NetApp

Configure network ports

ONTAP 9

NetApp
February 20, 2026

Table of Contents

Configure network ports

Combine physical ports to create ONTAP interface groups

An interface group, also known as a Link Aggregation Group (LAG), is created by combining two or more physical ports on the same node into a single logical port. The logical port provides increased resiliency, increased availability, and load sharing.

Interface group types

Three types of interface groups are supported on the storage system: single-mode, static multimode, and dynamic multimode. Each interface group provides different levels of fault tolerance. Multimode interface groups provide methods for load balancing network traffic.

Characteristics of single-mode interface groups

In a single-mode interface group, only one of the interfaces in the interface group is active. The other interfaces are on standby, ready to take over if the active interface fails.

Characteristics of a single-mode interface groups:

- For failover, the cluster monitors the active link and controls failover.
  Because the cluster monitors the active link, there is no switch configuration required.
- There can be more than one interface on standby in a single-mode interface group.
- If a single-mode interface group spans multiple switches, you must connect the switches with an Inter-Switch link (ISL).
- For a single-mode interface group, the switch ports must be in the same broadcast domain.
- Link-monitoring ARP packets, which have a source address of 0.0.0.0, are sent over the ports to verify that the ports are in the same broadcast domain.

The following figure is an example of a single-mode interface group. In the figure, e0a and e1a are part of the a0a single-mode interface group. If the active interface, e0a, fails, the standby e1a interface takes over and maintains the connection to the switch.
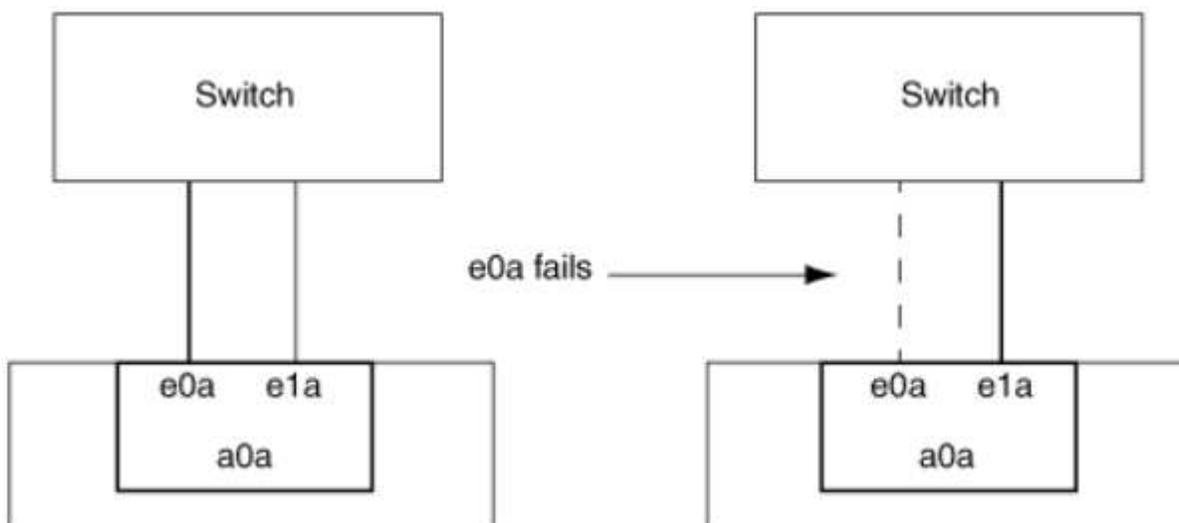
> ⓘ To accomplish single-mode functionality, the recommended approach is to instead use failover groups. By using a failover group, the second port can still be used for other LIFs and need not remain unused. Additionally, failover groups can span more than two ports and can span ports on multiple nodes.

**Characteristics of static multimode interface groups**

The static multimode interface group implementation in ONTAP complies with IEEE 802.3ad (static). Any switch that supports aggregates, but does not have control packet exchange for configuring an aggregate, can be used with static multimode interface groups.

Static multimode interface groups do not comply with IEEE 802.3ad (dynamic), also known as Link Aggregation Control Protocol (LACP). LACP is equivalent to Port Aggregation Protocol (PAgP), the proprietary link aggregation protocol from Cisco.

The following are characteristics of a static multimode interface group:

- All interfaces in the interface group are active and share a single MAC address.
  - Multiple individual connections are distributed among the interfaces in the interface group.
  - Each connection or session uses one interface within the interface group.
    When you use the sequential load balancing scheme, all sessions are distributed across available links on a packet-by-packet basis, and are not bound to a particular interface from the interface group.
- Static multimode interface groups can recover from a failure of up to "n-1" interfaces, where n is the total number of interfaces that form the interface group.
- If a port fails or is unplugged, the traffic that was traversing the failed link is automatically redistributed to one of the remaining interfaces.
- Static multimode interface groups can detect a loss of link, but they cannot detect a loss of connectivity to the client or switch misconfigurations that might impact connectivity and performance.
- A static multimode interface group requires a switch that supports link aggregation over multiple switch ports.
  The switch is configured so that all ports to which links of an interface group are connected are part of a single logical port. Some switches might not support link aggregation of ports configured for jumbo frames. For more information, see your switch vendor's documentation.
- Several load balancing options are available to distribute traffic among the interfaces of a static multimode interface group.

The following figure is an example of a static multimode interface group. Interfaces e0a, e1a, e2a, and e3a are part of the a1a multimode interface group. All four interfaces in the a1a multimode interface group are active.
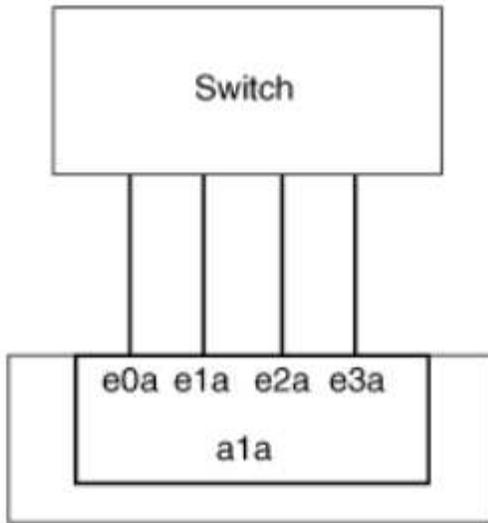
Several technologies exist that enable traffic in a single aggregated link to be distributed across multiple physical switches. The technologies used to enable this capability vary among networking products. Static multimode interface groups in ONTAP conform to the IEEE 802.3 standards. If a particular multiple switch link aggregation technology is said to interoperate with or conform to the IEEE 802.3 standards, it should operate with ONTAP.

The IEEE 802.3 standard states that the transmitting device in an aggregated link determines the physical interface for transmission. Therefore, ONTAP is only responsible for distributing outbound traffic, and cannot control how inbound frames arrive. If you want to manage or control the transmission of inbound traffic on an aggregated link, that transmission must be modified on the directly connected network device.

Dynamic multimode interface group

Dynamic multimode interface groups implement Link Aggregation Control Protocol (LACP) to communicate group membership to the directly attached switch. LACP enables you to detect the loss of link status and the inability of the node to communicate with the direct-attached switch port.

Dynamic multimode interface group implementation in ONTAP complies with IEEE 802.3 AD (802.1 AX). ONTAP does not support Port Aggregation Protocol (PAgP), which is a proprietary link aggregation protocol from Cisco.

A dynamic multimode interface group requires a switch that supports LACP.

ONTAP implements LACP in nonconfigurable active mode that works well with switches that are configured in either active or passive mode. ONTAP implements the long and short LACP timers (for use with nonconfigurable values 3 seconds and 90 seconds), as specified in IEEE 802.3 AD (802.1AX).

The ONTAP load balancing algorithm determines the member port to be used to transmit outbound traffic, and does not control how inbound frames are received. The switch determines the member (individual physical port) of its port channel group to be used for transmission, based on the load balancing algorithm configured in the switch's port channel group. Therefore, the switch configuration determines the member port (individual physical port) of the storage system to receive traffic. For more information about configuring the switch, see the documentation from your switch vendor.

If an individual interface fails to receive successive LACP protocol packets, then that individual interface is marked as "lag_inactive" in the output of "ifgrp status" command. Existing traffic is automatically rerouted to any remaining active interfaces.

The following rules apply when using dynamic multimode interface groups:

- Dynamic multimode interface groups should be configured to use the port-based, IP-based, MAC-based, or round robin load balancing methods.

- In a dynamic multimode interface group, all interfaces must be active and share a single MAC address.

The following figure is an example of a dynamic multimode interface group. Interfaces e0a, e1a, e2a, and e3a are part of the a1a multimode interface group. All four interfaces in the a1a dynamic multimode interface group are active.
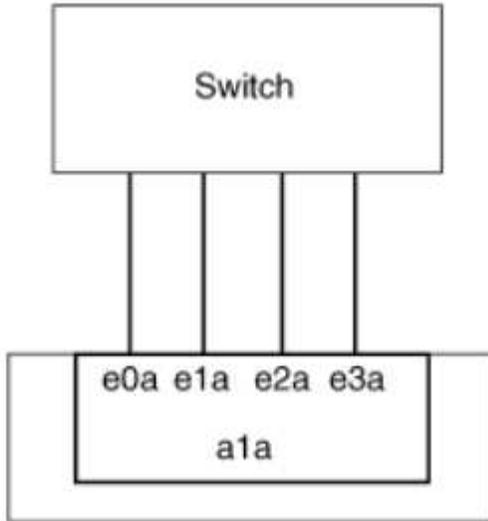
Load balancing in multimode interface groups

You can ensure that all interfaces of a multimode interface group are used equally for outgoing traffic by using the IP address, MAC address, sequential, or port-based load balancing methods to distribute network traffic equally over the network ports of a multimode interface group.

The load balancing method for a multimode interface group can be specified only when the interface group is created.

Best Practice: Port-based load balancing is recommended whenever possible. Use port-based load balancing unless there is a specific reason or limitation in the network that prevents it.

Port-based load balancing

Port-based load balancing is the recommended method.

You can equalize traffic on a multimode interface group based on the transport layer (TCP/UDP) ports by using the port-based load balancing method.

The port-based load balancing method uses a fast hashing algorithm on the source and destination IP addresses along with the transport layer port number.

IP address and MAC address load balancing

IP address and MAC address load balancing are the methods for equalizing traffic on multimode interface groups.

These load balancing methods use a fast hashing algorithm on the source and destination addresses (IP

address and MAC address). If the result of the hashing algorithm maps to an interface that is not in the UP link-state, the next active interface is used.

> (i) Do not select the MAC address load balancing method when creating interface groups on a system that connects directly to a router. In such a setup, for every outgoing IP frame, the destination MAC address is the MAC address of the router. As a result, only one interface of the interface group is used.

IP address load balancing works in the same way for both IPv4 and IPv6 addresses.

Sequential load balancing

You can use sequential load balancing to equally distribute packets among multiple links using a round robin algorithm. You can use the sequential option for load balancing a single connection's traffic across multiple links to increase single connection throughput.

However, because sequential load balancing may cause out-of-order packet delivery, extremely poor performance can result. Therefore, sequential load balancing is generally not recommended.

Create an interface group or LAG

You can create an interface group or LAG—single-mode, static multimode, or dynamic multimode (LACP)—to present a single interface to clients by combining the capabilities of the aggregated network ports.

The procedure you follow depends on the interface that you use—System Manager or the CLI:

System Manager

Use System Manager to create a LAG

Steps

1. Select **Network > Ethernet port > + Link Aggregation Group** to create a LAG.

2. Select the node from the drop-down list.

3. Choose from the following:

    a. ONTAP to **Automatically select broadcast domain (recommended)**.

    b. To manually select a broadcast domain.

4. Select the ports to form the LAG.

5. Select the mode:

    a. Single: Only one port is used at a time.

    b. Multiple: All ports can be used simultaneously.

    c. LACP: The LACP protocol determines the ports that can be used.

6. Select the load balancing:

    a. IP based

    b. MAC based

    c. Port

    d. Sequential

7. Save your changes.

CLI

Use the CLI to create an interface group

When creating a multimode interface group, you can specify any of the following load-balancing methods:

- `port`: Network traffic is distributed on the basis of the transport layer (TCP/UDP) ports. This is the recommended load-balancing method.

- `mac`: Network traffic is distributed on the basis of MAC addresses.

- `ip`: Network traffic is distributed on the basis of IP addresses.

- `sequential`: Network traffic is distributed as it is received.

> ⓘ The MAC address of an interface group is determined by the order of the underlying ports and how these ports initialize during bootup. You should therefore not assume that the ifgrp MAC address is persistent across reboots or ONTAP upgrades.

Step

Use the `network port ifgrp create` command to create an interface group.

Interface groups must be named using the syntax `a<number><letter>`. For example, a0a, a0b, a1c, and a2a are valid interface group names.

Learn more about `network port ifgrp create` in the ONTAP command reference.

The following example shows how to create an interface group named a0a with a distribution function of port and a mode of multimode:

```
network port ifgrp create -node cluster-1-01 -ifgrp a0a -distr-func port -mode
multimode
```

Add a port to an interface group or LAG

You can add up to 16 physical ports to an interface group or LAG for all port speeds.

The procedure you follow depends on the interface that you use—System Manager or the CLI:

System Manager
Use System Manager to add a port to a LAG

Steps

1. Select **Network > Ethernet port > LAG** to edit a LAG.
2. Select additional ports on the same node to add to the LAG.
3. Save your changes.

CLI
Use the CLI to add ports to an interface group

Step
Add network ports to the interface group:

```
network port ifgrp add-port
```

The following example shows how to add port e0c to an interface group named a0a:

```
network port ifgrp add-port -node cluster-1-01 -ifgrp a0a -port e0c
```

Beginning with ONTAP 9.8, interface groups are automatically placed into an appropriate broadcast domain about one minute after the first physical port is added to the interface group. If you do not want ONTAP to do this, and prefer to manually place the ifgrp into a broadcast domain, then specify the `-skip -broadcast-domain-placement` parameter as part of the `ifgrp add-port` command.

Learn more about `network port ifgrp add-port` and configuration restrictions that apply to port interface groups in the ONTAP command reference.

Remove a port from an interface group or LAG

You can remove a port from an interface group that hosts LIFs, as long as it is not the last port in the interface group. There is no requirement that the interface group must not host LIFs or that the interface group must not be the home port of a LIF considering that you are not removing the last port from the interface group. However, if you are removing the last port, then you must migrate or move the LIFs from the interface group first.

About this task

You can remove up to 16 ports (physical interfaces) from an interface group or LAG.

The procedure you follow depends on the interface that you use—System Manager or the CLI:

System Manager

Use System Manager to remove a port from a LAG

Steps

1. Select **Network > Ethernet port > LAG** to edit a LAG.
2. Select the ports to remove from the LAG.
3. Save your changes.

CLI

Use the CLI to remove ports from an interface group

Step

Remove network ports from an interface group:

```
network port ifgrp remove-port
```

Learn more about `network port ifgrp remove-port` in the ONTAP command reference.

The following example shows how to remove port e0c from an interface group named a0a:

```
network port ifgrp remove-port -node cluster-1-01 -ifgrp a0a -port e0c
```

Delete an interface group or LAG

You can delete interface groups or LAGs if you want to configure LIFs directly on the underlying physical ports or decide to change the interface group or LAG mode or distribution function.

Before you begin

- The interface group or LAG must not be hosting a LIF.
- The interface group or LAG must be neither the home port nor the failover target of a LIF.

The procedure you follow depends on the interface that you use—System Manager or the CLI:

> **System Manager**
>
> **Use System Manager to delete a LAG**
>
> **Steps**
>
> 1. Select **Network > Ethernet port > LAG** to delete a LAG.
> 2. Select the LAG you want to remove.
> 3. Delete the LAG.
>
> **CLI**
>
> **Use the CLI to delete an interface group**
>
> **Step**
>
> Use the `network port ifgrp delete` command to delete an interface group.
>
> Learn more about `network port ifgrp delete` in the ONTAP command reference.
>
> The following example shows how to delete an interface group named a0b:
>
> ```
> network port ifgrp delete -node cluster-1-01 -ifgrp a0b
> ```

Configure ONTAP VLANS over physical ports

You can use VLANs in ONTAP to provide logical segmentation of networks by creating separate broadcast domains that are defined on a switch port basis as opposed to the traditional broadcast domains, defined on physical boundaries.

A VLAN can span multiple physical network segments. The end-stations belonging to a VLAN are related by function or application.

For example, end-stations in a VLAN might be grouped by departments, such as engineering and accounting, or by projects, such as release1 and release2. Because physical proximity of the end- stations is not essential in a VLAN, you can disperse the end-stations geographically and still contain the broadcast domain in a switched network.

In ONTAP 9.14.1 and 9.13.1, untagged ports that are unused by any Logical Interfaces (LIFs) and lack native VLAN connectivity on the connected switch are marked as degraded. This is to help identify unused ports and does not indicate an outage. Native VLANs allow untagged traffic on the ifgrp base port, such as ONTAP CFM broadcasts. Configure native VLANs on the switch to prevent blocking untagged traffic.

You can manage VLANs by creating, deleting, or displaying information about them.

> ⓘ You should not create a VLAN on a network interface with the same identifier as the native VLAN of the switch. For example, if the network interface e0b is on native VLAN 10, you should not create a VLAN e0b-10 on that interface.

Create a VLAN

You can create a VLAN for maintaining separate broadcast domains within the same network domain by using System Manager or the `network port vlan create` command.

Before you begin

Confirm that the following requirements have been met:

- The switches deployed in the network must either comply with IEEE 802.1Q standards or have a vendor-specific implementation of VLANs.
- For supporting multiple VLANs, an end-station must be statically configured to belong to one or more VLANs.
- The VLAN is not attached to a port hosting a cluster LIF.
- The VLAN is not attached to ports assigned to the Cluster IPspace.
- The VLAN is not created on an interface group port that contains no member ports.

About this task

Creating a VLAN attaches the VLAN to the network port on a specified node in a cluster.

When you configure a VLAN over a port for the first time, the port might go down, resulting in a temporary disconnection of the network. Subsequent VLAN additions to the same port do not affect the port state.

> (i) You should not create a VLAN on a network interface with the same identifier as the native VLAN of the switch. For example, if the network interface e0b is on native VLAN 10, you should not create a VLAN e0b-10 on that interface.

The procedure you follow depends on the interface that you use—System Manager or the CLI:

System Manager

Use System Manager to create a VLAN

Beginning with ONTAP 9.12.0, you can automatically select the broadcast domain or manually select on from the list. Previously, broadcast domains were always automatically selected based on layer 2 connectivity. If you manually select a broadcast domain, a warning appears indicating that manually selecting a broadcast domain could result in loss of connectivity.

Steps

1. Select **Network > Ethernet port > + VLAN**.

2. Select the node from the drop-down list.

3. Choose from the following:

    a. ONTAP to **Automatically select broadcast domain (recommended)**.

    b. To manually select a broadcast domain from the list.

4. Select the ports to form the VLAN.

5. Specify the VLAN ID.

6. Save your changes.

CLI

Use the CLI to create a VLAN

In certain circumstances, if you want to create the VLAN port on a degraded port without correcting the hardware issue or any software misconfiguration, then you can set the `-ignore-health-status` parameter of the `network port modify` command as `true`.

Learn more about `network port modify` in the ONTAP command reference.

Steps

1. Use the `network port vlan create` command to create a VLAN.

2. You must specify either the `vlan-name` or the `port` and `vlan-id` options when creating a VLAN. The VLAN name is a combination of the name of the port (or interface group) and the network switch VLAN identifier, with a hyphen in between. For example, `e0c-24` and `e1c-80` are valid VLAN names.

The following example shows how to create a VLAN `e1c-80` attached to network port `e1c` on the node `cluster-1-01`:

```
network port vlan create -node cluster-1-01 -vlan-name e1c-80
```

Beginning with ONTAP 9.8, VLANs are automatically placed into appropriate broadcast domains about one minute after their creation. If you do not want ONTAP to do this, and prefer to manually place the VLAN into a broadcast domain, then specify the `-skip-broadcast-domain-placement` parameter as part of the `vlan create` command.

Learn more about `network port vlan create` in the ONTAP command reference.

Edit a VLAN

You can change the broadcast domain or disable a VLAN.

Use System Manager to edit a VLAN

Beginning with ONTAP 9.12.0, you can automatically select the broadcast domain or manually select on from the list. Previously broadcast domains were always automatically selected based on layer 2 connectivity. If you manually select a broadcast domain, a warning appears indicating that manually selecting a broadcast domain could result in loss of connectivity.

Steps

1. Select **Network > Ethernet port > VLAN**.
2. Select the edit icon.
3. Do one of the following:
   - Change the broadcast domain by selecting a different one from the list.
   - Clear the **Enabled** check box.
4. Save your changes.

Delete a VLAN

You might have to delete a VLAN before removing a NIC from its slot. When you delete a VLAN, it is automatically removed from all of the failover rules and groups that use it.

Before you begin

Make sure there are no LIFs associated with the VLAN.

About this task

Deletion of the last VLAN from a port might cause a temporary disconnection of the network from the port.

The procedure you follow depends on the interface that you use—System Manager or the CLI:

System Manager
Use System Manager to delete a VLAN

Steps

1. Select **Network > Ethernet port > VLAN**.

2. Select the VLAN you want to remove.

3. Click **Delete**.

CLI
Use the CLI to delete a VLAN

Step

Use the `network port vlan delete` command to delete a VLAN.

The following example shows how to delete VLAN `e1c-80` from network port `e1c` on the node `cluster-1-01`:

```
network port vlan delete -node cluster-1-01 -vlan-name e1c-80
```

Learn more about `network port vlan delete` in the ONTAP command reference.

Modify ONTAP network port attributes

You can modify the autonegotiation, duplex, flow control, speed, and health settings of a physical network port.

Before you begin
The port that you want to modify must not be hosting any LIFs.

About this task

*   It is not recommended to modify the administrative settings of the 100 GbE, 40 GbE, 10 GbE or 1 GbE network interfaces.

    The values that you set for duplex mode and port speed are referred to as administrative settings. Depending on network limitations, the administrative settings can differ from the operational settings (that is, the duplex mode and speed that the port actually uses).

*   It is not recommended to modify the administrative settings of the underlying physical ports in an interface group.

    The `-up-admin` parameter (available at the advanced privilege level) modifies the administrative settings of the port.

*   It is not recommended to set the `-up-admin` administrative setting to false for all ports on a node, or for the port that hosts the last operational cluster LIF on a node.

*   It is not recommended to modify the MTU size of the management port, `e0M`.

- The MTU size of a port in a broadcast domain cannot be changed from the MTU value that is set for the broadcast domain.
- The MTU size of a VLAN cannot exceed the value of the MTU size of its base port.

Steps

1. Modify the attributes of a network port:

   ```
   network port modify
   ```

2. You can set the `-ignore-health-status` field to true for specifying that the system can ignore the network port health status of a specified port.

   The network port health status is automatically changed from degraded to healthy, and this port can now be used for hosting LIFs. You should set the flow control of cluster ports to `none`. By default, the flow control is set to `full`.

The following command disables the flow control on port e0b by setting the flow control to none:

```
network port modify -node cluster-1-01 -port e0b -flowcontrol-admin none
```

Learn more about `network port modify` in the ONTAP command reference.

Create 10GbE ports for ONTAP networks by converting 40GbE NIC ports

You can convert the X1144A-R6 and the X91440A-R6 40GbE Network Interface Cards (NICs) to support four 10GbE ports.

If you are connecting a hardware platform that supports one of these NICs to a cluster that supports 10GbE cluster interconnect and customer data connections, the NIC must be converted to provide the necessary 10GbE connections.

Before you begin

You must be using a supported breakout cable.

About this task

For a complete list of platforms that support NICs, see the Hardware Universe.

ⓘ On the X1144A-R6 NIC, only port A can be converted to support the four 10GbE connections. Once port A is converted, port e is not available for use.

Steps

1. Enter maintenance mode.

2. Convert the NIC from 40GbE support to 10GbE support.

   ```
   nicadmin convert -m [40G | 10G] [port-name]
   ```

3. After using the convert command, halt the node.

4. Install or change the cable.

5. Depending on the hardware model, use the SP (Service Processor) or BMC (Baseboard Management Controller) to power-cycle the node for the conversion to take effect.

Configure UTA X1143A-R6 ports for the ONTAP network

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

When configured for Ethernet and FCoE, X1143A-R6 adapters support concurrent NIC and FCoE target traffic on the same 10-GBE port. When configured for FC, each two-port pair that shares the same ASIC can be individually configured for FC target or FC initiator mode. This means that a single X1143A-R6 adapter can support FC target mode on one two-port pair and FC initiator mode on another two-port pair. Port pairs connected to the same ASIC must be configured in the same mode.

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

To configure the unified target adapter (X1143A-R6), you must configure the two adjacent ports on the same chip in the same personality mode.

Steps

1. View the port configuration:

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

2. Configure the ports as needed for Fibre Channel (FC) or Converged Network Adapter (CNA):

```
system node hardware unified-connect modify -node <node_name> -adapter
<adapter_name> -mode {fcp|cna}
```

3. Attach the appropriate cables for FC or 10 Gb Ethernet.

4. Verify that you have the correct SFP+ installed:

```
network fcp adapter show -instance -node -adapter
```

For CNA, you should use a 10Gb Ethernet SFP. For FC, you should either use an 8 Gb SFP or a 16 Gb SFP, based on the FC fabric being connected to.

Convert the UTA2 port for use in the ONTAP network

You can convert your UTA2 port from Converged Network Adapter (CNA) mode to Fibre

Channel (FC) mode, or vice versa.

You should change the UTA2 personality from CNA mode to FC mode when you need to change the physical medium that connects the port to its network or to support the FC initiators and target.

From CNA mode to FC mode

Steps

1. Take the adapter offline:

```
network fcp adapter modify -node <node_name> -adapter <adapter_name>
-status-admin down
```

2. Change the port mode:

```
ucadmin modify -node <node_name> -adapter <adapter_name> -mode fcp
```

3. Reboot the node, and then bring the adapter online:

```
network fcp adapter modify -node <node_name> -adapter <adapter_name>
-status-admin up
```

4. Notify your admin or VIF manager to delete or remove the port, as applicable:

   ◦ If the port is used as a home port of a LIF, is a member of an interface group (ifgrp), or hosts VLANs, then an admin should do the following:

     ▪ Move the LIFs, remove the port from the ifgrp, or delete the VLANs, respectively.

     ▪ Manually delete the port by running the `network port delete` command. If the `network port delete` command fails, the admin should address the errors, and then run the command again.

   ◦ If the port is not used as the home port of a LIF, is not a member of an ifgrp, and does not host VLANs, then the VIF manager should remove the port from its records at the time of reboot. If the VIF manager does not remove the port, then the admin must remove it manually after the reboot by using the `network port delete` command.

     Learn more about `network port delete` in the ONTAP command reference.

5. Verify that you have the correct SFP+ installed:

```
network fcp adapter show -instance -node -adapter
```

For CNA, you should use a 10Gb Ethernet SFP. For FC, you should either use an 8 Gb SFP or a 16 Gb SFP, before changing the configuration on the node.

From FC mode to CNA mode

Steps

1. Take the adapter offline:

```
network fcp adapter modify -node <node_name> -adapter <adapter_name>
-status-admin down
```

2. Change the port mode:

```
ucadmin modify -node <node_name> -adapter <adapter_name> -mode cna
```

3. Reboot the node

4. Verify that you have the correct SFP+ installed.

   For CNA, you should use a 10Gb Ethernet SFP.

Convert the CNA/UTA2 optical modules for the ONTAP network

You should change the optical modules on the unified target adapter (CNA/UTA2) to support the personality mode you have selected for the adapter.

Steps

1. Verify the current SFP+ used in the card. Then, replace the current SFP+ with the appropriate SFP+ for the preferred personality (FC or CNA).

2. Remove the current optical modules from the X1143A-R6 adapter.

3. Insert the correct modules for your preferred personality mode (FC or CNA) optics.

4. Verify that you have the correct SFP+ installed:

```
network fcp adapter show -instance -node -adapter
```

Supported SFP+ modules and Cisco-branded Copper (Twinax) cables are listed in the NetApp Hardware Universe.

Remove NICs from ONTAP cluster nodes

You might have to remove a faulty NIC from its slot or move the NIC to another slot for maintenance purposes.

> (i) The procedure for removing a NIC is different in ONTAP 9.7 and earlier versions. If you need to remove a NIC from an ONTAP cluster node running ONTAP 9.7 and earlier, refer to the procedure Removing a NIC from the node (ONTAP 9.7 or earlier).

Steps

1. Power down the node.

2. Physically remove the NIC from its slot.

3. Power on the node.

4. Verify that the port has been deleted:

```
network port show
```

> (i) ONTAP automatically removes the port from any interface groups. If the port was the only member of an interface group, the interface group is deleted. Learn more about `network port show` in the [ONTAP command reference](#).

5. If the port had any VLANs configured on it, they are displaced. You can view displaced VLANs using the following command:

```
cluster controller-replacement network displaced-vlans show
```

> (i) The `displaced-interface show`, `displaced-vlans show`, and `displaced-vlans restore` commands are unique and do not require the fully qualified command name, which starts with `cluster controller-replacement network`.

6. These VLANs are deleted, but can be restored using the following command:

```
displaced-vlans restore
```

7. If the port had any LIFs configured on it, ONTAP automatically chooses new home ports for those LIFs on another port in the same broadcast domain. If no suitable home port is found on the same filer, those LIFs are considered displaced. You can view displaced LIFs using the following command:

```
displaced-interface show
```

8. When a new port is added to the broadcast domain on the same node, the home ports for the LIFs are automatically restored. Alternatively, you can either set the home port using `network interface modify -home-port -home-node` or use the `displaced- interface restore` command.

Related information

- [cluster controller-replacement network displaced-interface delete](#)
- [network interface modify](#)

Monitor network ports

Monitor the health of ONTAP network ports

ONTAP management of network ports includes automatic health monitoring and a set of health monitors to help you identify network ports that might not be suitable for hosting LIFs.

About this task

If a health monitor determines that a network port is unhealthy, it warns administrators through an EMS message or marks the port as degraded. ONTAP avoids hosting LIFs on degraded network ports if there are healthy alternative failover targets for that LIF. A port can become degraded because of a soft failure event, such as link flapping (links bouncing quickly between up and down) or network partitioning:

- Network ports in the cluster IPspace are marked as degraded when they experience link flapping or loss of layer 2 (L2) reachability to other network ports in the broadcast domain.
- Network ports in non-cluster IPspaces are marked as degraded when they experience link flapping.

You must be aware of the following behaviors of a degraded port:

- A degraded port cannot be included in a VLAN or an interface group.

  If a member port of an interface group is marked as degraded, but the interface group is still marked as healthy, LIFs can be hosted on that interface group.

- LIFs are automatically migrated from degraded ports to healthy ports.
- During a failover event, a degraded port is not considered as the failover target. If no healthy ports are available, degraded ports host LIFs according to the normal failover policy.
- You cannot create, migrate, or revert a LIF to a degraded port.

  You can modify the `ignore-health-status` setting of the network port to `true`. You can then host a LIF on the healthy ports.

Steps

1. Log in to the advanced privilege mode:

   ```
   set -privilege advanced
   ```

2. Check which health monitors are enabled for monitoring network port health:

   ```
   network options port-health-monitor show
   ```

   The health status of a port is determined by the value of health monitors.

   The following health monitors are available and enabled by default in ONTAP:

   - Link-flapping health monitor: Monitors link flapping

     If a port has link flapping more than once in five minutes, this port is marked as degraded.

   - L2 reachability health monitor: Monitors whether all ports configured in the same broadcast domain have L2 reachability to each other

     This health monitor reports L2 reachability issues in all IPspaces; however, it marks only the ports in the cluster IPspace as degraded.

   - CRC monitor: Monitors the CRC statistics on the ports

This health monitor does not mark a port as degraded but generates an EMS message when a very high CRC failure rate is observed.

Learn more about `network options port-health-monitor show` in the ONTAP command reference.

3. Enable or disable any of the health monitors for an IPspace as desired by using the `network options port-health-monitor modify` command.

Learn more about `network options port-health-monitor modify` in the ONTAP command reference.

4. View the detailed health of a port:

```
network port show -health
```

The command output displays the health status of the port, `ignore health status` setting, and list of reasons the port is marked as degraded.

A port health status can be `healthy` or `degraded`.

If the `ignore health status` setting is `true`, it indicates that the port health status has been modified from `degraded` to `healthy` by the administrator.

If the `ignore health status` setting is `false`, the port health status is determined automatically by the system.

Learn more about `network port show` in the ONTAP command reference.

Monitor the reachability of ONTAP network ports

Reachability monitoring is built into ONTAP 9.8 and later. Use this monitoring to identify when the physical network topology does not match the ONTAP configuration. In some cases, ONTAP can repair port reachability. In other cases, additional steps are required.

About this task

Use these commands to verify, diagnose, and repair network misconfigurations that stem from the ONTAP configuration not matching either the physical cabling or the network switch configuration.
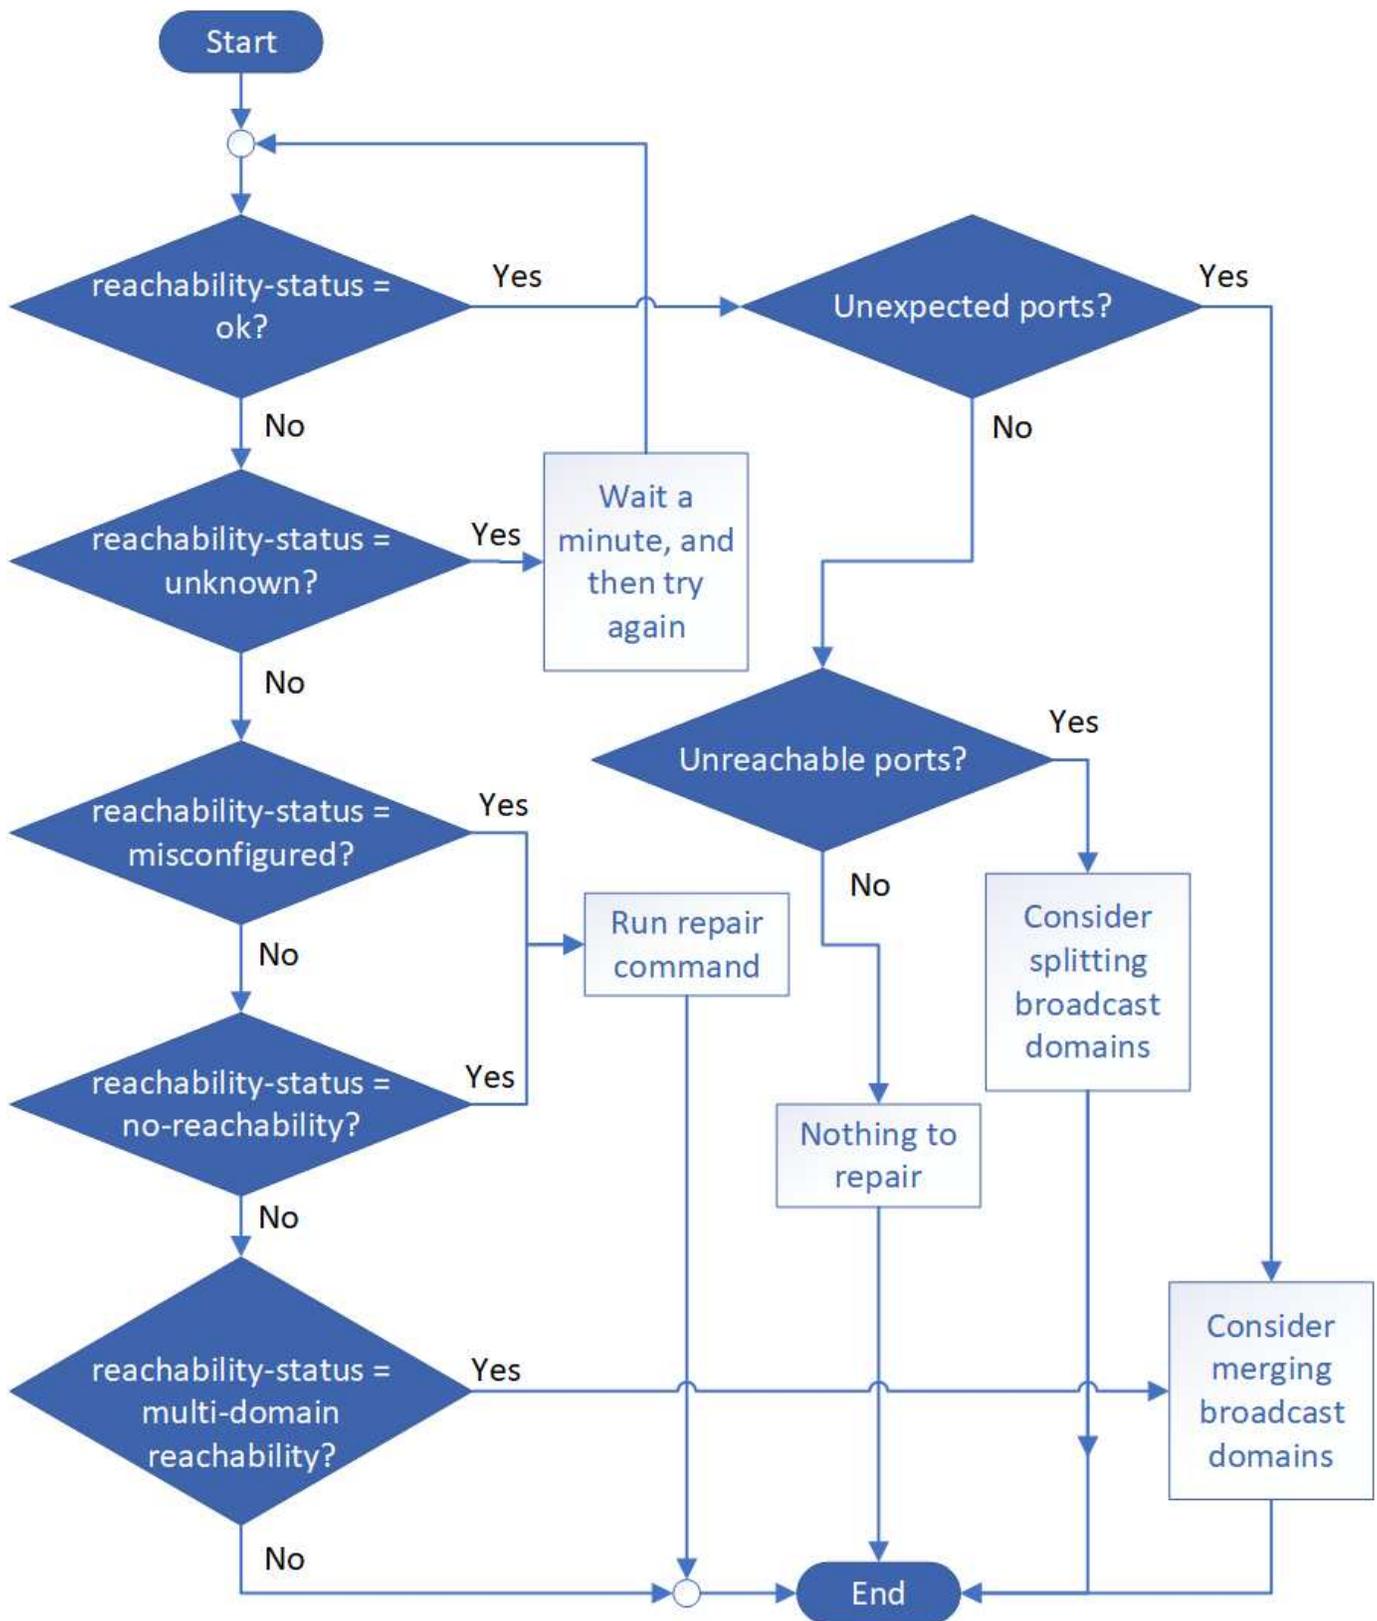
Step

1. View port reachability:

```
network port reachability show
```

Learn more about `network port reachability show` in the ONTAP command reference.

2. Use the following decision tree and table to determine the next step, if any.

Reachability-status	Description

ok	The port has layer 2 reachability to its assigned broadcast domain. If the reachability-status is "ok", but there are "unexpected ports", consider merging one or more broadcast domains. For more information, see the following *Unexpected ports* row. If the reachability-status is "ok", but there are "unreachable ports", consider splitting one or more broadcast domains. For more information, see the following *Unreachable ports* row. If the reachability-status is "ok", and there are no unexpected or unreachable ports, your configuration is correct.
Unexpected ports	The port has layer 2 reachability to its assigned broadcast domain; however, it also has layer 2 reachability to at least one other broadcast domain. Examine the physical connectivity and switch configuration to determine if it is incorrect or if the port's assigned broadcast domain needs to be merged with one or more broadcast domains. For more information, see Merge broadcast domains.
Unreachable ports	If a single broadcast domain has become partitioned into two different reachability sets, you can split a broadcast domain to synchronize the ONTAP configuration with the physical network topology. Typically, the list of unreachable ports defines the set of ports that should be split into another broadcast domain after you have verified that the physical and switch configuration is accurate. For more information, see Split broadcast domains.
misconfigured-reachability	The port does not have layer 2 reachability to its assigned broadcast domain; however, the port does have layer 2 reachability to a different broadcast domain. You can repair the port reachability. When you run the following command, the system will assign the port to the broadcast domain to which it has reachability: `network port reachability repair -node -port` For more information, see Repair port reachability.
no-reachability	The port does not have layer 2 reachability to any existing broadcast domain. You can repair the port reachability. When you run the following command, the system will assign the port to a new automatically created broadcast domain in the Default IPspace: `network port reachability repair -node -port` For more information, see Repair port reachability. Learn more about `network port reachability repair` in the ONTAP command reference.

multi-domain-reachability	The port has layer 2 reachability to its assigned broadcast domain; however, it also has layer 2 reachability to at least one other broadcast domain. Examine the physical connectivity and switch configuration to determine if it is incorrect or if the port's assigned broadcast domain needs to be merged with one or more broadcast domains. For more information, see Merge broadcast domains or Repair port reachability.
unknown	If the reachability-status is "unknown", then wait a few minutes and try the command again.

After you repair a port, you need to check for and resolve displaced LIFs and VLANs. If the port was part of an interface group, you also need to understand what happened to that interface group. For more information, see Repair port reachability.

Learn about port usage on the ONTAP network

Several well-known ports are reserved for ONTAP communications with specific services. Port conflicts occur if a port value in your storage network environment is the same as the value on an ONTAP port.

Inbound traffic

Inbound traffic on your ONTAP storage uses the following protocols and ports:

Protocol	Port	Purpose
All ICMP	All	Pinging the instance
TCP	22	Secure shell access to the IP address of the cluster management LIF or a node management LIF
TCP	80	Web page access to the IP address of the cluster management LIF
TCP/UDP	111	RPCBIND, remote procedure call for NFS
UDP	123	NTP, network time protocol
TCP	135	MSRPC, Microsoft remote procedure call
TCP	139	NETBIOS-SSN, NetBIOS service session for CIFS
TCP/UDP	161-162	SNMP, simple network management protocol
TCP	443	Secure web page access to the IP address of the cluster management LIF
TCP	445	MS Active Domain Services, Microsoft SMB/CIFS over TCP with NetBIOS framing
TCP/UDP	635	NFS mount to interact with a remote file system as if it were local
TCP	749	Kerberos

Protocol	Port	Purpose
UDP	953	Name daemon
TCP/UDP	2049	NFS server daemon
TCP	2050	NRV, NetApp remote volume protocol
TCP	3260	iSCSI access through the iSCSI data LIF
TCP/UDP	4045	NFS lock daemon
TCP/UDP	4046	Network status monitor for NFS
UDP	4049	NFS RPC Rquotad
UDP	4444	KRB524, Kerberos 524
UDP	5353	Multicast DNS
TCP	10000	Backup using Network Data Management Protocol (NDMP)
TCP	11104	Cluster peering, bi-directional management of intercluster communication sessions for SnapMirror
TCP	11105	Cluster peering, bi-directional SnapMirror data transfer using intercluster LIFs
SSL/TLS	30000	Accepts NDMP secure control connections between the DMA and NDMP server over secure sockets (SSL/TLS). Security scanners can report a vulnerability on port 30000.

Outbound traffic

Outbound traffic on your ONTAP storage can be set up using basic or advanced rules depending on business needs.

Basic outbound rules

All ports can be used for all outbound traffic over ICMP, TCP, and UDP protocols.

Protocol	Port	Purpose
All ICMP	All	All outbound traffic
All TCP	All	All outbound traffic
All UDP	All	All outbound traffic

Advanced outbound rules

If you need rigid rules for outbound traffic, you can use the following information to open only those ports that are required for outbound communication by ONTAP.

Active Directory

Protocol	Port	Source	Destination	Purpose

Protocol	Port	Source	Destination	Purpose
TCP	88	Node management LIF, data LIF (NFS, CIFS, iSCSI)	Active Directory forest	Kerberos V authentication
UDP	137	Node management LIF, data LIF (NFS, CIFS)	Active Directory forest	NetBIOS name service
UDP	138	Node management LIF, data LIF (NFS, CIFS)	Active Directory forest	NetBIOS datagram service
TCP	139	Node management LIF, data LIF (NFS, CIFS)	Active Directory forest	NetBIOS service session
TCP	389	Node management LIF, data LIF (NFS, CIFS)	Active Directory forest	LDAP
UDP	389	Node management LIF, data LIF (NFS, CIFS)	Active Directory forest	LDAP
TCP	445	Node management LIF, data LIF (NFS, CIFS)	Active Directory forest	Microsoft SMB/CIFS over TCP with NetBIOS framing
TCP	464	Node management LIF, data LIF (NFS, CIFS)	Active Directory forest	Change and set the Kerberos V password (SET_CHANGE)
UDP	464	Node management LIF, Data LIF (NFS, CIFS)	Active Directory forest	Kerberos key administration
TCP	749	Node management LIF, Data LIF (NFS, CIFS)	Active Directory forest	Change and set the Kerberos V password (RPCSEC_GSS)

AutoSupport

Protocol	Port	Source	Destination	Purpose
TCP	80	Node management LIF	support.netapp.com	AutoSupport (only if the transport protocol is changed from HTTPS to HTTP)

SNMP

Protocol	Port	Source	Destination	Purpose
TCP/UDP	162	Node management LIF	Monitor server	Monitoring by SNMP traps

SnapMirror

Protocol	Port	Source	Destination	Purpose
TCP	11104	Intercluster LIF	ONTAP intercluster LIFs	Management of intercluster communication sessions for SnapMirror

Other services

Protocol	Port	Source	Destination	Purpose
TCP	25	Node management LIF	Mail server	SMTP alerts, can be used for AutoSupport
UDP	53	Node management LIF and data LIF (NFS, CIFS)	DNS	DNS
UDP	67	Node management LIF	DHCP	DHCP server
UDP	68	Node management LIF	DHCP	DHCP client for first-time setup
UDP	514	Node management LIF	Syslog server	Syslog forward messages
TCP	5010	Intercluster LIF	Backup endpoint or restore endpoint	Back up and restore operations for the Backup to S3 feature
TCP	18600 to 18699	Node management LIF	Destination servers	NDMP copy

Learn about ONTAP internal ports

The following table lists the ports that ONTAP uses internally and their functions. ONTAP uses these ports for various functions, such as establishing intracluster LIF communication.

This list is not exhaustive and might vary in different environments.

Port/Protocol	Component/function
514	Syslog
900	NetApp Cluster RPC
902	NetApp Cluster RPC
904	NetApp Cluster RPC
905	NetApp Cluster RPC
910	NetApp Cluster RPC
911	NetApp Cluster RPC
913	NetApp Cluster RPC
914	NetApp Cluster RPC
915	NetApp Cluster RPC
918	NetApp Cluster RPC
920	NetApp Cluster RPC
921	NetApp Cluster RPC
924	NetApp Cluster RPC

925	NetApp Cluster RPC
927	NetApp Cluster RPC
928	NetApp Cluster RPC
929	NetApp Cluster RPC
930	Kernel services and management functions (KSMF)
931	NetApp Cluster RPC
932	NetApp Cluster RPC
933	NetApp Cluster RPC
934	NetApp Cluster RPC
935	NetApp Cluster RPC
936	NetApp Cluster RPC
937	NetApp Cluster RPC
939	NetApp Cluster RPC
940	NetApp Cluster RPC
951	NetApp Cluster RPC
954	NetApp Cluster RPC
955	NetApp Cluster RPC
956	NetApp Cluster RPC
958	NetApp Cluster RPC
961	NetApp Cluster RPC
963	NetApp Cluster RPC
964	NetApp Cluster RPC
966	NetApp Cluster RPC
967	NetApp Cluster RPC
975	Key Management Interoperability Protocol (KMIP)
982	NetApp Cluster RPC
983	NetApp Cluster RPC
5125	Alternate Control Port for disk
5133	Alternate Control Port for disk
5144	Alternate Control Port for disk
65502	Node scope SSH
65503	LIF Sharing
7700	Cluster Session Manager (CSM)
7810	NetApp Cluster RPC

7811	NetApp Cluster RPC
7812	NetApp Cluster RPC
7813	NetApp Cluster RPC
7814	NetApp Cluster RPC
7815	NetApp Cluster RPC
7816	NetApp Cluster RPC
7817	NetApp Cluster RPC
7818	NetApp Cluster RPC
7819	NetApp Cluster RPC
7820	NetApp Cluster RPC
7821	NetApp Cluster RPC
7822	NetApp Cluster RPC
7823	NetApp Cluster RPC
7824	NetApp Cluster RPC
7835-7839 and 7845-7849	TCP ports for intracluster communication
8023	Node Scope TELNET
8443	ONTAP S3 NAS port for Amazon FSx
8514	Node Scope RSH
9877	KMIP Client Port (Internal Local Host Only)
10006	TCP port for HA interconnect communication