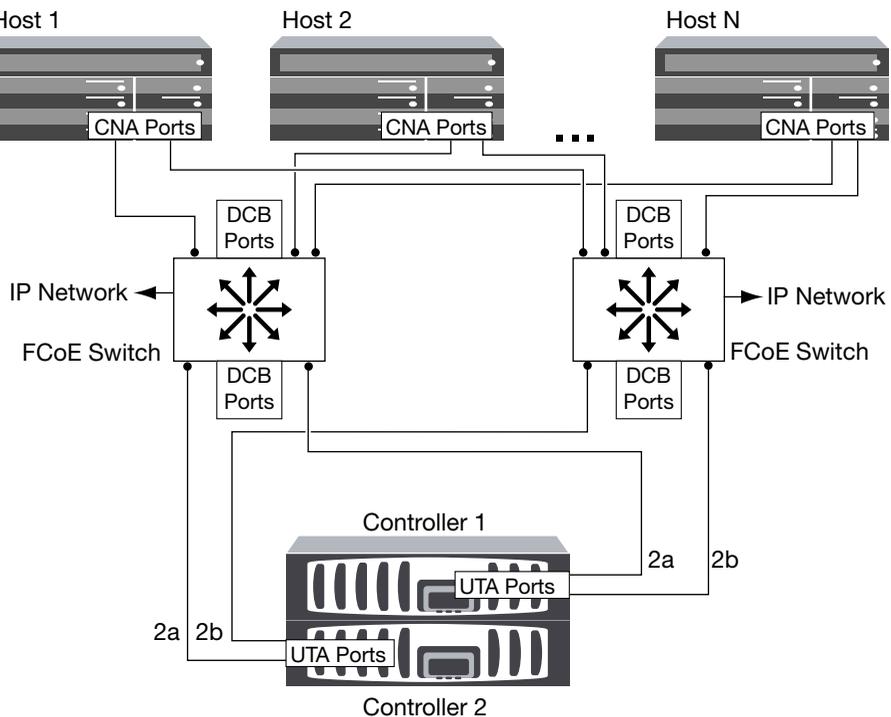
FCoE end-to-end configuration

You can connect hosts to both controllers in an HA pair. This connection is achieved using host 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 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 slots in which the FCoE target expansion adapters are installed.



Attribute	Value
Fully redundant	Yes
Type of fabric	Dual fabric

Attribute	Value
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

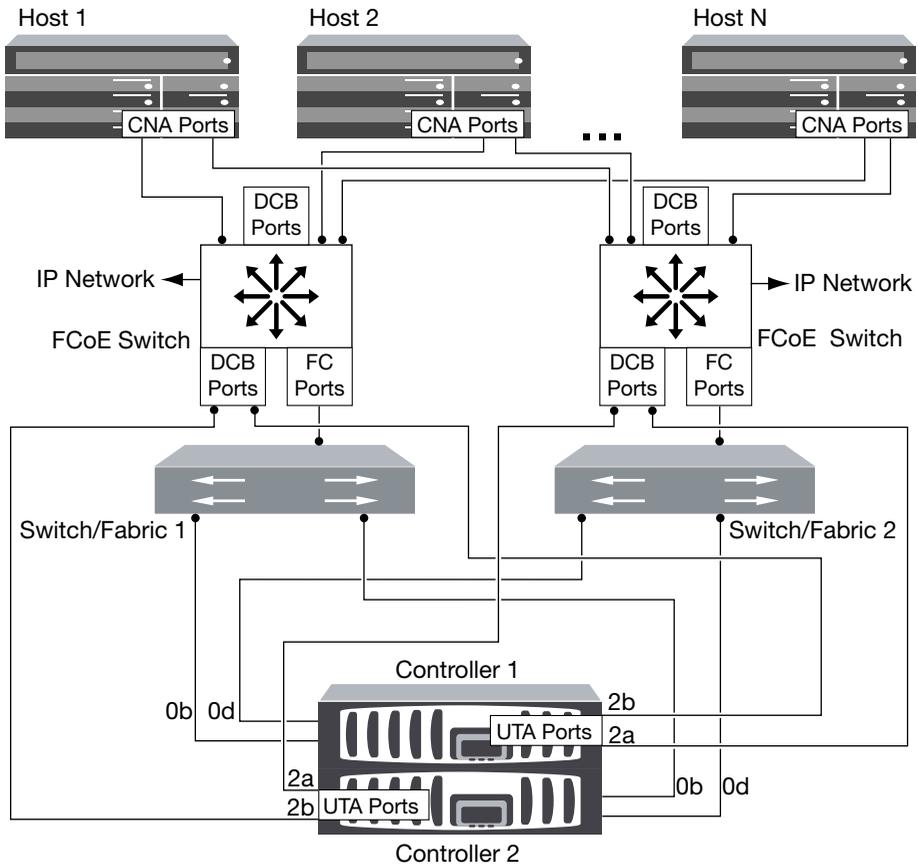
FCoE mixed with FC

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.



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, 8-Gb, or 16-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.

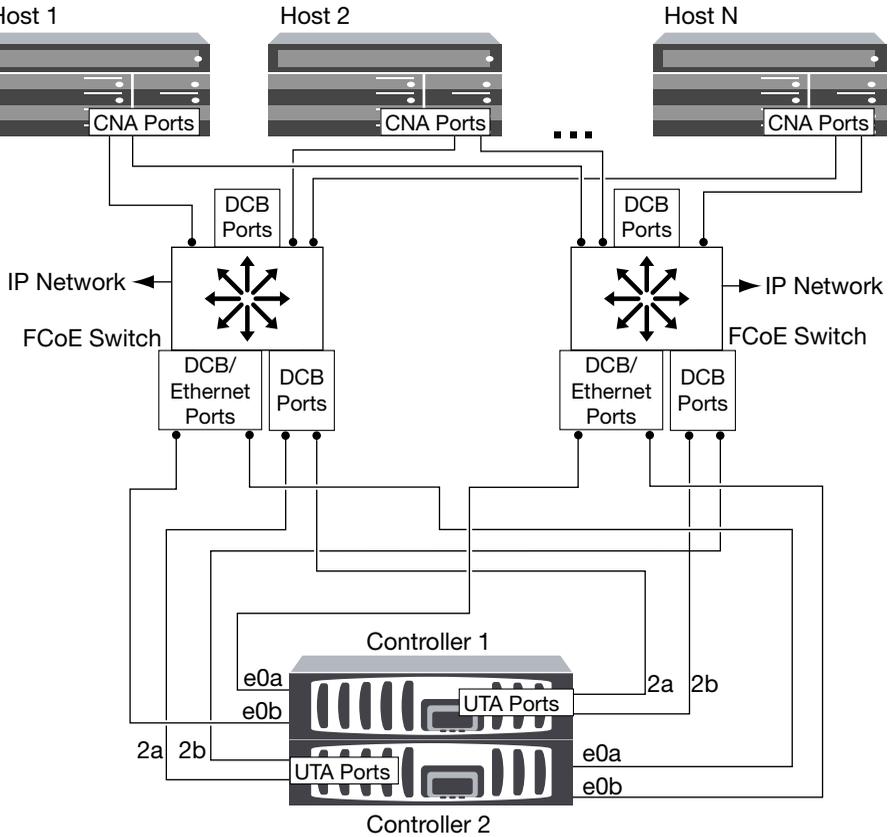
The FCoE target adapter is also called a unified target adapter or UTA. Like the CNA in a host, the UTA supports both FCoE and regular Ethernet traffic.

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 use 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.



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 configuration

Related information

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

Fibre Channel and FCoE zoning

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

Reasons for zoning

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

Recommendations for zoning

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

World Wide Name-based zoning

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

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

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

Individual zones

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

For Fibre Channel paths to nodes running clustered Data ONTAP, be sure the FC switches are zoned using the WWPNs of the target logical interfaces (LIFs), not the WWPNs of the physical ports on the node. The WWPNs of the physical ports start with "50" and the WWPNs of the LIFs start with "20."

Single-fabric zoning

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

Each initiator should have a minimum of one LIF from each node that the initiator can access. The zoning should allow at least one path from the host initiator to every node in the cluster to provide a path for LUN connectivity. This means that each initiator on the host might only have one target LIF per node in its zone configuration. If there is a requirement for multipathing to the same node, then each node will have multiple LIFs per node in its zone configuration. This enables the host to still access its LUNs if a node fails or a volume containing the LUN is moved to a different node.

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

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

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

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

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

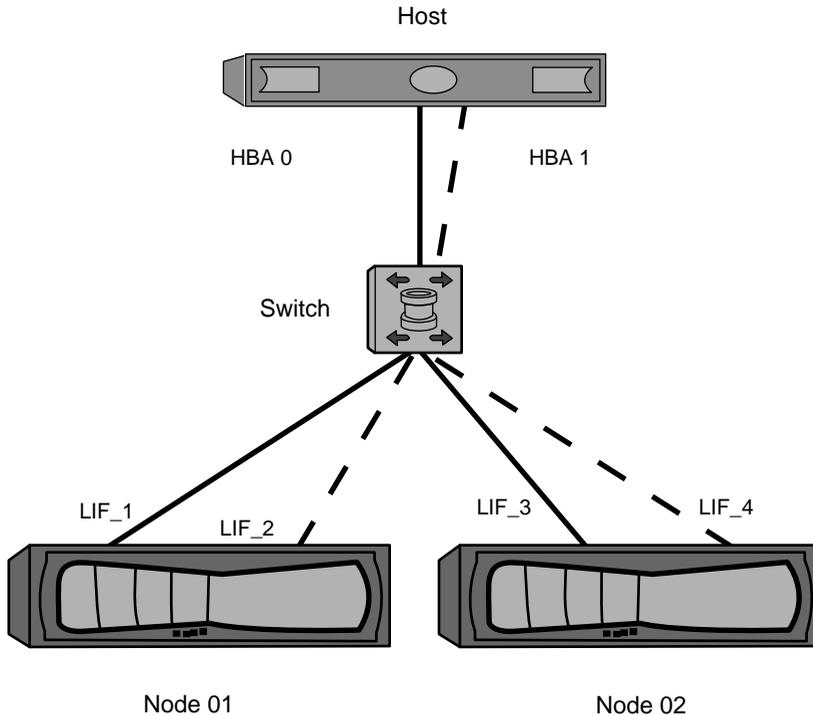


Figure 3: Single-fabric zoning

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

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

Note: The host operating system and multipathing software have to support the number of supported paths that are being used to access the LUNs on the nodes. To determine the number of paths used to access the LUNs on [nodes or storage controllers], see the configuration limits information elsewhere in this document.

Dual-fabric HA pair zoning

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

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

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

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

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

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

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

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

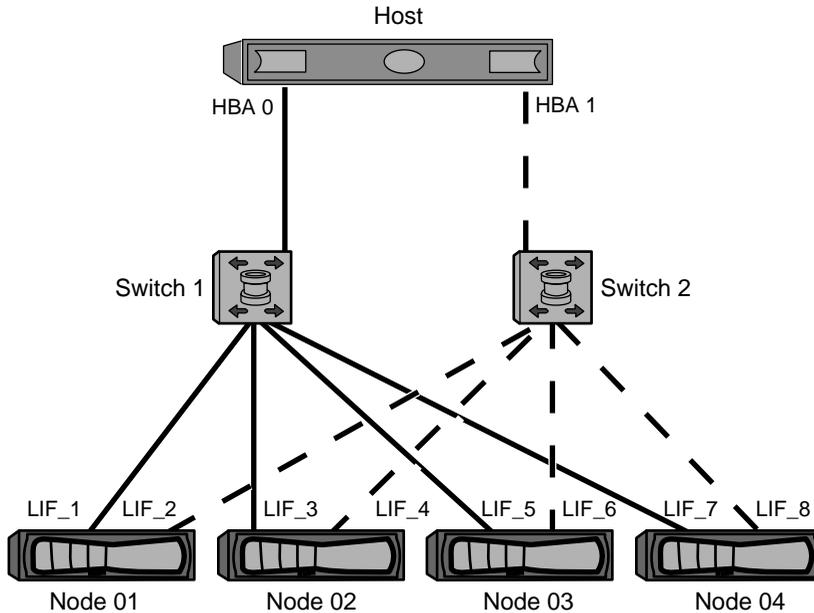


Figure 4: Dual-fabric zoning

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.

Support for shared SAN configurations

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

For information about the specific storage systems supported, and any steps you must take to ensure coexistence with other vendors' storage systems, see the Policies & Guidelines section of the NetApp Interoperability Matrix (IMT).

Related information

[NetApp Interoperability Matrix Tool](#)

ALUA Configurations

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

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

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

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

Related information

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

[NetApp Interoperability Matrix: support.netapp.com/NOW/products/interoperability/](https://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 Storage Virtual Machines (SVMs) with FlexVol volumes. Infinite Volumes are not supported for SAN.

In this deliverable, "volume" always means "FlexVol volume" and "SVM" always means SVM with FlexVol volumes.

SAN configuration limit parameters and definitions

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

Parameter	Definition
Visible target ports per host (iSCSI)	The maximum number of target iSCSI Ethernet ports that a host can see or access on iSCSI-attached controllers.
Visible target ports per host (FC)	The maximum number of FC adapters that a host can see or access on the attached FC controllers.
LUNs per host	The maximum number of LUNs that you can map from the controllers to a single host.
Maximum paths from host to LUN	The maximum number of paths from the host to a single LUN Note: Using the maximum number of paths is not recommended.
Maximum paths from host to storage solution	The maximum total number of paths from the host to the entire cluster.
Maximum LUN size	The maximum size of an individual LUN on the respective operating system.
Storage Virtual Machines (SVMs)	The total number of SVMs, including the default node SVM. There are limits on SVMs per node and SVMs per cluster.
Volumes per node	The total number of volumes supported on a single node.

Parameter	Definition
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 initiator-target nexus (ITNs) that can connect to a single FC port on a controller. Connecting the maximum number of ITNs is generally not recommended, and you might need to tune the FC queue depths on the host to achieve this maximum value.
FC LIFs per port	The maximum number of FC logical interfaces (LIFs) that can be defined on a single physical FC port.
iSCSI sessions per controller or node	The recommended maximum number of iSCSI sessions that you can connect to a single controller. The general formula to calculate this is as follows: Maximum sessions = $8 \times \text{System Memory} \div 512 \text{ MB}$.
IP LIFs per port	The maximum combined number of iSCSI, NFS, CIFS, and logical interfaces (LIFs) that can be defined on a single physical Ethernet port.
LIFs per port set	The maximum number of logical interfaces that can be assigned to a single port set.
FC ITNs per controller	The maximum number of initiator-target nexus (ITNs) supported per port per controller.
igroups per controller	The maximum number of initiator groups that you can configure per controller.
Initiators per igroup	The maximum number of FC initiators (HBA WWNs) or iSCSI initiators (host iqn/eui node names) that you can include in a single igroup.
LUN mappings per controller	The maximum number of LUN mappings per controller. For example, a LUN mapped to two igroups counts as two mappings.
LUN path name length	The maximum number of characters in a full LUN name. For example, /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.

Parameter	Definition
Ethernet ports per node	The maximum number of supported Ethernet ports per node.
FC target ports per controller or node	The maximum number of supported FC target ports per controller. FC initiator ports used for back-end disk connections, for example, connections to disk shelves, are not included in this number.
Port sets per node	The maximum number of port sets that can be created on a single node.

Related tasks

[Calculating queue depth](#) on page 57

Determining the number of supported nodes for SAN configurations

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

About this task

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

Steps

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

Related references

[SAN configuration limits](#) on page 48

Related information

[NetApp Hardware Universe](#)

Determining the number of supported hosts per cluster in FC configurations

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

About this task

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

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

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

Steps

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

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

The number of entries displayed at the bottom of the command output represents your number of ITNs for that LIF.
3. Record the number of ITNs displayed for each LIF.
4. Add the number of ITNs for each LIF on every node in your cluster.

This total represents the number of ITNs in your cluster.

Determining the supported number of hosts per cluster in iSCSI configurations

The maximum number of SAN hosts that can be connected to a cluster varies greatly based on your specific combination of multiple cluster attributes, such as the number of hosts connected to each cluster node, initiators per host, logins per host, and nodes in the cluster. You need to determine the number of iSCSI sessions per node to know whether you can add more hosts to your cluster.

About this task

As long as you are below the maximum number of iSCSI sessions per node, you can continue to add hosts to your cluster. The maximum number of iSCSI sessions per node varies based on the types of controllers in your cluster.

Steps

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

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

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

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

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

Related references

[SAN configuration limits](#) on page 48

Host operating system limits for SAN configurations

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

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

Note: AIX 6.0 Host Utilities do not support iSCSI.

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

Related references

[SAN configuration limit parameters and definitions](#) on page 42

SAN configuration limits

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

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

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

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

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

Parameters	Maximum per node
LUNs	1,024 for the FAS2220 2,048 for the following models: <ul style="list-style-type: none">• FAS2240• FAS2520• FAS2552• FAS2554• FAS3220• FAS3240 8,192 for the following models: <ul style="list-style-type: none">• FAS3250• 3270• 6210• 6220• 6240• FAS8020• FAS8040 12,288 for the following models: <ul style="list-style-type: none">• 6250• 6280• 6290• FAS8060• FAS8080

Parameters	Maximum per node
LUN mappings	1,024 for the FAS2220 2,048 for the following models: <ul style="list-style-type: none"> • FAS2240 • FAS2520 • FAS2552 • FAS2554 • FAS3220 • FAS3240 8,192 for the following models: <ul style="list-style-type: none"> • FAS3250 • 3270 • 6210 • 6220 • 6240 • FAS8020 • FAS8040 12,288 for the following models: <ul style="list-style-type: none"> • 6250 • 6280 • 6290 • FAS8060 • FAS8080
LUNs per volume	512 for all models
LUN size	16 TB for all models

Parameters	Maximum per node
iSCSI sessions per node	1,024 for the FAS2220 2,048 for the following models: <ul style="list-style-type: none"> • FAS2240 • FAS2520 • FAS2552 • FAS2554 • FAS3220 • FAS3240 4,096 for the following models: <ul style="list-style-type: none"> • FAS3250 • 3270 • 6210 • 6220 • 6240 • FAS8020 • FAS8040 8,192 for the following models: <ul style="list-style-type: none"> • 6250 • 6280 • 6290 • FAS8060 • FAS8080
iSCSI LIFs per node	256 for all models
iSCSI LIFs per port	32 for all models
iSCSI LIFs per port set	32 for all models

Parameters	Maximum per node
igroups	256 for the FAS2220 512 for the following models: <ul style="list-style-type: none"> • FAS2240 • FAS2520 • FAS2552 • FAS2554 • FAS3220 • FAS3240 2,048 for the following models: <ul style="list-style-type: none"> • FAS3250 • 3270 • 6210 • 6220 • 6240 • FAS8020 • FAS8040 4,096 for the following models: <ul style="list-style-type: none"> • 6250 • 6280 • 6290 • FAS8060 • FAS8080
Initiators per igroup	128 for the FAS2220 256 for all other models

Parameters	Maximum per node
Initiators per node	512 for the FAS2220 1,024 for the following models: <ul style="list-style-type: none">• FAS2240• FAS2520• FAS2552• FAS2554• FAS3220• FAS3240 2,048 for the following models: <ul style="list-style-type: none">• FAS3250• 3270• 6210• 6220• 6240• FAS8020• FAS8040 4,096 for the following models: <ul style="list-style-type: none">• 6250• 6280• 6290• FAS8060• FAS8080

Parameters	Maximum per node
Portsets	256 for the FAS2220 512 for the following models: <ul style="list-style-type: none"> • FAS2240 • FAS2520 • FAS2552 • FAS2554 • FAS3220 • FAS3240 2,048 for the following models: <ul style="list-style-type: none"> • FAS3250 • 3270 • 6210 • 6220 • 6240 • FAS8020 • FAS8040 4,096 for the following models: <ul style="list-style-type: none"> • 6250 • 6280 • 6290 • FAS8060 • FAS8080
Ethernet ports	See <i>Hardware Universe</i> for platform supported limits.
<p>Note: The following FC limits do not apply to the FAS2220 or the FAS2520. The FAS2220 and FAS2520 do not support FC. These FC limits are the same for all other platforms.</p>	
FC ports	See <i>Hardware Universe</i> for platform supported limits.
FC queue depth available per port	2,048 for all models
FC ITNs per port	512 for all models

Parameters	Maximum per node
FC ITNs per node	8,192 for the following models: <ul style="list-style-type: none"> • 6250 • 6280 • 6290 • FAS8060 • FAS8080 4,096 for the following models: <ul style="list-style-type: none"> • FAS3250 • 3270 • 6210 • 6220 • 6240 • FAS8020 • FAS8040 2,048 for the following models: <ul style="list-style-type: none"> • 6250 • 6280 • 6290 • FAS8060 • FAS8080
FC (FC/FCoE) LIFs per port	32 for all models
FC LIFs portset	32 for all models
FC LIFs per node	512 for all models

Related tasks

Determining the number of supported nodes for SAN configurations on page 44

Determining the number of supported hosts per cluster in FC configurations on page 45

Determining the supported number of hosts per cluster in iSCSI configurations on page 46

Related information

Hardware Universe: hwu.netapp.com

SAN configuration limits for Data ONTAP-v platforms

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

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

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

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

FC switch configuration limits

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

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

Brocade switch limits

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

Cisco Systems switch limits

You can find the configuration limits for Cisco switches in the *Cisco Configuration Limits* guide for your version of Cisco switch software.

Related information

[Cisco Configuration Limits - www.cisco.com/en/US/products/ps5989/products_installation_and_configuration_guides_list.html](http://www.cisco.com/en/US/products/ps5989/products_installation_and_configuration_guides_list.html)

Calculating queue depth

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

About this task

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

In a configuration with multiple initiators (hosts), all hosts should have similar queue depths. This prevents hosts with small queue depths from being starved by hosts with large queue depths.

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

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

Steps

1. Count the total number of FC initiators in all the hosts that connect to one FC target port.
2. Multiply by 128.
 - If the result is less than 2,048, set the queue depth for all initiators to 128.

Example

You have 15 hosts with one initiator connected to each of two target ports on the storage controller. $15 \times 128 = 1,920$. Because 1,920 is less than the total queue depth limit of 2,048, you can set the queue depth for all your initiators to 128.

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

Example

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

3. Choose one of the following options.
 - Option 1:
 - a. Add more FC target ports.
 - b. Redistribute your FC initiators.
 - c. Repeat steps 1 and 2.

Example

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

- Option 2:
 - a. Designate each host as "large" or "small" based on its expected I/O need.
 - b. Multiply the number of large initiators by 128.
 - c. Multiply the number of small initiators by 32.
 - d. Add the two results together.
 - e. If the result is less than 2,048, set the queue depth for "large" hosts to 128 and the queue depth for "small" hosts to 32.
 - f. If the result is still greater than 2,048 per port, reduce the queue depth per initiator until the total queue depth is less than or equal to 2,048.

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

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

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

Example

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

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

Example

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

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

Setting queue depths on AIX hosts

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

Examples:

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

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

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

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

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

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

Setting queue depths on HP-UX hosts

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

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

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

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

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

```
scsictl -a /dev/rdisk/c2t2d0
```

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

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

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

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

Setting queue depths on Solaris hosts

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

About this task

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

Steps

1. `# cd/kernel/drv`
2. `# vi lpfc.conf`
3. Search for `/tft-queue (/tgt-queue)`
`tgt-queue-depth=32`

Note: The default value is set to 32 at installation.

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

Setting queue depths on VMware hosts

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

To set maximum queue depth for a QLogic HBA

Steps

1. Log on to the service console as the root user.
2. Use the `#vmkload_mod -l` command to verify which Qlogic HBA module is currently loaded.
3. For a single instance of a Qlogic HBA, run the following command:
`#esxcfg-module -s ql2xmaxqdepth=64 qla2300_707`

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

4. Save your changes using the following command:

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

5. Reboot the server using the following command:

```
#reboot
```

6. Confirm the changes using the following commands:

```
a. #esxcfg-module -g qla2300_707
```

```
b. qla2300_707 enabled = 1 options = 'ql2xmaxqdepth=64'
```

To change the queue depth of an Emulex HBA

Steps

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

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

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

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

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

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

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

5. Enter the following command: `#esxcfg-boot -b`
6. Reboot using `#reboot`.

Setting queue depths on Windows hosts

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

To update Emulex HBA queue depths

Steps

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

Note: If you are setting **QueueDepth** greater than 150, the following Windows Registry value also need to be increased appropriately: `HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\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 list box.
4. Update the `Execution Throttle` parameter.

Copyright information

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

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

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

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

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

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

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

Trademark information

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

Cisco and the Cisco logo are trademarks of Cisco in the U.S. and other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

How to send comments about documentation and receive update notification

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

If you have suggestions for improving this document, send us your comments by email to 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.

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

You can also contact us in the following ways:

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

Index

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

A

- adapters
 - how to prevent loss of connectivity when using X1133A-R6 [14](#)
 - UTA2 also referred to as CNA [14](#)
- AIX hosts
 - setting queue depth for [60](#)
- ALUA configurations [41](#)
- asymmetric logical unit access (ALUA) configurations [41](#)

B

- best practices
 - FC switch configuration [17](#)
- Brocade switches
 - FC configuration limits [57](#)

C

- Cisco switches
 - FC and FCoE zoning requirement [39](#)
 - FC configuration limits [57](#)
- CNA adapters
 - also referred to as UTA2 adapters [14](#)
- CNA ports
 - also referred to as UTA2 ports [14](#)
- comments
 - how to send feedback about documentation [66](#)
- configuration limits
 - by SAN host operating system [46](#)
 - determining supported number of hosts per cluster in iSCSI [46](#)
 - FC switch [57](#)
 - FDvM200 storage systems [56](#)
 - hosts per node [45](#)
 - SAN [48](#)
 - SAN parameter, defined [42](#)

- volume and SVM requirements [42](#)
- configurations
 - FAS2240 single-fabric single-node FC [20](#)
 - iSCSI [5](#)
 - single-node, support for connecting hosts in [18](#)
 - support for FCoE initiator to FC target [28](#)
 - support for FCoE initiator to FCoE and FC mixed target [31](#)
 - support for FCoE initiator to FCoE target [30](#)
 - support for FCoE initiator to FCoE target mixed with IP traffic [33](#)
 - supported FC [11](#)
 - supported FCoE [28](#)
- configuring
 - unified target adapters [15](#)
- connectivity loss
 - how to prevent when using X1133A-R6 adapters [14](#)

D

- Data ONTAP-v
 - configuration limits [56](#)
- direct-attached configurations
 - example of iSCSI SAN single node [9](#)
- documentation
 - how to receive automatic notification of changes to [66](#)
 - how to send feedback about [66](#)
- dynamic VLANs [10](#)

E

- ESX
 - SAN host configuration limits [46](#)
- expansion FC adapter
 - supported port speed [15](#)

F

- fabrics
 - support for connecting hosts in single-node configuration using single or multiple FC [18](#)
- FAS2240
 - multifabric single-node configurations [23](#)
 - single-fabric single-node FC configurations [20](#)
- FAS80xx

- X1143A-R6 adapters [14](#)
- FC
 - multifabric switch zoning [38](#)
 - port speed [15](#)
 - single-fabric switch zoning [36](#)
 - supported configurations [11](#)
 - supported port speed [15](#)
 - switch configuration best practices [17](#)
 - switch hop count [12](#)
 - switch WWN zoning [36](#)
 - switch zoning with individual zones [36](#)
 - target port configuration recommendations [16](#)
 - zoning for [35](#)
- FC configurations
 - support for connecting hosts in single-node [18](#)
- FC fabrics
 - support for connecting hosts in single-node configurations using single or multiple [18](#)
- FC LIF
 - zoning restrictions for Cisco switches [39](#)
- FC ports
 - support for onboard and expansion adapter [13](#)
- FC switches
 - configuration limits [57](#)
- FCoE
 - initiator and target combinations [27](#)
 - maximum supported switch hop count [27](#)
 - overview [26](#)
 - supported configurations [28](#)
 - zoning for [35](#)
- FCoE configurations
 - support for FCoE initiator to FC target [28](#)
 - support for FCoE initiator to FCoE and FC mixed target [31](#)
 - support for FCoE initiator to FCoE target [30](#)
 - support for FCoE initiator to FCoE target mixed with IP traffic [33](#)
- FCoE LIF
 - zoning restrictions for Cisco switches [39](#)
- FDvM200
 - configuration limits [56](#)
- feedback
 - how to send comments about documentation [66](#)
- Fibre Channel over Ethernet
 - overview [26](#)
- FlexVol volumes
 - required for SAN [42](#)

H

- HA configurations
 - support for single-fabric [21](#)
- HA pairs
 - support for iSCSI SAN single-network configurations [7](#)
 - support for multi-network, in iSCSI SAN configurations [8](#)
- hop count
 - for FC switches [12](#)
- hop counts
 - maximum supported for FCoE switches [27](#)
- host multipathing software
 - when required [12](#)
- host utilities
 - AIX queue depths [60](#)
 - HP-UX queue depths [60](#)
 - Solaris queue depths [61](#)
 - VMware queue depths [61](#)
 - Windows queue depths [63](#)
- hosts per cluster
 - determining supported number in iSCSI configurations [46](#)
- hosts per node [45](#)
- HP-UX
 - SAN host configuration limits [46](#)
- HP-UX hosts
 - setting queue depth for [60](#)

I

- Infinite Volumes
 - not supported for SAN [42](#)
- information
 - how to send feedback about improving documentation [66](#)
- initiators
 - FCoE and FC combinations [27](#)
- IP traffic
 - in FCoE configurations, support for [33](#)
- iSCSI
 - configurations [5](#)
 - dynamic VLANs [10](#)
 - example of direct-attached single-node SAN configuration [9](#)
 - multi-network configuration [6](#)
 - single-network configuration [6](#)
 - static VLANs [10](#)

- support for multi-network HA pairs in iSCSI SAN configurations [8](#)
- support for single-network HA pair SAN configurations [7](#)
- using VLANs [10](#)
- iSCSI configurations
 - determining supported number of hosts per cluster [46](#)
- ITNs
 - per node limits [45](#)

L

- LIFs
 - FC and FCoE zoning restrictions for Cisco switches [39](#)
- Linux
 - SAN host configuration limits [46](#)
- LUNs
 - recommended number of paths to avoid single points of failure [5](#), [11](#), [26](#)

M

- MPIO software
 - when required [12](#)
- multifabric configurations [23](#)
- multifabric HA configurations
 - described and illustrated [24](#)
- multifabric single-node configurations
 - FAS2240 [23](#)
- multifabric single-node FC configurations
 - support for connecting hosts using [18](#)
- multipathing software
 - when required [12](#)

N

- nodes
 - determining number of supported, for SAN configurations [44](#)

O

- onboard FC port
 - supported port speed [15](#)
- overviews
 - FCoE [26](#)

P

- parameters
 - SAN configuration limit, defined [42](#)
- paths
 - recommended number to avoid single points of failure [5](#), [11](#), [26](#)
- port speed
 - supported for FC [15](#)
- ports
 - support for onboard and expansion adapter FC [13](#)
 - UTA2 also referred to as CNA [14](#)
 - X1143A-R6 [15](#)

Q

- queue depths
 - calculating [57](#)
 - setting for AIX hosts [60](#)
 - setting for HP-UX hosts [60](#)
 - setting for Solaris hosts [61](#)
 - setting for VMware hosts [61](#)
 - setting for Windows hosts [63](#)
 - tuning [57](#)

R

- recommended number of paths
 - for avoiding single points of failure [5](#), [11](#), [26](#)

S

- SAN
 - configuration limits [48](#)
- SAN configuration limits
 - by host operating system [46](#)
 - FC switch [57](#)
 - FDvM200 storage systems [56](#)
 - parameters defined [42](#)
 - volume and SVM requirements [42](#)
- SAN configurations
 - determining supported number of nodes [44](#)
 - support for iSCSI single-network HA pair configurations [7](#)
 - support for multi-network HA pairs in iSCSI [8](#)
 - support for shared [40](#)
- shared SAN configurations
 - support for [40](#)
- single fabric configurations [20](#)
- single points of failure

70 | SAN Configuration Guide

- recommended number of paths to avoid [5](#), [11](#), [26](#)
- single-fabric HA configurations
 - support for [21](#)
- single-fabric single-node FC configurations
 - FAS2240 [20](#)
 - support for connecting hosts using [18](#)
- single-node configurations
 - support for connecting hosts in [18](#)
- soft zoning
 - FC switch [36](#)
- Solaris
 - SAN host configuration limits [46](#)
- Solaris hosts
 - setting queue depth for [61](#)
- static VLANs [10](#)
- suggestions
 - how to send feedback about documentation [66](#)
- supported configurations
 - FCoE [28](#)
- supported nodes
 - determining number of, for SAN configurations [44](#)
- SVMs
 - restrictions for SAN [42](#)
- switch
 - FC and FCoE zoning requirement [39](#)
 - FC hop count [12](#)
 - FC multifabric zoning [38](#)
 - FC single-fabric zoning [36](#)
 - FC WWN zoning [36](#)
 - FC zoning with individual zones [36](#)
- switches
 - FC configuration best practices [17](#)
 - FC configuration limits [57](#)
 - maximum supported hop count for FCoE [27](#)

T

- target port configurations
 - FC recommendations for [16](#)
- targets
 - FCoE and FC combinations [27](#)
- twitter
 - how to receive automatic notification of documentation changes [66](#)

U

- unified target adapters

- default configuration [15](#)
 - x1143A-R6 [14](#)
- UTA2
 - also referred to as CNA adapters [14](#)
- UTA2 ports
 - also referred to as CNA ports [14](#)

V

- virtual LANs
 - reasons for using [10](#)
- VLANs
 - dynamic [10](#)
 - reasons for using [10](#)
 - static [10](#)
- VMware hosts
 - setting queue depth for [61](#)
- volumes
 - restrictions for SAN [42](#)

W

- Windows
 - SAN host configuration limits [46](#)
- Windows hosts
 - setting queue depth for [63](#)
- WWN zoning
 - FC switch [36](#)

X

- X1133A-R6 adapters
 - how to prevent loss of connectivity when using [14](#)
- X1143A-R6 adapters
 - ports [14](#)

Z

- zoning
 - FC switch by WWN [36](#)
 - FC switch multifabric [38](#)
 - FC switch single-fabric [36](#)
 - FC switch with individual zones [36](#)
 - for FC and FCoE [35](#)
 - requirements for Cisco switches [39](#)