IBM System Storage N series



# Clustered Data ONTAP 8.2 Network Management Guide

# Contents

Preface	<b>v</b> i
About this guide	
Supported features	
Websites	
Getting information, help, and service	
Before you call	
Using the documentation	
Hardware service and support.	
Firmware updates	
How to send your comments	
Understanding the network configuration	1
Networking components of a cluster	1
Network cabling guidelines	
Network configuration during setup (cluster administrators only)	
Network configuration after setup	
Operfigure in a state of the st	-
Configuring network ports (cluster administrators only)	
Types of network ports	
Network port naming conventions	
Roles for network ports.	
Default Ethernet port roles by platform	
Combining physical ports to create interface groups	
What interface groups are	
Types of interface groups	
Single-mode interface group	
Static multimode interface group	
Dynamic multimode interface group	
Load balancing in multimode interface groups	
IP address and MAC address load balancing	
Round-robin load balancing	
Port-based load balancing	
Restrictions on interface groups	
Creating an interface group	
Adding or removing a port from an interface group.	
Deleting an interface group	
Configuring VLANs over physical ports.	
How VLANs work	
How switches identify different VLANs	
Advantages of VLANs.	
How to use VLANs for tagged and untagged network traffic.	
Creating a VLAN	
Deleting a VI AN	10
Modifying network port attributes	20
Removing a NIC from the node	
Configuring IPv6 addresses	23
Supported and unsupported features of IPv6	
Enabling IPv6 on the cluster.	
Configuring LIFs (cluster administrators only).	
What LIFs are	
Roles for LIFs	
Characteristics of LIFs	

LIF limits	· · · ·		· · · · · · · · · · · · · · · · · · ·		     	· · · · · · · · · · · · · · · · · · ·		· · · · · · · · · · · · · · · · · · ·		· · · · · · · · · · · · · · · · · · ·		30 31 31 32 34 34 35 36
Configuring failover groups for LIFs (cluster	admi	nist	rato	ors	only	).		_				. 37
Scenarios that cause a LIF failover.						, ·						37
Types of failover groups		ż					ż					37
Relation between LIF roles and failover groups												38
Creating or adding a port to a failover group												38
Renaming a failover group												39
Removing a port from or deleting a failover group												39
Enabling or disabling failover of a LIF									•			40
Managing routing in a Vserver (cluster admi	nistra	tors	on	lv).								. 43
Creating a routing group				· <b>,</b> / ·								43
Deleting a routing group		•	· ·	•	•••	•••	•		•		•	44
Creating a route within a routing group												44
Deleting a static route												45
												47
Configuring nost-name resolution	• • •	• •	•	• •	•	• •	• •	·	• •	·	• •	. 47
Host-name resolution for the admin Vserver		•	• •	•	· ·	• •	•	• •	·	• •	•	47
Host-name resolution for a Vserver		•	• •	•	· ·	• •	•	• •	·	• •	•	47
Managing the hosts table (cluster administrators only).		•	• •	·	• •	• •	·	• •	•	• •	•	48
Commands for managing DNS nost name entries .		·	• •	·	• •	• •	·	• •	·	• •	·	48
Commands for managing DNS domains for host-name resolution		•	• •	·	• •	• •	•	• •	·	• •	•	48
Commands for managing DNS domain comigurations		•	• •	•	• •	• •	•	• •	•	• •	·	40
Balancing network loads to optimize user tra	affic (	clus	ter	adr	ninis	strat	ors	onl	y).			. 51
Load balancing types												51
Guidelines for assigning load balancing weights												51
Assigning a load balancing weight to a LIF												52
How DNS load balancing works												52
Creating a DNS load balancing zone												53
Adding or removing a LIF from a load balancing zone									•			53
How automatic LIF rebalancing works		•		•	• •	• •	•		•		•	55
Enabling or disabling automatic LIF rebalancing		•		•			•	• •	•		•	55
Combining load balancing methods in a Vserver accessible	le in a 1	multi	proto	ocol	envir	onme	ent	• •	•	• •	•	56
Managing SNMP on the cluster (cluster adm	inistra	ator	s or	ılv)								. 59
What MIBs are												59
Creating an SNMP community												60
Configuring SNMPv3 users in a cluster												61
SNMPv3 security parameters												62
Examples for different security levels												62
SNMP traps												64
Configuring traphosts												64
Commands for managing SNMP												65
Viewing notwork information												67
	•••	•••	•	• •	•	•••	• •	•	• •	·	• •	. 0/
Displaying network port information (cluster administrate	ors only	y).	• •	·	• •	• •	·	• •	·	• •	•	67
Displaying information about a VLAN (cluster administra	ators or	uy)	• •	·	• •	• •	·	• •	•	• •	•	68
Displaying interface group information (cluster administr	ators of	nly)	• •	•	• •	• •	·	• •	•	• •	•	69
Displaying LIF information		•	• •	·	• •	• •	·	• •	·	• •	•	69
Displaying routing information.		•	• •	·	• •	• •	•	• •	•	• •	·	71

Displaying host name entries (cluster administrators only).	. 72
Displaying DNS domain configurations	. 72
Displaying information about failover groups (cluster administrators only)	. 73
Viewing failover targets of LIFs	. 73
Viewing LIFs in a load balancing zone	. 75
Displaying cluster connections	. 76
Displaying active connections by client (cluster administrators only)	. 76
Displaying active connections by protocol (cluster administrators only)	. 77
Displaying active connections by service (cluster administrators only).	. 78
Displaying active connections by LIF on a node and Vserver	. 78
Displaying active connections in a cluster	. 79
Displaying listening connections in a cluster	. 80
Commands for diagnosing network problems	. 81
Using CDP to detect network connectivity	. 81
Considerations for using CDP	. 81
Enabling or disabling CDP	. 82
Configuring hold time for CDP messages	. 82
Setting the intervals for sending CDP advertisements	. 83
Viewing or clearing CDP statistics	. 83
Viewing neighbor information by using CDP	. 84
Copyright and trademark information	87
Trademark information	88
	. 00
Notices	89
Index	91

### Preface

### About this guide

This document applies to IBM N series systems running Data ONTAP, including systems with gateway functionality. If the terms *Cluster-Mode* or *clustered Data ONTAP* are used in this document, they refer to the Data ONTAP features and functionality designed for clusters, which are different from 7-Mode and prior Data ONTAP 7.1, 7.2, and 7.3 release families.

In this document, the term *gateway* describes IBM N series storage systems that have been ordered with gateway functionality. Gateways support various types of storage, and they are used with third-party disk storage systems—for example, disk storage systems from IBM, HP<sup>®</sup>, Hitachi Data Systems<sup>®</sup>, and EMC<sup>®</sup>. In this case, disk storage for customer data and the RAID controller functionality is provided by the back-end disk storage system. A gateway might also be used with disk storage expansion units specifically designed for the IBM N series models.

The term *filer* describes IBM N series storage systems that either contain internal disk storage or attach to disk storage expansion units specifically designed for the IBM N series storage systems. Filer storage systems do not support using third-party disk storage systems.

### Supported features

IBM System Storage N series storage systems are driven by NetApp Data ONTAP software. Some features described in the product software documentation are neither offered nor supported by IBM. Please contact your local IBM representative or reseller for further details.

Information about supported features can also be found on the N series support website (accessed and navigated as described in Websites).

### Websites

IBM maintains pages on the World Wide Web where you can get the latest technical information and download device drivers and updates. The following web pages provide N series information:

• A listing of currently available N series products and features can be found at the following web page:

www.ibm.com/storage/nas/

• The IBM System Storage N series support website requires users to register in order to obtain access to N series support content on the web. To understand how the N series support web content is organized and navigated, and to access the N series support website, refer to the following publicly accessible web page:

www.ibm.com/storage/support/nseries/

This web page also provides links to AutoSupport information as well as other important N series product resources.

• IBM System Storage N series products attach to a variety of servers and operating systems. To determine the latest supported attachments, go to the IBM N series interoperability matrix at the following web page:

www.ibm.com/systems/storage/network/interophome.html

• For the latest N series hardware product documentation, including planning, installation and setup, and hardware monitoring, service and diagnostics, see the IBM N series Information Center at the following web page:

publib.boulder.ibm.com/infocenter/nasinfo/nseries/index.jsp

### Getting information, help, and service

If you need help, service, or technical assistance or just want more information about IBM products, you will find a wide variety of sources available from IBM to assist you. This section contains information about where to go for additional information about IBM and IBM products, what to do if you experience a problem with your IBM N series product, and whom to call for service, if it is necessary.

#### Before you call

Before you call, make sure you have taken these steps to try to solve the problem yourself:

- Check all cables to make sure they are connected.
- Check the power switches to make sure the system is turned on.
- Use the troubleshooting information in your system documentation and use the diagnostic tools that come with your system.
- Refer to the N series support website (accessed and navigated as described in Websites) for information on known problems and limitations.

#### Using the documentation

The latest versions of N series software documentation, including Data ONTAP and other software products, are available on the N series support website (accessed and navigated as described in Websites).

Current N series hardware product documentation is shipped with your hardware product in printed documents or as PDF files on a documentation CD. For the latest N series hardware product documentation PDFs, go to the N series support website.

Hardware documentation, including planning, installation and setup, and hardware monitoring, service, and diagnostics, is also provided in an IBM N series Information Center at the following web page:

publib.boulder.ibm.com/infocenter/nasinfo/nseries/index.jsp

#### Hardware service and support

You can receive hardware service through IBM Integrated Technology Services. Visit the following web page for support telephone numbers:

www.ibm.com/planetwide/

#### Firmware updates

IBM N series product firmware is embedded in Data ONTAP. As with all devices, ensure that you run the latest level of firmware. Any firmware updates are posted to the N series support website (accessed and navigated as described in Websites).

**Note:** If you do not see new firmware updates on the N series support website, you are running the latest level of firmware.

Verify that the latest level of firmware is installed on your machine before contacting IBM for technical support.

# How to send your comments

Your feedback helps us to provide the most accurate and high-quality information. If you have comments or suggestions for improving this document, please send them by email to starpubs@us.ibm.com.

Be sure to include the following:

- Exact publication title
- Publication form number (for example, GC26-1234-02)
- Page, table, or illustration numbers
- A detailed description of any information that should be changed

## Understanding the network configuration

You need to understand how to configure networking components of the cluster during and after the setting up the cluster. You should follow some guidelines when cabling the nodes and switches in your network. After cabling and installing clustered Data ONTAP on the nodes, you need to set up the cluster and configure at least one functional Vserver.

### Networking components of a cluster

You should familiarize yourself with the networking components of a cluster, before or after setting up a cluster. Abstracting and virtualizing physical networking components of a cluster into logical components provides the flexibility and potential multi-tenancy in Data ONTAP.

The various networking components in a cluster are:

- Ports
  - Physical ports: Network interface cards (NICs) and HBAs provide physical (Ethernet and Fibre Channel) connections to the physical networks (management and data networks).
  - Virtual ports: VLANs and interface groups (ifgrps) constitute the virtual ports.
     While interface groups treat several physical ports as a single port, VLANs subdivide a physical port into multiple separate ports.
- Logical interfaces

A logical interface (LIF) is an IP address and is associated with attributes such as failover rule lists, firewall rules. A LIF communicates over the network through the port (physical or virtual) it is currently bound to.

Different types of LIFs in a cluster are data LIFs, cluster-management LIFs, node-management LIFs, intercluster LIFs, and cluster LIFs. The ownership of the LIFs depends on the Vserver where the LIF resides. Data LIFs are owned by data Vserver, node-management and cluster LIFs are owned by node Vserver, and cluster-management LIFs are owned by admin Vserver.

Routing groups

A routing group is a routing table. Each LIF is associated with a routing group and uses only the routes of that group. Multiple LIFs can share a routing group. Each routing group needs a minimum of one route to access clients outside the defined subnet.

DNS zones

DNS zone can be specified during the LIF creation, providing a name for the LIF to be exported through the cluster's DNS server. Multiple LIFs can share the same name, allowing the DNS load balancing feature to distribute IP addresses for the name according to load. A Vserver can have multiple DNS zones.

The following diagram illustrates how the different networking components are associated in a 4-node cluster:



For more information about the basic cluster concepts and Vservers, see the *Clustered Data ONTAP System Administration Guide for Cluster Administrators*.

### Network cabling guidelines

You should cable a cluster so that the cluster traffic is on a separate network from all other traffic. It is optional but a recommended practice to have network management separated from data and intercluster traffic. By maintaining separate networks, you can achieve better performance, ease of administration, and improved security and management access to the nodes.

The following diagram illustrates the network cabling of a 4-node cluster with two different networks for data and management, apart from the cluster interconnect:



**Note:** Apart from these networks, there is a separate network for ACP (Alternate Control Path) that enables Data ONTAP to manage and control a SAS disk shelf

storage subsystem. ACP uses a separate network (alternate path) from the data path. For more information about ACP, see the *Clustered Data ONTAP Physical Storage Management Guide*.

You should follow certain guidelines when cabling network connections:

• Each node should be connected to three distinct networks—one for management, one for data access, and one for intracluster communication. The management and data networks can be logically separated.

For setting up the cluster interconnect and the management network by using the supported Cisco switches, see the *Clustered Data ONTAP Switch Setup and Configuration Guide*.

For setting up the cluster interconnect and the management network by using the CN1610 switch, see the *Clustered Data ONTAP Switch Setup and Configuration Guide* 

- A cluster can be created without data network connections, but must include a cluster interconnect connection.
- There should always be two cluster connections to each node, but nodes on N3240 systems may be configured with a single 10-GbE cluster port.
- You can have more than one data network connection to each node for improving the client (data) traffic flow.

### Network configuration during setup (cluster administrators only)

You should be aware of some basic network configuration to be performed during the cluster setup. Basic network configuration includes setting up the management and cluster networks, and configuring network services for the admin Vserver. For a data Vserver, you should configure data LIFs and naming services.

You can perform the initial network configuration of the cluster and the Vserver by using the following wizards:

- Cluster Setup wizard
- Vserver Setup wizard

**Note:** The Cluster Setup and Vserver Setup wizards do not support IPv6. If you want data LIFs and management LIFs configured for IPv6, you can modify the LIFs after the cluster has been configured and is running. It is not possible to configure an IPv6-only cluster, because cluster LIFs and intercluster LIFs only support IPv4.

#### Networking configuration during the Cluster setup

You can perform the initial setup of the admin Vserver by using the Cluster Setup wizard. During the setup, you can either select the default values or customize your setup. The default values for setup are generated by using the *zero configuration networking* mechanism.

For setting up the	Types of information required
Cluster network	<ul> <li>Two ports for the private cluster network</li> <li>MTU size of 9000 for cluster ports</li> <li>IP address for the cluster interfaces</li> <li>When a new node joins a cluster, IP addresses are generated automatically. However, if you want to assign IP addresses manually to the cluster LIFs, you must ensure that the new IP addresses are in the same subnet range as the existing cluster LIFs.</li> <li>Netmask for the cluster interfaces</li> </ul>
Admin Vserver	<ul> <li>Node and port for the cluster-management LIF</li> <li>IP address for the cluster-management LIF</li> <li>Netmask for the cluster-management LIF</li> <li>Default gateway for the cluster-management LIF</li> <li>DNS domain name</li> <li>DNS name servers</li> </ul>
Node Vserver (for each node in the cluster)	<ul> <li>Port for the node-management LIF</li> <li>IP address for the node-management LIF</li> <li>Netmask for the node-management LIF</li> <li>Default gateway for the node-management LIF</li> <li>Note: If you choose not to configure a node-management LIF or specify a default gateway for the node-management LIF while using the Cluster Setup wizard, you must configure a node-management LIF on each node later by using the command-line interface. Otherwise, some operations might fail.</li> </ul>

### Network configuration during the Vserver setup

The Vserver Setup wizard guides you through the following configuration:

- Storage resources
- Data LIFs
- Naming services

For more information about the setup process using the Cluster Setup and Vserver Setup wizards, see the *Clustered Data ONTAP Software Setup Guide*.

#### Related concepts:

"Configuring host-name resolution" on page 47

#### Related reference:

"Commands for managing DNS domain configurations" on page 48

### Network configuration after setup

You can configure various logical network components after setting up the admin Vserver or data Vserver. For example, after creating the admin Vserver you can create LIFs, routing groups, failover groups, and so on.

If you are a cluster administrator, you can perform the following tasks:

- Configure network ports
- Configure LIFs
- Configure failover groups
- Configure routing groups
- Configure network services
- Balance network traffic

If you are a Vserver administrator, you can perform the following tasks:

- View LIFs
- View routing groups
- Create, modify, and manage DNS hosts table entries

# Configuring network ports (cluster administrators only)

Ports are either physical ports (NICs) or virtualized ports, such as interface groups or VLANs. A LIF communicates over the network through the port to which it is currently bound.

### Types of network ports

Ports are either physical ports (NICs), or virtualized ports such as interface groups or VLANs. Interface groups treat several physical ports as a single port, while VLANs subdivide a physical port into multiple separate virtual ports.

#### Interface group

A port aggregate containing two or more physical ports that act as a single trunk port. An interface group can be single-mode, multimode, or dynamic multimode.

VLAN A virtual port that receives and sends VLAN-tagged (IEEE 802.1Q standard) traffic. VLAN port characteristics include the VLAN ID for the port. The underlying physical port or interface group ports are considered to be VLAN trunk ports, and the connected switch ports must be configured to trunk the VLAN IDs.

**Note:** The underlying physical port or interface group ports for a VLAN port can continue to host LIFs, which transmit and receive untagged traffic.

#### Related concepts:

"Combining physical ports to create interface groups" on page 9 "Configuring VLANs over physical ports" on page 16

### Network port naming conventions

Port names consist of three characters that describe the port's type and location. You must be aware of certain conventions of naming the Ethernet ports on the network interfaces.

In Data ONTAP physical ports, the first character describes the port's type and is always *e* to represent Ethernet. The second character is a numeral identifying the slot in which the port adapter is located; the numeral 0 (zero) indicates that the port is on the node's motherboard. The third character indicates the port's position on a multiport adapter. For example, the port name e0b indicates the second Ethernet port on the motherboard, and the port name e3a indicates the first Ethernet port on an adapter in slot 3.

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

VLANs must be named by using the syntax *port\_name-vlan-id*, where *port\_name* specifies the physical port or interface group and *vlan-id* specifies the VLAN ID. For example, e1c-80 is a valid VLAN name.

### **Roles for network ports**

Network ports can have roles that define their purpose and their default behavior. Port roles limit the types of LIFs that can be bound to a port. Network ports can have four roles: node management, cluster, data, and intercluster.

Each network port has a default role. You can modify the roles for obtaining the best configuration.

#### Node management ports

The ports used by administrators to connect to and manage a node. These ports can be VLAN-tagged virtual ports where the underlying physical port is used for other traffic. The default port for node management differs depending on hardware platform.

Some platforms have a dedicated management port (e0M). The role of such a port cannot be changed, and these ports cannot be used for data traffic.

#### **Cluster ports**

The ports used for intracluster traffic only. By default, each node has two cluster ports on 10-GbE ports enabled for jumbo frames.

**Note:** In some cases, nodes on N3240 systems will be configured with a single 10-GbE cluster port.

You cannot create VLANs or interface groups on cluster ports.

#### Data ports

The ports used for data traffic. These ports are accessed by NFS, CIFS, FC, and iSCSI clients for data requests. Each node has a minimum of one data port.

You can create VLANs and interface groups on data ports. VLANs and interface groups have the data role by default, and the port role cannot be modified.

#### Intercluster ports

The ports used for cross-cluster communication. An intercluster port should be routable to another intercluster port or the data port of another cluster. Intercluster ports can be on physical ports or virtual ports.

#### **Related concepts:**

"Roles for LIFs" on page 26

#### Related reference:

"Default Ethernet port roles by platform"

#### Default Ethernet port roles by platform

During the configuration of a cluster, default roles are assigned to each network port. The network port for each role varies depending on the platform. You can modify these assignments later, depending on your needs.

You need at least two networks for cluster and node connectivity:

• A physically secure, dedicated network to connect the cluster ports on all nodes in the cluster

**Note:** The cluster ports on the nodes should be configured on a high-speed, high-bandwidth network, and the MTU should be set to 9000 bytes.

• A network to connect to the management and data ports on each node

Platform	Cluster ports	Node management port	Data ports
N3150	e0a, e0b	e0M	All other Ethernet ports are data ports
N3220 and N3240	e1a, e1b	e0M	All other Ethernet ports are data ports
N6000 series	e1a, e2a	e0M	All other Ethernet ports are data ports
N6200 series	e1a, e2a	e0M	All other Ethernet ports are data ports
N7000 series	е5а, еба	e0f	All other Ethernet ports are data ports
N7x50T series	e0c, e0e	e0M	All other Ethernet ports are data ports

For each hardware platform, the default role for each port is defined as follows:

### Combining physical ports to create interface groups

You can use interface groups to combine two or more physical ports and present them to clients as a single virtual port with higher throughput than a LIF associated with a single physical port.

### What interface groups are

An interface group is a port aggregate containing two or more physical ports that acts as a single trunk port. Expanded capabilities include increased resiliency, increased availability, and load sharing.

### Types of interface groups

You can create three different types of interface groups on your storage system: single-mode, static multimode, and dynamic multimode interface groups.

Each interface group provides different levels of fault tolerance. Multimode interface groups provide methods for load balancing network traffic.

#### Single-mode interface group

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. All interfaces in a single-mode interface group share a common MAC address.

There can be more than one interface on standby in a single-mode interface group. If an active interface fails, the cluster randomly picks one of the standby interfaces to be the next active link. The active link is monitored and link failover is controlled by the cluster; therefore, single-mode interface group does not require any switch configuration. Single-mode interface groups also do not require a switch that supports link aggregation.

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 (for example, a LAN or a

VLAN). Link-monitoring ARP packets, which have a source address of 0.0.0.0, are sent over the ports of a single-mode interface group to detect whether 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.



#### Static multimode interface group

The static multimode interface group implementation in Data 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 a few characteristics of a static multimode interface group:

• In a static multimode interface group, all interfaces in the interface group are active and share a single MAC address.

This logical aggregation of interfaces allows for multiple individual connections to be distributed among the interfaces in the interface group. Each connection or session uses one interface within the interface group and has a reduced likelihood of sharing that single interface with other connections. This effectively allows for greater aggregate throughput, although each individual connection is limited to the maximum throughput available in a single port.

When you use the round-robin 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.

For more information about this scheme, see the Round-robin load balancing.

• 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 in a static multimode interface group, the traffic that was traversing that failed link is automatically redistributed to one of the

remaining interfaces. If the failed or disconnected port is restored to service, traffic is automatically redistributed among all active interfaces, including the newly restored interface.

- 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 e0, e1, e2, and e3 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 Data 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 Data ONTAP.

The IEEE 802.3 standard states that the transmitting device in an aggregated link determines the physical interface for transmission. Therefore, Data 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 Data ONTAP complies with IEEE 802.3 AD (802.1 AX). Data 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.

Data ONTAP implements LACP in nonconfigurable active mode that works well with switches that are configured in either active or passive mode. Data 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 Data 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 re-routed 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 equally utilized for outgoing traffic by using the IP address, MAC address, round-robin, 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.

#### 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.

**Note:** Do not select the MAC address load balancing method when creating interface groups on a storage 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.

#### **Round-robin load balancing**

You can use round-robin for load balancing multimode interface groups. You should use the round-robin option for load balancing a single connection's traffic across multiple links to increase single connection throughput. However, this method might cause out-of-order packet delivery.

If the remote TCP endpoints do not handle TCP reassembly correctly or lack enough memory to store out-of-order packets, they might be forced to drop packets. Therefore, this might result in unnecessary retransmissions from the storage controller.

#### Port-based load balancing

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.

### **Restrictions on interface groups**

Interface groups have certain restrictions.

- All the ports in an interface group must be physically located on the same storage system, but do not need to be on the same network adapter in the storage system.
- There can be a maximum of 16 physical interfaces in an interface group.
- There can be a maximum of 4 physical interfaces if the interface group is made up of 10-GbE ports.
- A port that is already a member of an interface group cannot be added to another interface group.
- All ports in an interface group must have the same port role (data).
- Cluster ports and node management ports cannot be included in an interface group.
- A port to which a LIF is already bound cannot be added to an interface group.

- An interface group can be moved to the administrative up and down settings, but the administrative settings of the underlying physical ports cannot be changed.
- Interface groups cannot be created over VLANs or other interface groups.
- In static multimode and dynamic multimode (LACP) interface groups, the network ports used must have identical port characteristics. Some switches allow media types to be mixed in interface groups. However, the speed, duplex, and flow control should be identical.

The network ports should belong to network adapters of the same model. Support for hardware features such as TSO, LRO, and checksum offloading varies for different models of network adapters. If all ports do not have identical support for these hardware features, the feature might be disabled for the interface group.

**Note:** Using ports with different physical characteristics and settings can have a negative impact on multimode interface group throughput.

### Creating an interface group

You can create an interface group — single-mode, static multimode, or dynamic multimode (LACP) — to present a single interface to clients by combining the capabilities of the aggregated network ports. Interface groups cannot be created from other interface groups or VLANs.

#### About this task

- In a single-mode interface group, you can select the active port or designate a port as nonfavored by executing the **ifgrp** command from the nodeshell.
- While creating a multimode interface group, you can specify any of the following load balancing methods:
  - 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.
  - port: Network traffic is distributed on the basis of the transport layer (TCP/UDP) ports.

#### Procedure

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. For more information about this command, see the man pages. The following example shows how to create an interface group named a0a with a distribution function of ip and a mode of multimode:

cluster1::> network port ifgrp create -node cluster1-01 -ifgrp a0a -distr-func ip -mode multimode

### Adding or removing a port from an interface group

You can add a port to an interface group after creating the initial interface group. You can also remove a port from an interface group.

### Before you begin

To remove a port from an interface group, it must not be hosting any LIFs.

### About this task

You can add up to 16 ports (physical interfaces) to an interface group.

#### Procedure

Depending on whether you want to add or remove network ports from an interface group, enter the following command:

If you want to	Then, enter the following command		
Add network ports to an interface group	network port ifgrp add-port		
Remove network ports from an interface	network port ifgrp remove-port		
group			

For more information about these commands, see the man pages. The following example shows how to add ports e0c to an interface group named a0a:

cluster1::> network port ifgrp add-port -node cluster1-01 -ifgrp a0a -port e0c

The following example shows how to remove port e0d from an interface group named a0a:

cluster1::> network port ifgrp remove-port -node cluster1-01 -ifgrp a0a -port e0d

### Deleting an interface group

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

#### Before you begin

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

#### Procedure

Use the **network port ifgrp delete** command to delete an interface group. For more information about this command, see the man pages. The following example shows how to delete an interface group named a0b:

cluster1::> network port ifgrp delete -node cluster1-01 -ifgrp a0b

#### Related tasks:

"Modifying network port attributes" on page 20

"Displaying LIF information" on page 69

### Configuring VLANs over physical ports

VLANs provide logical segmentation of networks by creating separate broadcast domains. 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.

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

### How VLANs work

Traffic from multiple VLANs can traverse a link that interconnects two switches by using VLAN tagging. A VLAN tag is a unique identifier that indicates the VLAN to which a frame belongs. A VLAN tag is included in the header of every frame sent by an end-station on a VLAN.

On receiving a tagged frame, the switch inspects the frame header and, based on the VLAN tag, identifies the VLAN. The switch then forwards the frame to the destination in the identified VLAN. If the destination MAC address is unknown, the switch limits the flooding of the frame to ports that belong to the identified VLAN.



For example, in this figure, if a member of VLAN 10 on Floor 1 sends a frame for a member of VLAN 10 on Floor 2, Switch 1 inspects the frame header for the VLAN tag (to determine the VLAN) and the destination MAC address. The

destination MAC address is not known to Switch 1. Therefore, the switch forwards the frame to all other ports that belong to VLAN 10, that is, port 4 of Switch 2 and Switch 3. Similarly, Switch 2 and Switch 3 inspect the frame header. If the destination MAC address on VLAN 10 is known to either switch, that switch forwards the frame to the destination. The end-station on Floor 2 then receives the frame.

### How switches identify different VLANs

A network switch distinguishes between VLANs by associating end-stations to a specific VLAN. This is known as VLAN membership. An end-station must become a member of a VLAN before it can share the broadcast domain with other end-stations on that VLAN.

VLAN membership can be based on one of the following:

- Switch ports
- End-station MAC addresses
- Protocol

In Data ONTAP, VLAN membership is based on switch ports. With port-based VLANs, ports on the same or different switches can be grouped to create a VLAN. As a result, multiple VLANs can exist on a single switch. The switch ports can be configured to belong to one or more VLANs (static registration).

Any broadcast or multicast packets originating from a member of a VLAN are confined only among the members of that VLAN. Communication between VLANs, therefore, must go through a router.

The following figure illustrates how communication occurs between geographically dispersed VLAN members:



In this figure, VLAN 10 (Engineering), VLAN 20 (Marketing), and VLAN 30 (Finance) span three floors of a building. If a member of VLAN 10 on Floor 1 wants to communicate with a member of VLAN 10 on Floor 3, the communication occurs without going through the router, and packet flooding is limited to port 1 of Switch 2 and Switch 3 even if the destination MAC address to Switch 2 and Switch 3 is not known.

### Advantages of VLANs

VLANs provide a number of advantages, such as ease of administration, confinement of broadcast domains, reduced broadcast traffic, and enforcement of security policies.

VLANs provide the following advantages:

• VLANs enable logical grouping of end-stations that are physically dispersed on a network.

When users on a VLAN move to a new physical location but continue to perform the same job function, the end-stations of those users do not need to be reconfigured. Similarly, if users change their job functions, they need not physically move: changing the VLAN membership of the end-stations to that of the new team makes the users' end-stations local to the resources of the new team.

• VLANs reduce the need to have routers deployed on a network to contain broadcast traffic.

Flooding of a packet is limited to the switch ports that belong to a VLAN.

• Confinement of broadcast domains on a network significantly reduces traffic.

By confining the broadcast domains, end-stations on a VLAN are prevented from listening to or receiving broadcasts not intended for them. Moreover, if a router is not connected between the VLANs, the end-stations of a VLAN cannot communicate with the end-stations of the other VLANs.

### How to use VLANs for tagged and untagged network traffic

You can configure an IP address for an interface with VLANs. Any untagged traffic goes to the base interface and the tagged traffic goes to the respective VLAN.

You can configure an IP address for the base interface (physical port) of the VLAN. Any tagged frame is received by the matching VLAN interface. Untagged traffic is received by the native VLAN on the base interface.

**Note:** 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.

You cannot bring down the base interface that is configured to receive tagged and untagged traffic. You must bring down all VLANs on the base interface before you bring down the interface. However, you can delete the IP address of the base interface.

### **Creating a VLAN**

You can create a VLAN for maintaining separate broadcast domains within the same network domain by using the **network port vlan create** command. You cannot create a VLAN from an existing VLAN.

#### Before you begin

You must contact your network administrator to check if the following requirements are met:

- The switches deployed in the network either comply with IEEE 802.1Q standards or have vendor-specific implementation of VLANs.
- For supporting multiple VLANs, an end-station is statically configured to belong to one or more VLANs.

#### About this task

You cannot create a VLAN on cluster-management and node-management ports.

#### Procedure

Use the **network port vlan create** command to create a VLAN. For more information about this command, see the man pages. You must specify either the vlan-name or the port and vlan-id options when creating a VLAN.

**Note:** You cannot attach a VLAN to a cluster port. The following example shows how to create a VLAN e1c-80 attached to network port e1c on the node cluster1-01:

```
cluster1::> network port vlan create -node cluster1-01 -vlan-name e1c-80
```

### **Deleting 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 failover rules and groups that use it.

### Before you begin

Ensure that there are no LIFs associated with the VLAN.

#### About this task

Before removing a NIC from its slot, you have to delete all the physical ports and their associated VLANs.

#### Procedure

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 cluster1-01:

cluster1::> network port vlan delete -node cluster1-01 -vlan-name e1c-80

#### Related tasks:

"Displaying LIF information" on page 69

### Modifying network port attributes

You can modify the MTU, autonegotiation, duplex, flow control, and speed settings of a physical network or interface group. You can modify only the MTU settings and not other port settings of a VLAN.

#### Before you begin

The port to be modified must not be hosting any LIFs.

#### About this task

You should not modify the following characteristics of a network port:

- The administrative settings of either the 10-GbE or the 1-GbE network interfaces. The values that you can 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).
- The administrative settings of the underlying physical ports in an interface group.

**Note:** Use the -up-admin parameter (available at advanced privilege level) to modify the administrative settings of the port.

- The MTU size of the management port, e0M.
- The MTU size of a VLAN cannot exceed the value of the MTU size of its base port.

#### Procedure

Use the **network port modify** command to modify the attributes of a network port.

**Note:** You should set the flow control of all ports to **none**. By default, the flow control is set to **full**.

The following example shows how to disable the flow control on port e0b by setting it to **none**:

cluster1::> network port modify -node cluster1-01 -port e0b -flowcontrol-admin none

### Removing a NIC from the node

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

#### Before you begin

- All the LIFs hosted on the NIC ports must have been migrated or deleted.
- All the NICs ports must not be the home ports of any LIFs.
- You must have advanced privileges to delete the ports from a NIC.

#### Procedure

- 1. Use the **network port delete** command to delete the ports from the NIC. For more information about removing a NIC, see the *Moving or replacing a NIC in Data ONTAP 8.1 operating in Cluster-Mode* document.
- 2. Use the **network port show** to verify that the ports have been deleted.
- **3**. Repeat step 1, if the output of the **network port show** command still shows the deleted port.

IBM N series support website: www.ibm.com/storage/support/nseries

# **Configuring IPv6 addresses**

IPv6 increases the IP address size from 32 bits (in IPv4) to 128 bits. This larger address space provides expanded routing and addressing capabilities. Starting from clustered Data ONTAP 8.2, you can create LIFs with IPv6 addresses.

The following are some of the advantages of the IPv6 protocol:

- Large address header
- · Address auto-configuration
- Neighbor Discovery
- Path MTU discovery
- Built-in security

Although most of the IPv6 features have been implemented in clustered Data ONTAP 8.2, you should familiarize yourself with the unsupported features of IPv6 as well. You can enable IPv6 on the cluster before configuring various networking components with IPv6 addresses.

For detailed explanations about various IPv6 address states, address auto-configuration, and the neighbor discovery features of IPv6, see the relevant RFCs.

#### **Related information:**

- □ IPv6 addressing architecture (RFC4291)
- ➡ Neighbour Discovery for IP version 6 (RFC4861)
- IPv6 Stateless Address Configuration (RFC4862)

### Supported and unsupported features of IPv6

Starting from clustered Data ONTAP 8.2, you can create LIFs with IPv6 addresses. Although most of the functionality of IPv6 addressing is supported, some of the key features of IPv6, such as address auto-configuration, are not supported.

The following are the supported features of IPv6:

- Simultaneous support for IPv4 and IPv6
- Network administration commands, such as the traceroute6, ping6, and pktt commands (available from the nodeshell)
- File access protocols—CIFS, SMB2.x, SMB3.0, HTTP, NFSv3, NFSv4, and NFSv4.1
- SNMP access over IPv6
- SSH, RSH and Telnet over IPv6
- RLM IPv6 manual and auto-configured addresses
- Dump, restore, NDMP, and ndmpcopy operations over IPv6

The following are the unsupported features of IPv6:

- Address auto-configuration such as SLAAC and DHCPv6
- · Manual configuration of link-local addresses
- · Configuring cluster LIFs and intercluster LIFs with IPv6 addresses
- Fast path

- DNS load balancing
- Cluster setup and Vserver setup wizards
- MIB for TCP, UDP, ICMPv6, and IPv6

### Enabling IPv6 on the cluster

Starting from clustered Data ONTAP 8.2, you can enable IPv6 on the cluster. By enabling IPv6 on the cluster, you can manage various networking objects such as LIFs, routing groups, routes, and firewall policies.

#### Before you begin

All the nodes in the cluster must be running clustered Data ONTAP 8.2.

#### About this task

You cannot disable IPv6 after enabling it on the cluster.

However, if you want to disable IPv6, contact technical support for guidance.

#### Procedure

- 1. Use the network options ipv6 modify command to enable IPv6 on the cluster.
- 2. Use the **network options ipv6 show** command to verify that IPv6 is enabled in the cluster.

# Configuring LIFs (cluster administrators only)

A LIF represents a network access point to a node in the cluster. You can configure LIFs on ports over which the cluster sends and receives communications over the network.

A cluster administrator can create, view, modify, migrate, or delete LIFs. A Vserver administrator can only view the LIFs associated with the Vserver.

### What LIFs are

A LIF (logical interface) is an IP address with associated characteristics, such as a role, a home port, a home node, a routing group, a list of ports to fail over to, and a firewall policy. You can configure LIFs on ports over which the cluster sends and receives communications over the network.

LIFs can be hosted on the following ports:

- Physical ports that are not part of interface groups
- Interface groups
- VLANs
- Physical ports or interface groups that host VLANs

While configuring SAN protocols such as FC on a LIF, it will be associated with a WWPN.

For more information about configuring WWPN to LIFs while using the FC protocol, see the *Clustered Data ONTAP SAN Administration Guide*.

The following figure illustrates the port hierarchy in a clustered Data ONTAP system:



### **Roles for LIFs**

A LIF role determines the kind of traffic that is supported over the LIF, along with the failover rules that apply and the firewall restrictions that are in place. A LIF can have any one of the five roles: node management, cluster management, cluster, intercluster, and data.

#### Node-management LIF

The LIF that provides a dedicated IP address for managing a particular node and gets created at the time of creating or joining the cluster. These LIFs are used for system maintenance, for example, when a node becomes inaccessible from the cluster. Node-management LIFs can be configured on either node-management or data ports.

The node-management LIF can fail over to other data or node-management ports on the same node.

Sessions established to SNMP and NTP servers use the node-management LIF. AutoSupport requests are sent from the node-management LIF.

#### **Cluster-management LIF**

The LIF that provides a single management interface for the entire cluster. Cluster-management LIFs can be configured on node-management or data ports.
The LIF can fail over to any node-management or data port in the cluster. It cannot fail over to cluster or intercluster ports.

#### **Cluster LIF**

The LIF that is used for intracluster traffic. Cluster LIFs can be configured only on cluster ports.

Cluster LIFs must always be created on 10-GbE network ports.

**Note:** Cluster LIFs need not be created on 10-GbE network ports in N3150 storage systems.

These interfaces can fail over between cluster ports on the same node, but they cannot be migrated or failed over to a remote node. When a new node joins a cluster, IP addresses are generated automatically. However, if you want to assign IP addresses manually to the cluster LIFs, you must ensure that the new IP addresses are in the same subnet range as the existing cluster LIFs.

#### Data LIF

The LIF that is associated with a Vserver and is used for communicating with clients. Data LIFs can be configured only on data ports.

You can have multiple data LIFs on a port. These interfaces can migrate or fail over throughout the cluster. You can modify a data LIF to serve as a Vserver management LIF by modifying its firewall policy to **mgmt**.

For more information about Vserver management LIFs, see the *Clustered Data ONTAP System Administration Guide for Cluster Administrators*.

Sessions established to NIS, LDAP, Active Directory, WINS, and DNS servers use data LIFs.

#### Intercluster LIF

The LIF that is used for cross-cluster communication, backup, and replication. Intercluster LIFs can be configured on data ports or intercluster ports. You must create an intercluster LIF on each node in the cluster before a cluster peering relationship can be established.

These LIFs can fail over to data or intercluster ports on the same node, but they cannot be migrated or failed over to another node in the cluster.

## Related concepts:

"Roles for network ports" on page 8

# **Characteristics of LIFs**

LIFs with different roles have different characteristics. A LIF role determines the kind of traffic that is supported over the interface, along with the failover rules that apply, the firewall restrictions that are in place, the security, the load balancing, and the routing behavior for each LIF.

# Compatibility with port roles and port types

	Data LIF	Cluster LIF	Node- management LIF	Cluster management LIF	Intercluster LIF
Primary traffic types	NFS server, CIFS server, NIS client, Active Directory, LDAP, WINS, DNS client and server, iSCSI and FC server	Intracluster	SSH server, HTTPS server, NTP client, SNMP, AutoSupport client, DNS client, loading code updates	SSH server, HTTPS server	Cross-cluster replication
Compatible with port roles	Data	Cluster	Node-management, data	Data	Intercluster, data
Compatible with port types	All	No interface group or VLAN	All	All	All
Notes	SAN LIFs cannot fail over. These LIFs also do not support load balancing.	Unauthenticated, unencrypted; essentially an internal Ethernet "bus" of the cluster. All network ports in the cluster role in a cluster should have the same physical characteristics (speed).	In new node-management LIFs, the default value of the use-failover-group parameter is <b>disabled</b> The use-failover-group parameter can be set to either <b>system-defined</b> or <b>enabled</b> .		Traffic flowing over intercluster LIFs is not encrypted.

# Security

	Data LIF	Cluster LIF	Node- management LIF	Cluster management LIF	Intercluster LIF
Require private IP subnet?	No	Yes	No	No	No
Require secure network?	No	Yes	No	No	Yes
Default firewall policy	Very restrictive	Completely open	Medium	Medium	Very restrictive
Is firewall customizable?	Yes	No	Yes	Yes	Yes

# Failover

	Data LIF	Cluster LIF	Node- management LIF	Cluster management LIF	Intercluster LIF
Default behavior	Includes all data ports on home node as well as one alternate node	Must stay on node and uses any available cluster port	Default is none, must stay on the same port on the node	Default is failover group of all data ports in the entire cluster	Must stay on node, uses any available intercluster port
Is customizable?	Yes	No	Yes	Yes	Yes

# Routing

	Data LIF	Cluster LIF	Node- management LIF	Cluster management LIF	Intercluster LIF
When is a default route needed?	When clients or domain controller are on different IP subnet	Never	When any of the primary traffic types require access to a different IP subnet	When administrator is connecting from another IP subnet	When other intercluster LIFs are on a different IP subnet
When is static route to a specific IP subnet needed?	Rare	Never	Rare	Rare	When nodes of another cluster have their intercluster LIFs in different IP subnets
When is static host route to a specific server needed?	To have one of the traffic types listed under node-management LIF go through a data LIF rather than a node-management LIF. This requires a corresponding firewall change.	Never	Rare	Rare	Rare

# Automatic LIF rebalancing

	Data LIF	Cluster LIF	Node- management LIF	Cluster management LIF	Intercluster LIF
Automatic LIF rebalancing	Yes If enabled, LIFs automatically migrate to other failover ports based	Yes Cluster network traffic is automatically distributed across cluster LIFs based on load.	No	No	No
	on load provided no CIFS or NFSv4 connections are on them.				
DNS: use as DNS server?	Yes	No	No	No	No
DNS: export as zone?	Yes	No	No	No	No

# LIF limits

There are limits on each type of LIF that you should consider when planning your network. You should also be aware of the effect of the number of LIFs in your cluster environment.

The maximum number of LIFs that are supported on a node is 262. You can create additional cluster, cluster-management, and intercluster LIFs, but creating these LIFs requires a reduction in the number of data LIFs.

LIF type	Minimum	Maximum	Effect of increasing the number of LIFs
Data LIFs	1 per Vserver	<ul> <li>128 per node with failover enabled</li> <li>256 per node without failover enabled</li> </ul>	<ul> <li>Increased client-side resiliency and availability if configured across the NICs of the cluster</li> <li>Increased granularity for load balancing</li> </ul>
Cluster LIFs	2 per node	NA	Increased cluster-side bandwidth if configured on an additional NIC
Node- management LIFs	1 per node	1 per port and per subnet	Negligible
Cluster- management LIFs	1 per cluster	NA	Negligible
Intercluster LIFs	<ul> <li>0 without cluster peering</li> <li>1 per node if cluster peering is enabled</li> </ul>	NA	Increased intercluster bandwidth if configured on an additional NIC

# **Guidelines for creating LIFs**

There are certain guidelines that you should consider before creating a LIF.

Consider the following points while creating a LIF:

- In data LIFs used for file services, the default data protocol options are NFS and CIFS.
- In node-management LIFs, the default data protocol option is set to **none** and the firewall policy option is automatically set to **mgmt**.

You can use such a LIF as a Vserver management LIF. For more information about using a Vserver management LIF to delegate Vserver management to Vserver administrators, see the *Clustered Data ONTAP System Administration Guide for Cluster Administrators*.

- In cluster LIFs the default data protocol option is set to **none** and the firewall policy option is automatically set to **cluster**
- You use FlexCache to enable caching to a 7-Mode volume that exists outside the cluster.

Caching within the cluster is enabled by default and does not require this parameter to be set. For information about caching a FlexVol volume outside the cluster, see the *Data ONTAP Storage Management Guide for 7-Mode*.

- FC LIFs can be configured only on FC ports. iSCSI LIFs cannot coexist with any other protocols. For more information about configuring the SAN protocols, see the *Clustered Data ONTAP SAN Administration Guide*.
- NAS and SAN protocols cannot coexist on the same LIF.
- The firewall policy option associated with a LIF is defaulted to the role of the LIF except for a Vserver management LIF.

For example, the default firewall policy option of a data LIF is **data**. For more information about firewall policies, see the *Clustered Data ONTAP System Administration Guide for Cluster Administrators*.

• Avoid configuring LIFs with addresses in the 192.168.1/24 and 192.168.2/24 subnets. Doing so might cause the LIFs to conflict with the private iWARP interfaces and prevent the LIFs from coming online after a node reboot or LIF migration.

# Guidelines for creating LIFs with IPv6 addresses

You should be aware of some guidelines before you create LIFs with IPv6 addresses.

- IPv6 must be enabled on the cluster.
- The IPv6 addresses must be unique, unicast addresses.
- The prefix for the IPv6 address should be ::/64
- You cannot configure a LIF with any of the following IPv6 addresses:
  - Multicast addresses

Multicast addresses begin with FF.

- Link-local addresses

Link-local addresses always begin with FE80. With the 64-bit interface identifier, the prefix for link-local addresses is always FE80::/64

- IPv4-compatible addresses
   0:0:0:0:0:0:w.x.y.z or ::w.x.y.z (where w.x.y.z is the dotted decimal representation of a public IPv4 address)
- IPv4-mapped addresses

0:0:0:0:0:FFFF:w.x.y.z or ::FFFF:w.x.y.z. It is used to represent an IPv4-only node to an IPv6 node.

- Unspecified addresses
   0:0:0:0:0:0:0:0 or ::
- Loop back addresses
   0:0:0:0:0:0:0:1 or ::1

# **Creating a LIF**

A LIF is an IP address associated with a physical port. If there is any component failure, a LIF can fail over to or be migrated to a different physical port, thereby continuing to communicate with the cluster.

# Before you begin

- The underlying physical network port must have been configured to the administrative **up** status.
- You should have considered the guidelines for creating LIFs: Guidelines for creating LIFs
- If you want to create LIFs with IPv6 addresses, you should have considered the guidelines for assigning IPv6 addresses: Guidelines for assigning IPv6 addresses for LIFs

## About this task

- You can create both IPv4 and IPv6 LIFs on the same network port.
- You cannot assign NAS and SAN protocols to a LIF.

The supported protocols are CIFS, NFS, FlexCache, iSCSI, and FCP.

• The data-protocol option must be specified when the LIF is created, and cannot be modified later.

If you specify **none** as the value for the data-protocol option, the LIF does not support any data protocol.

• A cluster LIF should not be on the same subnet as a management LIF or a data LIF.

## Procedure

1. Use the **network interface create** command to create a LIF.

```
cluster1::> network interface create -vserver vs1 -lif datalif1 -role data -home-node node-4 -home-port elc -address 192.0.2.145 -netmask 255.255.255.0 -firewall-policy data -auto-revert true
```

- 2. Optional: If you want to assign an IPv6 address in the -address option, then perform the following steps:
  - a. Use the **ndp** -**p** command to view the list of RA prefixes learned on various interfaces. The **ndp** -**p** command is available from the node shell.
  - b. Use the format **prefix** :: **id** to construct the IPv6 address manually. **prefix** is the prefix learned on various interfaces. For deriving the **id**, choose one of the following:
    - A random, 64-bit hexadecimal number
    - LLA address configured on the interface
- **3**. Use the **network interface show** command to verify that LIF has been created successfully.

The following example demonstrates different LIFs created in the cluster:

<pre>cluster1::&gt; network interface show</pre>						
	Logical	Status	Network	Current	Current	Is
Vserver	Interface	Admin/Oper	Address/Mask	Node	Port	Home
cluster1						
	cluster_mgm	ıt up∕up	192.0.2.3/24	node-1	ela	true
node-1	aluc1	un /un	102 0 2 12/24	nodo 1	o <b>0</b> o	+ 2010
	clus2	up/up up/up	192.0.2.12/24	node=1	e0a e0h	true
	mamt1	up/up up/up	192.0.2.68/24	node-1	ela	true
node-2	5					
	clus1	up/up	192.0.2.14/24	node-2	e0a	true
	clus2	up/up	192.0.2.15/24	node-2	e0b	true
nada 2	mgmt1	up/up	192.0.2.69/24	node-2	ela	true
noue-s	clus1	un/un	192 0 2 17/24	node-3	e0a	true
	clus2	up/up up/up	192.0.2.18/24	node-3	e0b	true
	mgmt1	up/up	192.0.2.68/24	node-3	e1a	true
node-4						
	clus1	up/up	192.0.2.20/24	node-4	e0a	true
	Clus2	up/up	192.0.2.21/24	node-4	e⊍b	true
vs1	IIIgIIItI	up/up	192.0.2./0/24	noue-4	era	true
131	datalif1	up/down	192.0.2.145/30	node-4	e1c	true
14 entries v	were display	ed.				

The following example demonstrates data LIFs named datalif3 and datalif4 configured with IPv4 and IPv6 addresses respectively:

cluster1::> network interface show						
Vserver	Logical Interface	Status Admin/Oper	Network Address/Mask	Current Node	Current Port	Is Home
cluster1						
	cluster_mg	mt up/up	192.0.2.3/24	node-1	e1a	true
node-1						
	clus1	up/up	192.0.2.12/24	node-1	e0a	true
	clus2	up/up	192.0.2.13/24	node-1	e0b	true
	mgmt1	up/up	192.0.2.68/24	node-1	ela	true
node-2		,				
	clusl	up/up	192.0.2.14/24	node-2	e0a	true
	clus2	up/up	192.0.2.15/24	node-2	e0b	true
	mgmt1	up/up	192.0.2.69/24	node-2	ela	true
node-3		,				
	clusl	up/up	192.0.2.17/24	node-3	e0a	true
	clus2	up/up	192.0.2.18/24	node-3	e0b	true
	mgmt1	up/up	192.0.2.68/24	node-3	ela	true
node-4		,				
	clusl	up/up	192.0.2.20/24	node-4	e0a	true
	clus2	up/up	192.0.2.21/24	node-4	e0b	true
	mgmt1	up/up	192.0.2.70/24	node-4	ela	true
vs1						
	datalif1	up/down	192.0.2.145/30	node-4	elc	true
vs3		,				
	data1113	up/up	192.0.2.146/30	node-3	e0c	true
	datalif4	up/up	2001::2/64	node-3	e⊍c	true
16 entries	were displa	yed.				

- 4. Use the **network ping** command to verify that the configured IPv4 addresses are reachable.
- 5. Use the **ping6** command (available for the nodeshell) to verify that the IPv6 addresses are reachable. All the name mapping and host-name resolution services, such as DNS, NIS, LDAP, and Active Directory, must be reachable from the data, cluster-management, and node-management LIFs of the cluster.

#### Related concepts:

"Roles for LIFs" on page 26 **Related tasks**: "Creating or adding a port to a failover group" on page 38 "Displaying LIF information" on page 69

# Modifying a LIF

You can modify a LIF by changing the attributes such as the home node or the current node, administrative status, IP address, netmask, failover policy, or the firewall policy. You can also modify the address family of a LIF from IPv4 to IPv6. However, you cannot modify the data protocol that is associated with a LIF when the LIF was created.

## About this task

LIF.

- To modify a data LIF with NAS protocols to also serve as a Vserver management LIF, you must modify the data LIF's firewall policy to **mgmt**.
- You cannot modify the data protocols used by a LIF. To modify the data protocols used by a LIF, you must delete and re-create the
- You cannot modify either the home node or the current node of a node-management LIF.
- To modify the address family of a LIF from IPv4 to IPv6, you must do the following:
  - Use the colon notation for the IPv6 address.
  - Add a new value for the -netmask-length parameter.
- You cannot modify the auto-configured link-local IPv6 addresses.
- You cannot change the routing group of a LIF belonging to the IPv4 address family to a routing group assigned to an IPv6 LIF.

## Procedure

1. Use the **network interface modify** command to modify a LIF's attributes. The following example shows how to modify a LIF datalif1 that is located on the Vserver vs0. The LIF's IP address is changed to 172.19.8.1 and its network mask is changed to 255.255.0.0.

cluster1::> network interface modify -vserver vs0 -lif datalif1 -address 172.19.8.1 -netmask 255.255.0.0
-auto-revert true

- 2. Use the **network ping** command to verify that the IPv4 addresses are reachable.
- **3**. Use the **ping6** command to verify that the IPv6 addresses are reachable. The ping6 command is available from the node shell.

# Migrating a LIF

You might have to migrate a LIF to a different port on the same node or a different node within the cluster, if the port is either faulty or requires maintenance.

## Before you begin

- The destination node and ports must be operational and must be able to access the same network as the source port.
- Failover groups must have been set up for the LIFs.

# About this task

- You must migrate LIFs hosted on the ports belonging to a NIC to other ports in the cluster, before removing the NIC from the node.
- You must execute the command for migrating a cluster LIF from the node where the cluster LIF is hosted.
- You can migrate a node-management LIF to any data or node-management port on the home node, even when the node is out of quorum.

For more information about quorum, see the *Clustered Data ONTAP System Administration Guide for Cluster Administrators*.

Note: A node-management LIF cannot be migrated to a remote node.

• You cannot migrate iSCSI LIFs from one node to another node.

To overcome this problem, you must create an iSCSI LIF on the destination node. For information about guidelines for creating an iSCSI LIF, see the *Clustered Data ONTAP SAN Administration Guide*.

• VMware VAAI copy offload operations fail when you migrate the source or the destination LIF.

For more information about VMware VAAI, see the *Clustered Data ONTAP File Access and Protocols Management Guide*.

## Procedure

Depending on whether you want to migrate a specific LIF or all the LIFs, perform the appropriate action:

If you want to migrate	Enter the following command
A specific LIF	network interface migrate
All the data and cluster-management LIFs on a node	network interface migrate-all

## Example

The following example shows how to migrate a LIF named datalif1 on the Vserver vs0 to the port e0d on node0b:

cluster1::> network interface migrate -vserver vs0 -lif datalif1 -dest-node node0b -dest-port e0d

The following example shows how to migrate all the data and cluster-management LIFs from the current (local) node:

( cluster1::> network interface migrate-all -node local

# Reverting a LIF to its home port

You can revert a LIF to its home port after it fails over or is migrated to a different port either manually or automatically. If the home port of a particular LIF is unavailable, the LIF remains at its current port and is not reverted.

## About this task

• If you administratively bring the home port of a LIF to the **up** state before setting the automatic revert option, the LIF is not returned to the home port.

- The node-management LIF does not automatically revert unless the value of the auto revert option is set to **true**.
- Cluster LIFs always revert to their home ports irrespective of the value of the auto revert option.

## Procedure

Depending on whether you want to revert a LIF to its home port manually or automatically, perform one of the following steps:

If you want to revert a LIF to its home port	Then enter the following command
Manually	<pre>network interface revert -vserver vserver_name -lif lif_name</pre>
Automatically	<pre>network interface modify -vserver vserver_name -lif lif_name -auto-revert true</pre>

*vserver\_name* is the name of the Vserver. *lif\_name* is the LIF name.

## Related tasks:

"Displaying LIF information" on page 69

# **Deleting a LIF**

You can delete an LIF that is not required.

#### Before you begin

LIF or LIFs to be deleted must not be in use.

#### Procedure

1. Use the **network interface delete** command to do the following:

If you want to	Enter the command		
Delete a LIF	network interface delete -lif lifname		
Delete all the LIFs	network interface delete -lif *		

cluster1::> network interface delete -vserver vs1 -lif mgmtlif2

2. Use the **network interface show** command to confirm that the LIF is deleted and the routing group associated with the LIF is not deleted.

### Related tasks:

"Displaying LIF information" on page 69

# Configuring failover groups for LIFs (cluster administrators only)

*LIF failover* refers to the automatic migration of a LIF in response to a link failure on the LIF's current network port. When such a port failure is detected, the LIF is migrated to a working port.

A failover group contains a set of network ports (physical, VLANs, and interface groups) on one or more nodes. A LIF can subscribe to a failover group. The network ports that are present in the failover group define the failover targets for the LIF.

You can manage failover groups by adding ports to them, removing ports from them, renaming them, and displaying information about them.

# Scenarios that cause a LIF failover

LIF failover occurs in scenarios such as port failure, network interface failure, or cable failure. LIFs can be associated with failover rules that enable you to reroute the network traffic to other available ports in the cluster.

LIF failover occurs in the following scenarios:

- When there is a power failure
- When automatic revert is enabled on a LIF and that LIF's home port reverts to the administrative **up** status

The LIF automatically migrates back to the home port.

- When the port hosting a LIF is in the administrative **down** status The LIFs move to another port.
- When a node reboots or falls out of quorum

The LIFs on that node fail over to the ports on other nodes. If a node returns to quorum, the LIFs automatically revert to the ports on the node, provided the ports are the home ports for the LIF and automatic revert is enabled on the LIF.

For more information about quorum, see the *Clustered Data ONTAP System Administration Guide for Cluster Administrators*.

• When automatic revert is enabled on a LIF and that LIF's home port reverts to the administrative **up** status the LIF automatically migrates back to the home port.

# Types of failover groups

Failover groups for LIFs can be system-defined or user-defined. Additionally, a failover group called *clusterwide* exists and is maintained automatically.

Failover groups are of the following types:

• System-defined failover groups: Failover groups that automatically manage LIF failover targets on a per-LIF basis.

This is the default failover group for data LIFs in the cluster.

For example, when the value of the failover-group option is **system-defined**, the system will automatically manage the LIF failover targets for that LIF, based on the home node or port of the LIF.

**Note:** All the network ports should be assigned correct port roles, and all the network ports should be in the same subnet.

- User-defined failover groups: Customized failover groups that can be created when the system-defined failover groups do not meet your requirements.
   For example, you can create a failover group consisting of all 10-GbE ports that enables LIFs to fail over only to the high-bandwidth ports.
- Clusterwide failover group: Failover group that consists of all the data ports in the cluster.

This is the default failover group for the cluster-management LIFs only.

For example, when the value of the failover-group option is **cluster-wide**, every data port in the cluster will be defined as the failover targets for that LIF.

# Relation between LIF roles and failover groups

The purpose and the default behavior assigned to any LIF are described by the role associated with that LIF. A LIF can subscribe to a failover group, which will automatically configure the LIF with a list of failover targets for each physical port in the failover group.

LIF role	Failover group	Failover target role	Failover target nodes	
Cluster LIF	system-defined (default)	cluster	home node	
Node management	system-defined (default)	node-management	home node	
LIF	user-defined	node-management or data	home node	
	cluster-wide (default)		any node	
Cluster management	system-defined	node-management or data	home node or any	
	user-defined		node	
Data LIF	system-defined (default)	data	home node or any	
	user-defined		node	
Intercluster LIF	system-defined (default)	intercluster	home node	
	user-defined	intercluster or data		

The relation between LIF roles and failover groups is described in the following table.

#### **Related concepts:**

"Roles for LIFs" on page 26

# Creating or adding a port to a failover group

You can create failover groups or add a port to a failover group by using the **network interface failover-groups create** command.

## About this task

- If you have LIFs in different VLANs or broadcast domains, you must configure failover groups for each VLAN or broadcast domain.
  - You must then configure the LIFs hosted on a particular VLAN or broadcast domain to subscribe to the corresponding failover group.
- Failover groups do not apply in a SAN iSCSI or FC environment.

## Procedure

Use the **network interface failover-groups create** command to create a failover group or add a port to an existing failover group. For more information about this command, see the man page.

cluster1::> network interface failover-groups create -failover-group failover-group\_2 -node cluster1-01 -port ele

# Renaming a failover group

To rename a failover group, you can use the **network interface failover-groups rename** command.

#### Procedure

Use the **network interface failover-groups rename** command to rename a failover group. For more information about this command, see the man page.

cluster1::> network interface failover-group rename -failover-group clusterwide -new-name clyde

## Removing a port from or deleting a failover group

To remove a port from a failover group or to delete an entire failover group, you use the **network interface failover-groups delete** command.

### Before you begin

For deleting an entire failover group, the failover group must not be used by any LIF.

## Procedure

Depending on whether you want to remove a port from a failover group or delete a failover group, complete the applicable step:

If you want to	Then, enter the following command
Remove a port from a failover group	network interface failover-groups delete -failover-group <i>failover_group_name</i> -node <i>node_name</i> -port <i>port</i> <b>Note:</b> If you delete all ports from the failover group, the failover group is also deleted.
Delete a failover group	<pre>network interface failover-groups delete -failover-group failover_group_name [-node   -port] *</pre>

failover\_group\_name specifies the name of the user-defined failover group. port

specifies the failover target port. *node\_name* specifies the node on which the port resides.

## Example

The following example shows how to delete port e1e from the failover group named failover-group\_2:

cluster1::> network interface failover-groups delete -failover-group failover-group\_2 -node cluster1-01 -port ele

# Enabling or disabling failover of a LIF

You can enable a LIF to fail over by specifying whether the LIF should subscribe to the system-defined or user-defined failover group. You can also disable a LIF from failing over.

## About this task

The values of the following parameters in the **network interface modify** command together determine the failover behavior of LIFs:

 -failover-policy: Enables you to specify the order in which the network ports are chosen during a LIF failover and enables you to prevent a LIF from failing over.

This parameter can have one of the following values:

- **nextavail** (default): Enables a LIF to fail over to the next available port, preferring a port on the current node.

In some instances, a LIF configured with the **nextavail** failover policy selects a failover port on a remote node, even though a failover port is available on the local node. No outages will be seen in the cluster, because the LIFs continue to be hosted on valid failover ports.

- priority: Enables a LIF to fail over to the first available port specified in the user-defined failover group (failover targets can be shown with the network interface show -failover command).
- disabled: Disables (prevents) a LIF from failing over.
- -failover-group: Specifies the failover behavior configured for the LIF. The value can be set to:
  - **system-defined** specifies that the LIF uses the implicit system-defined failover behavior for the LIF's role.
  - [empty] specifies that the LIF is not configured to use a failover group.
  - [**user-defined failover group**] specifies that the LIF is configured to fail over to any available port present in the failover group.

## Procedure

Depending on whether you want a LIF to fail over, complete the appropriate action:

If you want to	Enter the following command	
Enable a LIF to	<pre>network interface modify -vserver vserver_name -lif lif_name</pre>	
use a	-failover-policy {nextavail   priority} -failover-group	
user-defined	failover_group_name	
failover group		

If you want to	Enter the following command
Enable a LIF to use the system-defined failover group	<pre>network interface modify -vserver vserver_name -lif lif_name -failover-policy {nextavail   priority} -failover-group system-defined Note: When you create a LIF, it uses the system-defined failover group by default.</pre>
Disable a LIF from failing over	<pre>network interface modify -vserver vserver_name -lif lif_name -failover-policy disabled</pre>

# Managing routing in a Vserver (cluster administrators only)

You can control how LIFs in a Vserver use your network for outbound traffic by configuring routing groups and static routes. A set of common routes are grouped in a routing group that makes the administration of routes easier.

#### **Routing group**

A routing table. Each LIF is associated with one routing group and uses only the routes of that group. Multiple LIFs can share a routing group.

**Note:** For backward compatibility, if you want to configure a route per LIF, you should create a separate routing group for each LIF.

#### Static route

A defined route between a LIF and a specific destination IP address; the route can use a gateway IP address.

# Creating a routing group

When a LIF is created, an associated routing group is automatically created. You cannot modify or rename an existing routing group; therefore, you might have to create a routing group.

## About this task

If you want to segregate the data LIFs from the management LIFs, you must create different routing groups for each kind of LIFs.

The following rules apply when creating routing groups:

- The routing group and the associated LIFs should be in the same subnet.
- All LIFs sharing a routing group must be on the same IP subnet.
- All next-hop gateways must be on that same IP subnet.
- A Vserver can have multiple routing groups, but a routing group belongs to only one Vserver.
- The routing group name must be unique in the cluster and should not contain more than 64 characters.
- You can create a maximum of 256 routing groups per node.

## Procedure

Use the **network routing-groups create** command to create a routing group. You can use the optional -metric parameter with this command to specify a hop count for the routing group. The default settings for this parameter are 10 for management interfaces, 20 for data interfaces, and 30 for cluster interfaces. This parameter is used for source-IP address-selection of user-space applications such as NTP.

For more information about this command, see the man page.

cluster1::> network routing-groups create -vserver vs1 -routing-group d192.0.2.166 -subnet 192.0.2.165/24 -role data -metric 20

## What to do next

You should assign a static route to the created routing group. **Related tasks**: "Creating a LIF" on page 32 "Modifying a LIF" on page 34 "Displaying routing information" on page 71

# Deleting a routing group

To delete a routing group, you can use the **network routing-groups delete** command.

## Before you begin

- You must have modified the LIFs that are using the routing group to use a different routing group.
- You must have deleted the routes within the routing group.

## Procedure

Use the **network routing-groups delete** command as shown in the following example to delete a routing group. For more information about this command, see the man page.

cluster1::> network routing-groups delete -vserver vs1 -routing-group d192.0.2.165/24

#### Related tasks:

"Deleting a LIF" on page 36

## Creating a route within a routing group

You cannot modify the various parameters of an existing route such as changing its name, destination IP address, network mask, or gateway IP address. Instead, you must delete the route and create a new route with the desired parameter values.

#### About this task

In all installations, routes are not needed for routing groups of cluster LIFs.

**Note:** You should only modify routes for routing groups of node-management LIFs when you are logged into the console.

#### Procedure

1. Use the **network routing-groups create** command to create a routing group.

cluster1::> network routing-groups create -vserver vs1 -routing-group d192.0.2.166/24 -subnet 192.0.2.0/24 -role
data -metric 10

2. Create a route within the routing group by using the **network routing-groups** route create command.

The following example shows how to create a route for a LIF named mgmtif2. The routing group uses the destination IP address 0.0.0.0, the network mask 255.255.255.0, and the gateway IP address 192.0.2.1.

cluster1::> network routing-groups route create -vserver vs0 -routing-group d192.0.2.166/24 -destination 0.0.0.0/0
-gateway 192.0.2.1 -metric 10

#### Related tasks:

"Creating a routing group" on page 43

"Displaying routing information" on page 71

# Deleting a static route

You can delete an unneeded static route to a LIF from a routing group.

## Procedure

Use the **network routing-groups route delete** command to delete a static route. For more information about this command, see the man page. The following example deletes a static route associated with a LIF named mgmtif2 from routing group d192.0.2.167/24. The static route has the destination IP address 192.40.8.1.

cluster1::> network routing-groups route delete -vserver vs0 -routing-group d192.0.2.167/24 -destination 0.0.0.0/0

# **Configuring host-name resolution**

Clustered Data ONTAP supports two methods for host-name resolution: DNS and hosts table. Cluster administrators can configure DNS and hosts file naming services for host-name lookup in the admin Vserver. Cluster administrators and Vserver administrators can configure DNS for host-name lookup in the Vserver.

# Host-name resolution for the admin Vserver

Cluster administrators can configure only DNS and hosts table for host-name lookup in the admin Vserver. All applications except CIFS discovery use the host-name configuration of the admin Vserver. You cannot use NIS configuration for the admin Vserver.

Host-name resolution for the admin Vserver is configured when the cluster is created. This configuration is propagated to each node as it joins the cluster.

## Hosts table configuration for the admin Vserver

You can use the **vserver services dns hosts** command for configuring the hosts file that resides in the root volume of the admin Vserver.

By default, the order of lookup for the admin Vserver is hosts file first and then DNS.

## DNS configuration for the admin Vserver

It is best to configure DNS on the admin Vserver at the time of cluster creation either by using the Cluster Setup wizard or the **cluster create** command.

If you want to configure DNS later, you should use the **vserver services dns create** command.

# Host-name resolution for a Vserver

A cluster administrator or a Vserver administrator can configure DNS for host-name lookup in a Vserver.

Starting with Data ONTAP 8.1, each Vserver has its own DNS configuration. Each Vserver communicates with its DNS server in the Vserver's context, which includes the Vserver's LIF and routing tables. Client requests are authenticated and authorized by using the specific Vserver network configuration.

DNS configuration is mandatory when CIFS is used for data access.

DNS services can also be configured on a Vserver for FlexVol volumes by using the Vserver Setup wizard. If you want to configure DNS later, you must use the **vserver services dns create** command. For more information about the Vserver Setup wizard, see the *Clustered Data ONTAP Software Setup Guide*.

# Managing the hosts table (cluster administrators only)

A cluster administrator can add, modify, delete, and view the host name entries in the hosts table of the admin Vserver. A Vserver administrator can configure the host name entries only for the assigned Vserver.

# Commands for managing DNS host name entries

You can use the **vserver services dns hosts** commands to create, modify, or delete DNS host table entries.

While creating or modifying the DNS host name entries, you can specify multiple alias addresses separated by commas.

If you want to	Use this command
Create a DNS host name entry	<b>vserver services dns hosts create</b> <b>Note:</b> The name of the admin Vserver is same as the cluster name.
Modify a DNS host name entry	vserver services dns hosts modify
Delete a DNS host name entry	vserver services dns hosts delete

For more information, see the man pages for the **vserver services dns hosts** commands.

#### **Related concepts:**

"Host-name resolution for a Vserver" on page 47

## Managing DNS domains for host-name resolution

You can create, modify, view, or remove DNS configurations for the admin Vserver and each Vserver on the cluster. Vserver administrators can also create, modify, view, or remove the DNS configuration on their respective Vserver.

A cluster administrator can configure DNS on the admin Vserver. A cluster administrator can also log in to a particular Vserver context and configure DNS on the Vserver.

DNS services can also be configured on a Vserver with FlexVol volumes by using the Vserver Setup wizard. You can configure DNS services on a Vserver with Infinite Volume by using the **vserver services dns create** command.

For more information about the Vserver Setup wizard, see the *Clustered Data ONTAP Software Setup Guide*.

# Commands for managing DNS domain configurations

You can use the **vserver services dns** commands to create, modify, or delete a DNS domain configuration.

If you want to	Enter this command
Configure a DNS domain configuration	vserver services dns create
Modify a DNS domain configuration	vserver services dns modify
Delete a DNS domain configuration	vserver services dns delete

For more information, see the man pages for the **vserver services dns** commands. **Related concepts**:

"Host-name resolution for the admin Vserver" on page 47

"Host-name resolution for a Vserver" on page 47

# Balancing network loads to optimize user traffic (cluster administrators only)

You can configure your cluster to serve client requests from appropriately loaded LIFs and to automatically migrate these LIFs to under-utilized ports. This results in a more balanced utilization of LIFs and ports, which in turn allows for better utilization and performance of the cluster.

# Load balancing types

DNS load balancing and Automatic LIF rebalancing methods aid in selecting an appropriately loaded data LIF and balancing user network traffic across all available ports (physical or interface groups).

# **DNS load balancing**

With DNS load balancing, you can create a DNS load balancing zone on the Vserver that returns the least-loaded LIF based on the network traffic and the availability of the port resources (CPU usage, throughput, open connections, and so on). By configuring a DNS load balancing zone, you can balance new client connections better across available resources. This leads to improved performance for the entire cluster. Also, no manual intervention is required for deciding which LIFs to use when mounting a particular Vserver. You can use the DNS load balancing method to balance load for only new share connections. DNS load balancing works with NFSv3, NFSv4, NFSv4.1, CIFS, SMB 2.0, SMB 2.1, and SMB 3.0.

## Automatic LIF rebalancing

With automatic load balancing, LIFs are dynamically migrated to ports with low utilization, based on the failover rules. Automatic LIF rebalancing works only with NFSv3 connections. Automatic LIF rebalancing provides the following benefits:

- Different client connections use different bandwidth; therefore, LIFs can be migrated based on the load capacity.
- When new nodes are added to the cluster, LIFs can be migrated to the new ports.

# Guidelines for assigning load balancing weights

Data ONTAP automatically assigns weights to data LIFs by collecting periodic statistics on the current node and port resources (CPU usage, throughput, open connections, and so on). To override the automatic assignment, you must consider certain guidelines for manually assigning load balancing weights to LIFs.

The following guidelines should be kept in mind when assigning load balancing weights:

• The load balancing weight is inversely related to the load on a LIF.

A data LIF with a high load balancing weight is made available for client requests more frequently than one that has a low load balancing weight. For example, lif1 has a weight of 10 and lif2 has a weight of 1. For any mount request, lif1 is returned 10 times more than lif2.

- If all LIFs in a load balancing zone have the same weight, LIFs are selected with equal probability.
- When manually assigning load balancing weights to LIFs, you must consider conditions such as load, port capacity, client requirements, CPU usage, throughput, open connections, and so on.

For example, in a cluster having 10-GbE and 1-GbE data ports, the 10-GbE ports can be assigned a higher weight so that it is returned more frequently when any request is received.

• When a LIF is disabled, it is automatically assigned a load balancing weight of 0.

# Assigning a load balancing weight to a LIF

Data ONTAP automatically assigns load balancing weights to data LIFs. You can override the automatic assignment of the load balancing weights to LIFs.

## Before you begin

- You must be logged in at the advanced privilege level or higher.
- You must have read the guidelines for manually assigning the load balancing weights to LIFs: Considerations for assigning load balancing weights

## About this task

You might want to modify the automatic load balancing weights, for example, if a cluster has both 10-GbE and 1-GbE data ports, the 10-GbE ports can be assigned a higher weight so that it is returned more frequently when any request is received.

## Procedure

To assign or modify the weight of a LIF, enter the following command: network interface modify -vserver vserver\_name -lif lif\_name -lb-weight weight

*vserver\_name* specifies the node or Vserver on which the LIF is to be modified. *lif\_name* specifies the name of the LIF that is to be modified. *weight* specifies the weight of the LIF. A valid load balancing weight is any integer between 0 and 100.

```
cluster1::*> network interface modify -vserver vs0 -lif lif3 -lb-weight 3
```

#### **Related tasks:**

"Modifying a LIF" on page 34

# How DNS load balancing works

A DNS load balancing zone is a DNS server inside the cluster that dynamically evaluates the load on all LIFs and returns an appropriately loaded LIF. DNS load balancing works by assigning a weight (metric), which in turn is based on the load, to each LIF.

Clients mount a Vserver by specifying an IP address (associated with a LIF) or a host name (with multiple IP addresses as managed by a site-wide DNS server). LIFs are selected by the site-wide DNS server in a round-robin manner, which could result in overloading of some LIFs. Instead, you can create a DNS load balancing zone that handles the host-name resolution in a Vserver. By configuring

a DNS load balancing zone, you can ensure better balance of the new client connections across available resources, thus leading to improved performance of the cluster. Also, no manual intervention is required for deciding which LIF to use when mounting a particular Vserver.

The DNS load balancing zone dynamically calculates the load on all LIFs. Every LIF is assigned a weight based on its port load and CPU utilization of its home node. LIFs that are on less-loaded ports have a higher probability of being returned in a DNS query. Weights can also be manually assigned. The commands to manually assign weights are available only at the advanced privilege level.

# Creating a DNS load balancing zone

You can create a DNS load balancing zone to facilitate the dynamic selection of a LIF based on the load, that is, the number of clients mounted on a LIF. You can create a load balancing zone while creating a data LIF. In clustered Data ONTAP 8.2, DNS load balancing is not supported for LIFs with IPv6 addresses.

## Before you begin

You must have configured the DNS forwarder on the site-wide DNS server to forward all requests for the load balancing zone to the configured LIFs.

## About this task

- Any data LIF can respond to DNS queries for a DNS load balancing zone name.
- A DNS load balancing zone must have a unique name in the cluster, and the zone name must meet the following requirements:
  - It should not exceed 256 characters.
  - It should include at least one period.
  - The first and the last character should not be a period or any other special character.
  - It cannot include any spaces between characters.
  - Each label in the DNS name should not exceed 63 characters.

A label is the text appearing before or after the period. For example, the DNS zone named storage.company.com has three labels.

## Procedure

Use the **network interface create** command with the *zone\_domain\_name* option to create a DNS load balancing zone. For more information about the command, see the man pages. The following example demonstrates how to create a DNS load balancing zone named storage.company.com while creating the LIF lif1 :

cluster1::> network interface create -vserver vs0 -lif lif1 -role data -home-node node1 -home-port e0c -address 192.0.2.129 -netmask 255.255.255.128 -dns-zone storage.company.com

#### Related tasks:

"Creating a LIF" on page 32

# Adding or removing a LIF from a load balancing zone

You can add or remove a LIF from the DNS load balancing zone of a Vserver to ensure that all the LIFs are evenly used. You can also remove all the LIFs simultaneously from a load balancing zone.

# Before you begin

• All the LIFs in a load balancing zone should belong to the same Vserver.

Note: A LIF can be a part of only one DNS load balancing zone.

• Failover groups for each subnet must have been set up, if the LIFs belong to different subnets.

## About this task

• A LIF that is in the administrative **down** status is temporarily removed from the DNS load balancing zone.

When the LIF comes back to the administrative **up** status, the LIF is added automatically to the DNS load balancing zone.

• Load balancing is not supported for LIFs hosted on a VLAN port.

## Procedure

Depending on whether you want to add or remove a LIF, perform the appropriate action:

If you want to	Then, enter the following command	
Add a LIF to a load balancing zone	<pre>network interface modify -vserver vserver_name -lif lif_name -dns-zone zone_name</pre>	
	Example: cluster1::> network interface modify -vserver vs1 -lif data1 -dns-zone cifs.company.com	
Remove a LIF from a load balancing zone	<pre>network interface modify -vserver vserver_name -lif lif_name -dns-zone none Example:</pre>	
Remove all LIFs from a load balancing zone	<pre>network interface modify -vserver vserver_name -lif * -dns-zone none Example:</pre>	

*vserver\_name* specifies the node or Vserver on which the LIF is created. *lif\_name* specifies the name of the LIF. *zone\_name* specifies the name of the DNS load balancing zone.

### Related tasks:

"Modifying a LIF" on page 34

# How automatic LIF rebalancing works

In automatic LIF rebalancing, LIFs are automatically and periodically migrated to a less-utilized port based on the configured failover rules. Automatic LIF rebalancing allows even distribution of current load.

LIFs are migrated based on the weights assigned to the LIFs.

When new NICs are added to the cluster, add the ports to the failover group to which the automatically rebalancing LIFs belong. The network ports are then automatically included the next time that load is calculated dynamically, and each time thereafter.

Automatic LIF rebalancing is available only under the advanced privilege level of operation.

# Enabling or disabling automatic LIF rebalancing

You can enable or disable automatic LIF rebalancing in a Vserver. By enabling automatic LIF rebalancing, LIFs can be migrated to a less-utilized port on another node based on the failover rules.

## Before you begin

You must be logged in at the advanced privilege level.

#### About this task

- By default, automatic LIF rebalancing is disabled on a LIF.
- Automatic LIF rebalancing gets disabled if a LIF is enabled for automatically reverting to the home port (by enabling the auto-revert option in the **network interface modify** command).
- You can also restrict a LIF with automatic load balancing enabled to only fail over within the ports specified in a user-defined failover group.

#### Procedure

1. To enable or disable automatic LIF rebalancing on a LIF, use the **network interface modify** command.

The following example shows how to enable automatic LIF rebalancing on a LIF and also restrict the LIF to fail over only to the ports in the failover group failover-group\_2:

cluster1::\*>network interface modify -vserver vs1 -lif data1 -failover-policy priority -failover-group failover-group\_2 -allow-lb-migrate true

**2**. To verify whether automatic LIF rebalancing is enabled for the LIF, use the **network interface show** command.



#### Related tasks:

"Modifying a LIF" on page 34

# Combining load balancing methods in a Vserver accessible in a multiprotocol environment

In a Vserver that is accessible from multiple protocols, such as CIFS and NFS, you can use DNS load balancing and automatic LIF rebalancing simultaneously.

## Before you begin

You must have configured the DNS site-wide server to forward all DNS requests for NFS and CIFS traffic to the assigned LIFs.

## About this task

You should not create separate DNS load balancing zones for each protocol when the following conditions are true:

- Automatic LIF rebalancing is not configured (disabled by default)
- Only NFSv3 protocol is configured

Because automatic LIF rebalancing can be used only with NFSv3 connections, you should create separate DNS load balancing zones for CIFS and NFSv3 clients. Automatic LIF rebalancing on the zone used by CIFS clients is disabled automatically, as CIFS connections cannot be nondisruptively migrated. This enables the NFS connections to take advantage of automatic LIF rebalancing.

#### Procedure

1. Use the **network interface modify** command to create a DNS load balancing zone for the NFS connections and assign LIFs to the zone. The following

example shows how to create a DNS load balancing zone named nfs.company.com and assign LIFs named lif1, lif2, and lif3 to the zone:

cluster1::> network interface modify -vserver vs0 -lif lif1..lif3 -dns-zone nfs.company.com

2. Use the **network interface modify** command to create a DNS load balancing zone for the CIFS connections and assign LIFs to the zone. The following example shows how to create a DNS load balancing zone named cifs.company.com and assign LIFs named lif4, lif5, and lif6 to the zone.

cluster1::> network interface modify -vserver vs0 -lif lif4..lif6 -dns-zone cifs.company.com

**3**. Use the **set -privilege advanced** command to log in at the advanced privilege level. The following example shows how to enter the advanced privilege mode:

cluster1::> set -privilege advanced

Warning: These advanced commands are potentially dangerous; use them only when directed to do so by technical support. Do you want to continue?  $\{y \mid n\}$ : y

4. Use the **network interface modify** command to enable automatic LIF rebalancing on the LIFs that are configured to serve NFS connections. The following example shows how to enable automatic LIF rebalancing on LIFs named lif1, lif2, and lif3 in the DNS load balancing zone created for NFS connections.

cluster1::\*> network interface modify -vserver vs0 -lif lif1..lif3 -allow-lb-migrate true

**Note:** Because automatic LIF rebalancing is disabled for CIFS, automatic LIF rebalancing should not be enabled on the DNS load balancing zone that is configured for CIFS connections.

#### Results

NFS clients can mount by using nfs.company.com and CIFS clients can map CIFS shares by using cifs.company.com. All new client requests are directed to a LIF on a less-utilized port. Additionally, the LIFs on nfs.company.com are migrated dynamically to different ports based on the load.

# Managing SNMP on the cluster (cluster administrators only)

You can configure SNMP to monitor the cluster to avoid issues before they occur and to respond to issues when they occur. Managing SNMP involves configuring SNMP users and configuring SNMP traphosts for specific events. You can create and manage read-only SNMP users in the data Vserver. You can also configure data LIFs to receive SNMP requests on the Vserver.

SNMP network management workstations or managers can query your Vserver's SNMP agent for information. The SNMP agent gathers information and forwards it to the SNMP managers. The SNMP agent also generates trap notifications whenever specific events occur. The SNMP agent on the Vserver has read-only privileges—that is, it cannot be used for any set operations or for taking a corrective action in response to a trap. Data ONTAP provides an SNMP agent compatible with SNMP versions v1, v2c, and v3. SNMPv3 offers advanced security by using passphrases and encryption.

## Related tasks:

"Creating a LIF" on page 32

# What MIBs are

A MIB file is a text file that describes SNMP objects and traps. MIBs describe the structure of the management data of the storage system and they use a hierarchical namespace containing object identifiers (OIDs). Each OID identifies a variable that can be read by using SNMP.

Because MIBs are not configuration files and Data ONTAP does not read these files, SNMP functionality is not affected by MIBs. Data ONTAP provides two MIB files:

- A custom MIB
- An Internet SCSI (iSCSI) MIB

Data ONTAP supports IPv6 (RFC 2465), TCP (RFC 4022), UDP (RFC 4113), and ICMP (RFC 2466) MIBs, which show both IPv4 and IPv6 data, are supported.

Data ONTAP also provides a short cross-reference between object identifiers (OIDs) and object short names in the traps.dat file.

**Note:** The latest versions of the Data ONTAP MIBs and traps.dat files are available online on the N series support website (accessed and navigated as described in Websites). However, the versions of these files on the web site do not necessarily correspond to the SNMP capabilities of your Data ONTAP version. These files are provided to help you evaluate SNMP features in the latest Data ONTAP version.

IBM N series support website: www.ibm.com/storage/support/nseries/

# Creating an SNMP community

You can create an SNMP community that acts as an authentication mechanism between the management station and the cluster or the Vserver when using SNMPv1 and SNMPv2c. By creating SNMP communities in the data Vserver, you can execute commands such as **snmpwalk** and **snmpget** on the data LIFs.

## About this task

- In new installations of Data ONTAP, SNMPv1 and SNMPv2c are disabled by default.
  - SNMPv1 and SNMPv2c are enabled when an SNMP community is created.
- Data ONTAP supports read-only communities.
- By default, a firewall data policy has SNMP service set to deny. Create a new data policy with SNMP service set to allow when creating an SNMP user for a data Vserver.
- You can create SNMP communities for the SNMPv1 and SNMPv2c users for both the admin Vserver and the data Vserver.

#### Procedure

1. Use the **system snmp community add** command to create an SNMP community. The following example shows how you can create an SNMP community in the admin Vserver:

cluster1::> system snmp community add -type ro -community-name comty1

2. Use the vserver option of the **system snmp community add** to create an SNMP community in the Vserver. The following example shows how you can create an SNMP community in the data Vserver, vs0:

cluster1::> system snmp community add -type ro -community-name comty2 -vserver vs0

**3**. Use the **system snmp community show** command to verify that the communities have been created. The following example shows different communities created for SNMPv1 and SNMPv2c:

```
cluster1::> system snmp community show
cluster1 ro comty1
vs0 ro comty2
2 entries were displayed.
```

4. Use the **firewall policy show** -service snmp command to verify if SNMP is allowed as a service in data firewall policy. The following example shows that the snmp service is allowed in the data firewall policy:

cluster1::> fire (system service	wall policy es firewall	show -s policy	service snmp show)
Policy	Service	Action	IP-List
cluster			
	snmp	allow	0.0.0/0
data	snmp	allow	0.0.0.0/0, ::/0
Intercluster	snmp	allow	0.0.0/0
mgmt 4 entries were d	snmp isplayed.	allow	0.0.0.0/0, ::/0

5. Optional: If the value of the snmp service is **deny**, use the **system services firewall policy create** to create a data firewall policy with the value of the snmp service as **allow**. The following example shows how you can create a new data firewall policy data1 with value of the snmp service as allow, and verify if this has been created successfully:

```
cluster1::> system services firewall policy create -policy data1 -service
  snmp -action allow -ip-list 0.0.0/0
cluster1::> firewall policy show -service snmp
  (system services firewall policy show)
Policy
                Service Action IP-List
____.
cluster
                snmp
                           allow 0.0.0/0
data
                 snmp
                           deny 0.0.0.0/0, ::/0
intercluster
                           allow 0.0.0/0
                 snmp
mgmt
                           allow 0.0.0.0/0, ::/0
                 snmp
data1
                 snmp
                           allow 0.0.0.0/0, ::/0
5 entries were displayed.
```

6. Use the **network interface modify** command with the -firewall-policy option to put into effect a firewall policy for a data LIF. The following example shows how the created data firewall policy data1 can be assigned to a LIF datalif1:

cluster1::> network interface modify -vserver vs1 -lif datalif1 -firewall-policy data1

# Configuring SNMPv3 users in a cluster

SNMPv3 is a secure protocol when compared to SNMPv1 and SNMPv2c. To use SNMPv3, you must configure an SNMPv3 user to run the SNMP utilities from the SNMP manager.

## Procedure

Use the **security login create** command to create an SNMPv3 user. You are prompted to provide the following information:

- Engine ID: Default and recommended value is local EngineID
- Authentication protocol
- Authentication password
- Privacy protocol

Privacy protocol password

## Results

The SNMPv3 user can log in from the SNMP manager by using the user name and password and run the SNMP utility commands.

# SNMPv3 security parameters

SNMPv3 includes an authentication feature that, when selected, requires users to enter their names, an authentication protocol, an authentication key, and their desired security level when invoking a command.

Parameter	Command-line option	Description
engineID	-e EngineID	Engine ID of the SNMP agent. Default value is local EngineID (recommended).
securityName	-u Name	User name must not exceed 31 characters.
authProtocol	-a { <b>MD5</b>   <b>SHA</b> }	Authentication type can be MD5 or SHA.
authKey	-A PASSPHRASE	Passphrase with a minimum of eight characters.
securityLevel	-l {authNoPriv   AuthPriv   noAuthNoPriv}	Security level can be Authentication, No Privacy; Authentication, Privacy; or no Authentication, no Privacy.
privProtocol	-x { none   des}	Privacy protocol can be <b>none</b> or <b>des</b>
privPassword	-X password	Password with a minimum of eight characters.

The following table lists the SNMPv3 security parameters:

# Examples for different security levels

This example shows how an SNMPv3 user created with different security levels can use the SNMP client-side commands, such as **snmpwalk**, to query the cluster objects.

For better performance, you should retrieve all objects in a table rather than a single object or a few objects from the table.

**Note:** You must use **snmpwalk** 5.3.1 or later when the authentication protocol is SHA.

## Security level: authPriv

The following output shows the creation of an SNMPv3 user with the **authPriv** security level.
cluster1::> security login create -username snmpv3user -application snmp -authmethod usm
Please enter the authoritative entity's EngineID [local EngineID]:
Please choose an authentication protocol (none, md5, sha) [none]:sha
Please enter authentication protocol password (minimum 8 characters long):
Please enter authentication protocol password again:
Please choose a privacy protocol (none, des) [none]: des
Please enter privacy protocol password (minimum 8 characters long):
Please enter privacy protocol password again:

The following output shows the SNMPv3 user running the **snmpwalk** command:

```
$ snmpwalk -v 3 -u snmpv3user -a SHA -A password1! -x DES -X password1! -1
authPriv 192.0.2.62 .1.3.6.1.4.1.789.1.5.8.1.2
enterprises.789.1.5.8.1.2.1028 = "vol0"
enterprises.789.1.5.8.1.2.1032 = "vol0"
enterprises.789.1.5.8.1.2.1038 = "root_vs0"
enterprises.789.1.5.8.1.2.1042 = "root_vstrap"
enterprises.789.1.5.8.1.2.1064 = "vol1"
```

#### Security level: authNoPriv

The following output shows the creation of an SNMPv3 user with the **authNoPriv** security level.

cluster1::> security login create -username snmpv3user1 -application snmp -authmethod usm -role admin
Please enter the authoritative entity's EngineID [local EngineID]:
Please choose an authentication protocol (none, md5, sha) [none]: md5
Please enter authentication protocol password (minimum 8 characters long):
Please enter authentication protocol password again:
Please choose a privacy protocol (none, des) [none]: none

The following output shows the SNMPv3 user running the **snmpwalk** command:

```
$ snmpwalk -v 3 -u snmpv3user1 -a MD5 -A password1! -1 authNoPriv 192.0.2.62 .1.3.6.1.4.1.789.1.5.8.1.2
enterprises.789.1.5.8.1.2.1028 = "vol0"
enterprises.789.1.5.8.1.2.1032 = "vol0"
enterprises.789.1.5.8.1.2.1038 = "root_vs0"
enterprises.789.1.5.8.1.2.1042 = "root_vstrap"
enterprises.789.1.5.8.1.2.1064 = "vol1"
```

#### Security level: noAuthNoPriv

The following output shows the creation of an SNMPv3 user with the **noAuthNoPriv** security level.

cluster1::> security login create -username snmpv3user2 -application snmp -authmethod usm -role admin

Please enter the authoritative entity's EngineID [local EngineID]:

Please choose an authentication protocol (none, md5, sha) [none]: none

The following output shows the SNMPv3 user running the **snmpwalk** command:

```
$ snmpwalk -v 3 -u snmpv3user2 -l noAuthNoPriv 192.0.2.62 .1.3.6.1.4.1.789.1.5.8.1.2
enterprises.789.1.5.8.1.2.1028 = "vol0"
enterprises.789.1.5.8.1.2.1038 = "root_vs0"
enterprises.789.1.5.8.1.2.1042 = "root_vstrap"
enterprises.789.1.5.8.1.2.1064 = "vol1"
```

## SNMP traps

SNMP traps capture system monitoring information that is sent as an asynchronous notification from the SNMP agent to the SNMP manager. There are three types of SNMP traps: standard, built-in, and user-defined. User-defined traps are not supported in clustered Data ONTAP.

A trap can be used to check periodically for operational thresholds or failures that are defined in the MIB. If a threshold is reached or a failure is detected, the SNMP agent sends a message (trap) to the traphosts alerting them of the event.

#### Standard SNMP traps

These traps are defined in RFC 1215. There are five standard SNMP traps that are supported by Data ONTAP: coldStart, warmStart, linkDown, linkUp, and authenticationFailure.

**Note:** The authenticationFailure trap is disabled by default. You must use the **system snmp authtrap** command to enable the trap. See the man pages for more information.

#### **Built-in SNMP traps**

Built-in traps are predefined in Data ONTAP and are automatically sent to the network management stations on the traphost list if an event occurs. These traps, such as diskFailedShutdown, cpuTooBusy, and volumeNearlyFull, are defined in the custom MIB.

Each built-in trap is identified by a unique trap code.

## **Configuring traphosts**

You can configure the traphost (SNMP manager) to receive notifications (SNMP trap PDUs) when SNMP traps are generated. You can specify either the hostname or the IP address (IPv4 or IPv6) of the SNMP traphost.

#### Before you begin

SNMP and SNMP traps must be enabled on the cluster.

Note: By default SNMP and SNMP traps are enabled.

- DNS must be configured on the cluster for resolving the traphost names.
- IPv6 should be enabled on the cluster to configure SNMP traphosts with IPv6 addresses.

#### Procedure

Use the system snmp traphost add command to add SNMP traphosts.

**Note:** Traps can only be sent when at least one SNMP management station is specified as a traphost.

The following example illustrates addition of a new SNMP traphost named yyy.example.com:

cluster1::> system snmp traphost add -peer-address yyy.example.com

The following example demonstrates how an IPv6 address can be added to configure a traphost :

cluster1::> system snmp traphost add -peer-address 2001:0db8:1:1:209:6bff:feae:6d67

# **Commands for managing SNMP**

You can use the **system snmp** command to manage SNMP, traps, and traphosts. You can use the **security** command to manage SNMP users. You can use the **event** command to manage events related to SNMP traps. Starting from Data ONTAP 8.1.1, you can use commands for managing SNMP users of the data Vserver as well.

- Commands for configuring SNMP
- Commands for managing SNMPv3 users
- Commands for providing courtesy information
- Commands for managing SNMP communities
- Command for displaying SNMP option values
- Commands for managing SNMP traps and traphosts
- Commands for managing events related to SNMP traps

#### Commands for configuring SNMP

If you want to	Use this command
Enable SNMP on the cluster	options -option-name snmp.enable -option-value on Note: The SNMP service must be allowed under the management firewall policy. You can verify whether SNMP is allowed by using the system services firewall policy show command.
Disable SNMP on the cluster	options -option-name snmp.enable -option-value off

#### Commands for managing SNMP v1, v2c, and v3 users

If you want to	Use this command
Configure SNMP users	security login create
Display SNMP users	security snmpusers
Delete SNMP users	security login delete
Modify the access-control role name of a login method for SNMP users	security login modify

# Commands for providing contact and location information

If you want to	Use this command
Display or modify the contact details of the cluster	system snmp contact
Display or modify the location details of the cluster	system snmp location

# Commands for managing SNMP communities

If you want to	Use this command
Add a read-only (ro) community	system snmp community add
Delete a community or all communities	system snmp community delete
Display the list of all communities	system snmp community show

# Command for displaying SNMP option values

If you want to	Use this command
Display the current values of all SNMP options, such as contact, location, init, traphosts, and community	system snmp show

# Commands for managing SNMP traps and traphosts

If you want to	Use this command
Enable SNMP traps sent from the cluster	system snmp init -init 1
Disable SNMP traps sent from the cluster	system snmp init -init 0
Add a traphost that receives notification for specific events in the cluster	system snmp traphost add
Delete a traphost	system snmp traphost delete
Display the list of traphosts	system snmp traphost show

# Commands for managing events related to SNMP traps

If you want to	Use this command
Display the events for which SNMP traps (built-in) are generated	event route show Note: You use the snmp-support parameter to view the SNMP-related events. You use the instance parameter to view the corrective action for an event.
Display a list of SNMP trap history records, which are event notifications that have been sent to SNMP traps	event snmphistory show
Delete an SNMP trap history record	event snmphistory delete

For more information about the **system snmp** and **event** commands, see the man pages.

# Viewing network information

You can view information related to LIFs, routes, failover rules, failover groups, DNS, NIS, and connection information. This information might be useful in situations such as troubleshooting the cluster. However, if you are Vserver administrator, you can view the information related to the assigned Vserver.

# Displaying network port information (cluster administrators only)

You can display information about a specified port or about all ports in the cluster. You can view information such as network port role (cluster, data, or node-management), Link status (up or down), MTU setting, Autonegotiation setting (true or false), Duplex mode and operational status (half or full), speed setting and operational status (1-Gbps or 10-Gbps), and port's interface group, if applicable.

## About this task

After removing a NIC from a node, use the **network port show** command to verify that all the NIC has been removed from its slot.

## Procedure

To display port information, enter the following command: network port show The command displays the following information:

- Node name
- Port name
- Port role (cluster, data, or node-management)
- Link status (up or down)
- MTU setting
- Autonegotiation setting (true or false)
- Duplex mode and operational status (half or full)
- Port speed setting and operational status (1 gigabit or 10 gigabits per second)
- The port's interface group, if applicable
- The port's VLAN tag information, if applicable

If data for a field is not available (for instance, the operational duplex and speed for an inactive port), the field is listed as undef.

**Note:** You can get all available information by specifying the -instance parameter or get only the required fields by specifying the fields parameter. The following example displays information about all network ports in a cluster containing four nodes:

cluste	r1::> n	etwork port s	how				
					Auto-Negot	Duplex	Speed (Mbps)
Node	Port	Role	Link	MTU	Admin/Oper	Admin/Oper	Admin/Oper
node1							
nouer	e0a	cluster	un	1500	true/true	ful1/ful1	auto/1000
	e0b	cluster	up	1500	true/true	full/full	auto/1000
	eOc	data	up	1500	true/true	full/full	auto/1000
	e0d	data	up	1500	true/true	full/full	auto/1000
	ela	node-mamt	up	1500	true/true	full/full	auto/1000
	e1b	data	down	1500	true/true	full/half	auto/10
	e1c	data	down	1500	true/true	full/half	auto/10
	e1d	data	down	1500	true/true	full/half	auto/10
node2							
	e0a	cluster	up	1500	true/true	full/full	auto/1000
	e0b	cluster	up	1500	true/true	full/full	auto/1000
	e0c	data	up	1500	true/true	full/full	auto/1000
	e0d	data	up	1500	true/true	full/full	auto/1000
	e1a	node-mgmt	up	1500	true/true	full/full	auto/1000
	e1b	data	down	1500	true/true	full/half	auto/10
	e1c	data	down	1500	true/true	full/half	auto/10
	e1d	data	down	1500	true/true	full/half	auto/10
node3							
	e0a	cluster	up	1500	true/true	full/full	auto/1000
	e0b	cluster	up	1500	true/true	full/full	auto/1000
	e0c	data	up	1500	true/true	full/full	auto/1000
	e0d	data	up	1500	true/true	full/full	auto/1000
	ela	node-mgmt	up	1500	true/true	full/full	auto/1000
	e1b	data	down	1500	true/true	full/half	auto/10
	elc	data	down	1500	true/true	full/half	auto/10
	eld	data	down	1500	true/true	full/half	auto/10
node4						6 33 (6 33	
	e0a	cluster	up	1500	true/true	full/full	auto/1000
	e0b	cluster	up	1500	true/true	full/full	auto/1000
	e0c	data	up	1500	true/true	full/full	auto/1000
	e⊍d	data	up	1500	true/true		auto/1000
	ela	node-mgmt	up	1500	true/true	TUII/TUII	auto/1000
	elb	udīd	down	1500	true/true	iull/nalT	auto/10
	elc	uata	down	1500	true/true	full/half	auto/10
	610	udīd	aown	1200	true/true	iuii/nait	dUT0/10

# Displaying information about a VLAN (cluster administrators only)

To display information about a VLAN, you can use the  ${\tt network\ port\ vlan\ show}$  command.

#### About this task

You can get all available information by specifying the -instance parameter or get only the required fields by specifying the fields parameter.

#### Procedure

Use the **network port vlan show** command to view information about VLANs. To customize the output, you can enter one or more optional parameters. For more information about this command, see the man page.

		Network	Network	
Node	VLAN Name	Port	VLAN ID	MAC Address
cluster	~1-01			
	a0a-10	a0a	10	02:a0:98:06:10:b2
	a0a-20	a0a	20	02:a0:98:06:10:b2
	a0a-30	a0a	30	02:a0:98:06:10:b2
	a0a-40	a0a	40	02:a0:98:06:10:b2
	a0a-50	a0a	50	02:a0:98:06:10:b2
cluster	^1 <b>-</b> 02			
	a0a-10	a0a	10	02:a0:98:06:10:ca
	a0a-20	a0a	20	02:a0:98:06:10:ca
	a0a-30	a0a	30	02:a0:98:06:10:ca
	a0a-40	a0a	40	02:a0:98:06:10:ca
	a0a-50	a0a	50	02:a0:98:06:10:ca

# Displaying interface group information (cluster administrators only)

You can display information about an interface group to determine its configuration.

## Procedure

To display information about interface groups, enter the following command: network port ifgrp show

The command displays the following information:

- The node on which the interface group is located
- The interface group's name
- Distribution function (MAC, IP, port, or sequential)
- The interface group's Media Access Control (MAC) address
- Whether all aggregated ports are active (full participation), whether some are inactive (partial participation), or whether none are active
- The list of network ports included in the interface group

**Note:** You can get all available information by specifying the -instance parameter or get only the required fields by specifying the fields parameter. The following example displays information about all interface groups in the cluster:

cluster1: Node	:> network Port IfGrp	port ifgrp sł Distribution Function	now MAC Address	Active Ports	Ports
cluster1-	·01				
	a0a	ip	02:a0:98:06:10:b2	full	e7a, e7b
cluster1-	02				
	a0a	sequential	02:a0:98:06:10:ca	full	e7a, e7b
cluster1-	·03				
	a0a	port	02:a0:98:08:5b:66	full	e7a, e7b
cluster1-	-04				
	a0a	mac	02:a0:98:08:61:4e	full	e7a, e7b
4 entries	were displ	ayed.			

# **Displaying LIF information**

You might have to view the information about a LIF in instances such as diagnosing basic problems with LIFs, finding out the routing group of a LIF. If you are a Vserver admin, you can view the information about the LIFs associated with the Vserver.

## About this task

You might to have the view the information about a LIF in the following scenarios:

- Verifying the IP address associated with the LIF
- Verifying the administrative status of the LIF
- Verifying the operational status of the LIF

The operational status of data LIFs is determined by the status of the Vserver with which the data LIFs are associated. When a Vserver is stopped, the operational status of LIF changes to **down**. When the Vserver is started again, the operational status changes to **up** 

- Determining the node and the port the LIF resides on
- Finding out the routing group that LIF belongs to
- Checking for duplicate IP addresses
- Verifying if the network port belongs to the correct subnet

#### Procedure

To view the LIF information, use the **network interface show** command. You can get all available information by specifying the -instance parameter or get only the required fields by specifying the fields parameter. If data for a field is not available, the field is listed as undef..

The following example displays general information about all LIFs in a cluster:

vs1::> netwo Vserver	ork interfac Logical Interface	e show Status Admin/Oper	Network Address/Mask	Current Node	Current Port	Is Home
xample	1:61		100 0 0 100/00			
	1111	up/up	192.0.2.129/22	node-01	e0d	false
node	cluster_mgm	t up/up	192.0.2.3/20	node-02		
ode-01					e0c	false
	clus1	up/up	192.0.2.65/18	node-01	e0a	true
	clus2	up/up	192.0.2.66/18	node-01	e0b	true
	mgmt1	up/up	192.0.2.1/20	node-01	e0c	true
ode-02	clus1	un/un	192 0 2 67/18	node-02		
	clus?	up/up	102 0 2 68/18	nodo 02	e0a	true
	crusz	up/up	192.0.2.00/10	node 02	e0b	true
	IIIgIIItZ	up/up	192.0.2.2/20	1006-02		
vs1					e0d	true
	dl	up/up	192.0.2.130/21	node-01	e0d	false
	d2	up/up	192.0.2.131/21	node-01	e0d	true
	data3	up/up	192.0.2.132/20	node-02	e0c	true
1 entries w	were display	ed.				

The following example shows detailed information about a LIF:

```
cluster1::> network interface show -lif data1 -instance
                    Vserver Name: vs1
              Logical Interface: data1
                           Role: data
                  Data Protocol: nfs,cifs,fcache
                      Home Node: node-1
                      Home Port: e0c
                   Current Node: node-3
                   Current Port: eOc
              Operational Status: up
                Extended Status: -
                        Is Home: false
                Network Address: 10.72.34.39
                        Netmask: 255.255.192.0
                IPv4 Link Local: -
            Bits in the Netmask: 18
             Routing Group Name: d10.72.0.0/18
          Administrative Status: up
                Failover Policy: nextavail
                Firewall Policy: data
                    Auto Revert: false
             Use Failover Group: system-defined
                       DNS Zone: xxx.example.com
            Failover Group Name:
                       FCP WWPN:
```

#### Related tasks:

"Creating a LIF" on page 32 "Modifying a LIF" on page 34 "Migrating a LIF" on page 34 "Reverting a LIF to its home port" on page 35 "Deleting a LIF" on page 36

# **Displaying routing information**

You can display information about all routing groups and static routes.

#### Procedure

Depending on the information that you want to view, enter the applicable command:

If you want to display information about	Then, enter the following command
Static routes	network routing-groups route show
Routing groups	network routing-groups show

Note: You can get all available information by specifying the -instance parameter.

#### Example

The following example displays static routes within a routing group originating from Vserver vs1:

vs1::> ne	etwork rou Routing	ting-groups route	show	
Vserver	Group	Destination	Gateway	Metric
vs1				
	d172.17.	176.0/24		
		0.0.0.0/0	172.17.176.1	20

#### **Related tasks:**

"Creating a routing group" on page 43

"Deleting a routing group" on page 44

"Creating a route within a routing group" on page 44

"Deleting a static route" on page 45

# Displaying host name entries (cluster administrators only)

You can use the vserver services dns hosts show command to view the host name and IP addresses mappings (with alias addresses, if any) in the hosts table of the admin Vserver.

#### **Procedure**

To view the host name entries in the hosts table of the admin Vserver, enter the following command: vserver services dns hosts show

#### Example

The following sample output shows the hosts table:

```
cluster1::> vserver services dns hosts show
Vserver
        Address
                   Hostname
                                Aliases
_____ __ ___
cluster1
        10.72.219.36 lnx219-36
cluster1
        10.72.219.37 lnx219-37 lnx219-37.example.com
2 entries were displayed.
```

# **Displaying DNS domain configurations**

You can display the DNS domain configuration of one or more Vservers. You can filter the output by specifying specific parameters based on the output fields that you want to view.

#### Procedure

To view the DNS domain configurations, enter the following command: vserver services dns show

## Example

The following example shows the output of the **vserver services dns show** command:

cluster1::>	vserver servi	ices dns show	Name
Vserver	State	Domains	Servers
cluster1	enabled	xyz.company.com	192.56.0.129,
vs1	enabled	xyz.company.com	192.56.0.130 192.56.0.129,
vs2	enabled	xyz.company.com	192.56.0.130 192.56.0.129,
vs3	enabled	XVZ COMPANY COM	192.56.0.130 192.56.0.129
		xyz.company.com	192.56.0.130
4 entries we	ere displayed.		

# Displaying information about failover groups (cluster administrators only)

You can view information about failover groups including the list of nodes and ports in each failover group.

#### Procedure

Use the **network interface failover-groups show** command to display information about failover groups.

The following example displays information about all failover groups on a two-node cluster:

```
cluster1::> network interface failover-groups show
Failover
                            Port
Group
               Node
----- ------
clusterwide
               node-02
                             e0c
               node-01
                             e0c
               node-01
                             e0d
fg1
               node-02
                             e0c
                             e0c
               node-01
5 entries were displayed.
```

#### **Related tasks**:

"Creating or adding a port to a failover group" on page 38

"Renaming a failover group" on page 39

"Enabling or disabling failover of a LIF" on page 40

"Removing a port from or deleting a failover group" on page 39

# Viewing failover targets of LIFs

You might have to check whether the failover rules of a LIF are configured correctly. In order to prevent misconfiguration of the failover rules, you can view the failover target for a LIF or all LIFs.

## About this task

By viewing the failover targets for a LIF, you can check for the following:

- If the LIFs are configured correctly for failover
- If the failover groups contain LIFs to fail over
- If the failover target of a data LIF is not a management port (e0M)

#### Procedure

1. Use the failover option of the **network interface show** command to view the failover targets of a LIF. The following example displays information about the failover targets of the different LIFs in the cluster:

cluster1	::> network inte	rface show -failove	er Failover	Failover
Vserver	Interface	Node·Port	Policy	Group
vs1				
	clus1	node1:e0a	nextavail	system-defined
	clus2	node1:e0b	nextavail	system-defined
	mgmt1	nodel:ela	nextavail	system-defined
		Failover Targets:	nodel:ela	
vs2				
	clus1	node2:e0a	nextavail	system-defined
	clus2	node2:e0b	nextavail	system-defined
	mgmt1	node2:e1a	nextavail	system-defined
		Failover Targets:	node2:e1a	
	mgmt2	node2:e1a	nextavail	system-defined
		Failover Targets:	node2:e1a	
vs3				
	clus1	node3:e0a	nextavail	system-defined
	clus2	node3:e0b	nextavail	system-defined
	mgmt1	node3:e1a	nextavail	system-defined
		Failover Targets:	node3:e1a	
vs4				
	clus1	node4:e0a	nextavail	system-defined
	clus2	node4:e0b	nextavail	system-defined
	mgmt1	node4:e1a	nextavail	system-defined
		Failover Targets:	ie3070-4:e1a	

The failover targets row shows the (prioritized) list of node-port combinations for a given LIF.

2. Optional: Use the failover-targets option of the **network interface show** command to view the failover targets of a LIF. The following example shows the failover targets of each LIF in the cluster:

clus (n vser	terl::> etwork i ver lif	net int show -fields failover-targets nterface show) failover-targets
vs1	clus1	
vs1	clus2	
vs1	mgmt1	- nodel·ela
vs2	clus1	
vs2	clus2	
vs2	mgmt1	- nada2+a1a
vs2	mgmt2	
vs3	clus1	nouez:ela
vs3	clus2	
vs3	mgmt1	-
vs4	clus1	nodes:ela
vs4	clus2	
vs4	mgmt1	- node4:e1a

# Viewing LIFs in a load balancing zone

To verify whether a load balancing zone is configured correctly, you can view all the LIFs that belong to a load balancing zone. You can also view the load balancing zone of a particular LIF.

## Procedure

Depending on the LIFs and details that you want to view, perform the appropriate action:

If you want to view	Then, enter the following command
LIFs in a load balancing zone	network interface show -dns-zone zone_name
	<i>zone_name</i> specifies the name of the load balancing zone.
Load balancing zone of a particular LIF	network interface show -lif <i>lif</i> -fields dns-zone
	<i>lif</i> specifies the name of the LIF
All LIFs and their corresponding load balancing zones	network interface show -fields dns-zone

## Example

The following example shows the details of all the LIFs in the load balancing zone storage.company.com:

cluster1::> net int show -vserver vs0 -dns-zone storage.company.com

Vserver	Logical Interface	Status Admin/Oper	Network Address/Mask	Current Node	Current Port	Is Home
VSU						
	lif3	up/up	10.98.226.225/20	ndeux-11	e0c	true
	lif4	up/up	10.98.224.23/20	ndeux-21	e0c	true
	lif5	up/up	10.98.239.65/20	ndeux-11	e0c	true
	lif6	up/up	10.98.239.66/20	ndeux-11	e0c	true
	lif7	up/up	10.98.239.63/20	ndeux-21	e0c	true
	lif8	up/up	10.98.239.64/20	ndeux-21	e0c	true
6 entries w	ere display	ed.				

The following example shows the DNS zone details of the LIF lif1:

cluster1::> network interface show -lif data3 -fields dns-zone Vserver lif dns-zone -----vs0 data3 storage.company.com

The following example shows the list of all LIFs in the cluster and their corresponding DNS zones:

#### Related tasks:

"Displaying LIF information" on page 69

## **Displaying cluster connections**

You can display all the active connections in the cluster or a count of active connections on the node by client, logical interface, protocol, or service. You can also display all the listening connections in the cluster.

# Displaying active connections by client (cluster administrators only)

You can view the active connections by client to verify the node that a specific client is using and to view imbalances between client counts per node.

#### About this task

The count of active connections by client is useful in the following scenarios:

- Finding a busy or overloaded node because you can view the number of clients that are being serviced by each node.
- Determining why a particular client's access to a volume is slow. You can view details about the node that the client is accessing and then compare it with the node on which the volume resides. If accessing the volume

requires traversing the cluster network, clients might experience decreased performance because of the remote access to the volume on an oversubscribed remote node.

- Verifying that all nodes are being used equally for data access.
- Finding clients that have an unexpectedly high number of connections.
- Verifying if certain expected clients do not have connections to a node.

#### Procedure

Use the **network connections active show-clients** command to display a count of the active connections by client on a node.

cluster	<pre>r1::&gt; network connectio</pre>	ns active
Node	Client IP Address	Count
node0	192.0.2.253	1
	192.0.2.252	2
	192.0.2.251	5
node1	192.0.2.250	1
	192.0.2.252	3
	192.0.2.253	4
node2	customer.example.com	1
	192.0.2.245	3
	192.0.2.247	4
node3	192.0.2.248	1
	customer.example.net	3
	customer.example.org	4

For more information about this command, see the man pages.

# Displaying active connections by protocol (cluster administrators only)

You can use the **network connections active show-protocols** command to display a count of the active connections by protocol (TCP or UDP) on a node. You can use this command to compare the usage of protocols within the cluster.

#### About this task

The **network connections active show-protocols** command is useful in the following scenarios:

• Finding the UDP clients that are losing their connection.

If a node is near its connection limit, UDP clients are the first to be dropped.

• Verifying that no other protocols are being used.

#### Procedure

Use the **network connections active show-protocols** command to display a count of the active connections by protocol on a node. For more information about this command, see the man pages.

cluster1:	::> network	connections	active	show-protocols	
Node	Protocol	Count			
node0					
	UDP	19			
	ТСР	11			
node1					
	UDP	17			
	ТСР	8			
node2					
	UDP	14			
	ТСР	10			
node3					
	UDP	18			
	ТСР	4			

# Displaying active connections by service (cluster administrators only)

You can use the **network connections active show-services** command to display a count of the active connections by service (for example, by NFS, CIFS, mount, and so on) on a node. You can use this command to compare the usage of services within the cluster, which helps to determine the primary workload of a node.

#### About this task

The **network connections active services** command is useful in the following scenarios:

- Verifying that all nodes are being used for the appropriate services and that the load balancing for that service is working.
- Verifying that no other services are being used.

#### Procedure

Use the **network connections active show-services** command to display a count of the active connections by service on a node. For more information about this command, see the man pages.

```
cluster1::> network connections active show-services
Node
         Service Count
_____
          -----
node0
                         3
          mount
                        14
          nfs
          nlm v4
                         4
                         3
          cifs srv
          port_map
                        18
          rclopcp
                        27
node1
          cifs srv
                         3
          rclopcp
                        16
node2
                        13
          rclopcp
node3
          cifs srv
                        17
          rclopcp
```

# Displaying active connections by LIF on a node and Vserver

You can use the **network connections active show-lifs** command to display a count of active connections for each LIF by node and Vserver, and verify connection imbalances between LIFs within the cluster.

## About this task

The **network connections active show-lifs** command is useful in the following scenarios:

- Finding an overloaded LIF by comparing the number of connections on each LIF.
- Verifying that DNS load balancing is working for all data LIFs.
- Comparing the number of connections to the various Vservers to find the Vservers that are used the most.

## Procedure

Use the **network connections active show-lifs** command to display a count of active connections for each LIF by Vserver and node. For more information about this command, see the man pages.

```
cluster1::> network connections active show-lifs
Node
          Vserver Name Interface Name Count
node0
                                              3
          vs0
                        datalif1
          vs0
                        cluslif1
                                              6
                        cluslif2
                                              5
          vs0
node1
          vs0
                        datalif2
                                              3
          vs0
                        cluslif1
                                              3
                                              5
          vs0
                        cluslif2
node2
          vs1
                        datalif2
                                              1
                        cluslif1
                                              5
          vs1
          vs1
                        cluslif2
                                              3
node3
                        datalif1
          vs1
                                              1
          vs1
                        cluslif1
                                              2
```

# Displaying active connections in a cluster

You can use the **network connections active** command to display information about the active connections in a cluster. You can use this command to view the LIF, port, remote IP, service, and protocol used by individual connections.

## About this task

The **network connections active** command is useful in the following scenarios:

- Verifying that individual clients are using the correct protocol and service on the correct node.
- If a client is having trouble accessing data using a certain combination of node, protocol, and service, you can use this command to find a similar client for configuration or packet trace comparison.

#### Procedure

Use the **network connections active** command to display the active connections in a cluster. For more information about this command, see the man pages.

Vserver Name	Interface Name:Local Port	Remote IP Address:Port	Protocol/Service
node0	cluslif1:7070	192.0.2.253:48621	UDP/rclopcp
node0	cluslif1:7070	192.0.2.253:48622	UDP/rclopcp
node0	cluslif2:7070	192.0.2.252:48644	UDP/rclopcp
node0	cluslif2:7070	192.0.2.250:48646	UDP/rclopcp
node0	cluslif1:7070	192.0.2.245:48621	UDP/rclopcp
node0	cluslif1:7070	192.0.2.245:48622	UDP/rclopcp
node0	cluslif2:7070	192.0.2.251:48644	UDP/rclopcp
node0	cluslif2:7070	192.0.2.251:48646	UDP/rclopcp
node0	cluslif1:7070	192.0.2.248:48621	UDP/rclopcp
node0	cluslif1:7070	192.0.2.246:48622	UDP/rclopcp
node0	cluslif2:7070	192.0.2.252:48644	UDP/rclopcp
node0	cluslif2:7070	192.0.2.250:48646	UDP/rclopcp
node0	cluslif1:7070	192.0.2.254:48621	UDP/rclopcp
node0	cluslif1:7070	192.0.2.253:48622	UDP/rclopcp

# Displaying listening connections in a cluster

You can use the **network connections listening show** command to display information about the listening connections in a cluster. You can use this command to view the LIFs and ports that are accepting connections for a given protocol and service.

#### About this task

The **network connections listening show** command is useful in the following scenarios:

- Verifying that the desired protocol or service is listening on a LIF if client connections to that LIF fail consistently.
- Verifying that a UDP/rclopcp listener is opened at each cluster LIF if remote data access to a volume on one node through a LIF on another node fails.
- Verifying that a UDP/rclopcp listener is opened at each cluster LIF if SnapMirror transfers between two nodes in the same cluster fail.
- Verifying that a TCP/ctlopcp listener is opened at each intercluster LIF if SnapMirror transfers between two nodes in different clusters fail.

#### Procedure

Use the **network connections listening show** command to display the listening connections.

```
cluster1::> network connections listening show
Server Name Interface Name:Local Port Protocol/Service
node0 cluslif1:7700
node0 cluslif2:7700
                               UDP/rclopcp
                                UDP/rclopcp
        cluslif1:7700
node1
                                UDP/rclopcp
node1
         cluslif2:7700
                                 UDP/rclopcp
         cluslif1:7700
                                 UDP/rclopcp
node2
        cluslif2:7700
node2
                                 UDP/rclopcp
     cluslif1:7700
node3
                                 UDP/rclopcp
node3
         cluslif2:7700
                                 UDP/rclopcp
8 entries were displayed.
```

# Commands for diagnosing network problems

You can diagnose problems on your network by using commands such as **ping**, **traceroute**, **ndp**, and **pktt**. You can also use commands such as **ping6** and **traceroute6** to diagnose IPv6 problems.

If you want to	Enter this command
Test whether your node can reach other hosts on your network	network ping
Trace the route that the IPv4 packets take to a network node.	network traceroute
Test whether the node can reach other hosts on your IPv6 network	run -node {nodename} ping6 <b>Note:</b> This command is available from the nodeshell.
Trace the route that the IPv6 packets take to a network node	<pre>run -node {nodename} traceroute6 Note: This command is available from the nodeshell.</pre>
Control the address mapping table used by Neighbor Discovery Protocol (NDP)	<pre>run -node {nodename} ndp Note: This command is available from the nodeshell.</pre>
Trace the packets sent and received in the network	<pre>run -node {nodename} pktt start Note: This command is available from the nodeshell.</pre>
View the CDP neighbors of the node. Data ONTAP only supports CDPv1 advertisements.	<pre>run -node {nodename} cdpd show-neighbors Note: This command is available from the nodeshell.</pre>

For more information about these commands, see the man pages.

## Using CDP to detect network connectivity

In a data center, you can use Cisco Discovery Protocol (CDP) to view network connectivity between a pair of physical or virtual systems and their network interfaces.

CDP enables you to automatically discover and view information about directly connected CDP-enabled devices in a network. Each device advertises identification, capabilities, and connectivity information. This information is transmitted in Ethernet frames to a multicast MAC address and is received by all neighboring CDP-enabled devices.

Neighboring devices of the storage system that are discovered by using CDP are called *CDP neighbors*. For two devices to become CDP neighbors, each must have the CDP protocol enabled and correctly configured. The functionality of CDP is limited to directly connected networks. CDP neighbors include CDP-enabled devices such as switches, routers, bridges, and so on.

# **Considerations for using CDP**

By default, Cisco devices and CDP-compliant devices send CDPv2 advertisements. CDP-compliant devices send CDPv1 advertisements only when they receive CDPv1 advertisements. Data ONTAP supports only CDPv1. Therefore, when the storage system sends CDPv1 advertisements, the CDP-compliant neighboring devices send back CDPv1 advertisements. You should consider the following information before enabling CDP on the node:

- You can execute the CDP commands only from the nodeshell.
- CDP is supported for all port roles.
- CDP advertisements are sent and received by ports that are configured with the LIFs and in the **up** state.
- CDP must be enabled on both the transmitting and receiving devices for sending and receiving CDP advertisements.
- CDP advertisements are sent at regular intervals, and you can configure the time interval involved.
- When IP addresses are changed at the storage system side, the storage system sends the updated information in the next CDP advertisement.

**Note:** Sometimes when LIFs are changed on the node, the CDP information is not updated at the receiving device side (for example, a switch). If you encounter such a problem, you should configure the network interface of the node to the **down** status and then to the **up** status.

- Only IPv4 addresses are advertised in CDP advertisements.
- For physical network ports with VLANs, all the LIFs configured on the VLANs on that port are advertised.
- For physical ports that are part of an interface group, all the IP addresses configured on that interface group are advertised on each physical port.
- For an interface group that hosts VLANs, all the LIFs configured on the interface group and the VLANs are advertised on each of the network ports.
- For packets with MTU size equal to or greater than 1,500 bytes, only the number of LIFs that can fit into a 1500 MTU-sized packet is advertised.

## Enabling or disabling CDP

To discover and send advertisements to CDP-compliant neighboring devices, CDP must be enabled on each node of the cluster.

#### About this task

When the cdpd.enable option is set to **on**, CDPv1 is enabled on all physical ports of the node from which the command is run. Starting from Data ONTAP 8.2, CDP is enabled by default. If you change the value of the cdpd.enable option to **off**, the cluster network traffic might not be optimized.

#### Procedure

To enable or disable CDP, enter the following command from the nodeshell: options cdpd.enable {on|off} on—Enables CDP off—Disables CDP

## Configuring hold time for CDP messages

Hold time is the period of time for which all CDP advertisements are stored in a cache in the neighboring CDP-compliant devices. Hold time is advertised in each CDPv1 packet and is updated whenever a CDPv1 packet is received by the storage system.

#### About this task

• The value of the cdpd.holdtime option applies to both nodes of a High Availability (HA) pair.

- The default value of hold time is 180 seconds.
- If an IP address is removed before the hold time expires, the CDP information is cached till the hold time expires.

## Procedure

To configure the hold time, enter the following command from the nodeshell: options cdpd.holdtime *holdtime holdtime* is the time interval, in seconds, for which the CDP advertisements are cached in the neighboring CDP-compliant devices. You can enter values ranging from 10 seconds to 255 seconds.

# Setting the intervals for sending CDP advertisements

CDP advertisements are sent at periodic intervals. You can increase or decrease the intervals between the sending of each CDP advertisement, depending on the network traffic and change in the network topology. You can use the cdpd.interval option to configure the time interval for sending CDP advertisements.

## About this task

The value of the cdpd.interval option applies to both the nodes of an HA pair.

## Procedure

To configure the interval for sending CDP advertisements, enter the following command from the nodeshell:

options cdpd.interval interval

*interval* is the time interval after which CDP advertisements should be sent. The default interval is 60 seconds. The time interval can be set between the range of 5 seconds and 900 seconds.

# Viewing or clearing CDP statistics

You can analyze the CDP statistics to detect any network connectivity issues. You can use the **cdpd show-stats** command to view the CDP send and receive statistics. CDP statistics are cumulative from the time they were last cleared. To clear the CDP statistics, you can use the **cdpd zero-stats** command.

## Before you begin

CDP must be enabled.

## Procedure

Depending on whether you want to view or clear the CDP statistics, complete the following step:

If you want to	Then, enter the following command from the nodeshell
View the CDP statistics	cdpd show-stats
Clear the CDP statistics	cdpd zero-stats

## Example of showing the statistics before and after clearing them

The following example shows the CDP statistics before they are cleared:

system1> cdpd show	-stats					
RECEIVE Packets: Invalid length: Missing TLVs:	9116 0 0	Csum Errors:  Malformed:  Cache overflow:	0 0 0	Unsupported Vers: Mem alloc fails: Other errors:	4561 0 0	
TRANSMIT Packets: Packet truncated:	4557 0	Xmit fails: Mem alloc fails:	0 0	No hostname: Other errors:	0 0	,

This output displays the total packets that are received from the last time the statistics were cleared.

The following command clears the CDP statistics:

system1> cdpd zero-stats

The following output shows the statistics after they are cleared:

system1> cdpd show-s	tats		
RECEIVE Packets: Invalid length: Missing TLVs:	0   Csum Errors: 0   Malformed: 0   Cache overflow:	0   Unsupported Vers: 0   Mem alloc fails: 0   Other errors:	0 0 0
TRANSMIT Packets: Packet truncated:	0   Xmit fails: 0   Mem alloc fails:	0   No hostname: 0   Other errors:	0 0
OTHER Init failures:	0		

After the statistics are cleared, the statistics get added from the time the next CDP advertisement is sent or received.

# Viewing neighbor information by using CDP

You can view information about the neighboring devices connected to each port of your storage system, provided that the port is connected to a CDP-compliant device. You can use the **cdpd show-neighbors** command to view neighbor information.

#### Before you begin

CDP must be enabled.

#### **Procedure**

To view information about all CDP-compliant devices connected to your storage system, enter the following command from the nodeshell: cdpd show-neighbors

# Example

The following example shows the output of the **cdpd show-neighbors** command:

system Local	1> cdpd show-nei Remote	ghbors Remote Interface	Remote	Hold Time	Remote
e0a	sw-215-cr(4C2)	GigabitEthernet1/17	cisco WS-C4948	125	RSI
e0b	sw-215-11(4C5)	GigabitEthernet1/15	cisco WS-C4948	145	SI
e0c	sw-215-11(4C5)	GigabitEthernet1/16	cisco WS-C4948	145	SI

The output lists the Cisco devices that are connected to each port of the storage system. The "Remote Capability" column specifies the capabilities of the remote device that are connected to the network interface. The following capabilities are available:

- R—Router
- T—Transparent bridge
- B—Source-route bridge
- S—Switch
- H—Host
- I—IGMP
- r—Repeater
- P—Phone

# Copyright and trademark information

This section includes copyright and trademark information, and important notices.

## **Copyright information**

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

Portions copyright © 2013 IBM Corporation. All rights reserved.

US Government Users Restricted Rights - Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

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.

References in this documentation to IBM products, programs, or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM's product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any of IBM's or NetApp's intellectual property rights may be used instead of the IBM or NetApp product, program, or service. Evaluation and verification of operation in conjunction with other products, except those expressly designated by IBM and NetApp, are the user's responsibility.

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.A. 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**

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

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

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

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

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

NetApp NetCache is certified RealSystem compatible.

# Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe on any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing to:

IBM Director of Licensing IBM Corporation North Castle Drive Armonk, N.Y. 10504-1785 U.S.A.

For additional information, visit the web at: http://www.ibm.com/ibm/licensing/contact/

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM web sites are provided for convenience only and do not in any manner serve as an endorsement of those web sites. The materials at those web sites are not part of the materials for this IBM product and use of those web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

If you are viewing this information in softcopy, the photographs and color illustrations may not appear.

# Index

# Α

active connections displaying 76 active connections by client displaying 76 admin Vserver host-name resolution 47 automatic LIF rebalancing 55 disabling 55 enabling 55 automatic load balancing 51 assigning weights 51

# С

CDP configuring hold time 82 configuring periodicity 83 Data ONTAP support 82 disabling 82 enabling 82 viewing neighbor information 84 viewing statistics 83 CDP (Cisco Discovery Protocol) 81, 82 CIFS 56 Cisco Discovery Protocol (CDP) 81, 82 cluster interconnect cabling guidelines 2 Cluster default port assignments 8 cluster connections displaying 76 commands 48 managing DNS domain configuration 48 snmp traps 64 system snmp 65 vserver services dns hosts show 72 configuring 23 DNS 47 host-name resolution 47 connections active, displaying count by client on node 76 active, displaying count by LIF on node 79 active, displaying count by protocol on node 77 active, displaying count by service on node 78 active, displaying information about 79 listening, displaying information about 80 copyright and trademark information 87 copyright information 87 creating interface groups 14 LIF 32 route 44

# D

Data ports default assignments 8 displaying DNS domain configurations 72 displaying (continued) failover groups 73 host name entries 72 interface groups 69 load balancing zones 75 network information 67 network ports 67 routing groups 71 static routes 71 VLANs 68 DNS configuration 47 DNS domain configurations displaying 72 DNS domains managing 48 DNS load balancing 51, 56 DNS load balancing zone about 52 creating 53

# Ε

enabling, on the cluster 24 Ethernet ports default assignments 8

# F

failover disabling, of a LIF 40 enabling, of a LIF 40 failover groups clusterwide 37 configuring 37 creating or adding entries 39 deleting 39 displaying information 73 LIFs, relation 38 removing ports from 39 renaming 39 system-defined 37 types 37 user-defined 37 failover targets viewing 74

# G

guidelines cluster interconnect cabling 2 creating LIFs 31

# Η

host name entries displaying 72 viewing 72 host-name resolution 47 host-name resolution (*continued*) admin Vserver 47 configuring 47 hosts table 48 hosts table managing 48

# 

interface group dynamic multimode 9, 12 load balancing 13 load balancing, IP address based 13 load balancing, MAC address based 13 single-mode 9 static multimode 9, 10 types 9 interface groups 9 creating 14 deleting 15 ports 9 ports, adding 15 ports, displaying information 69 ports, removing 15 ports, restrictions on 13 interfaces logical, deleting 36 logical, modifying 34 logical, reverting to home port 35 IPv6 23, 24 supported features 23 unsupported features 23 IPv6 addresses guidelines 31 IPv6 addresses creating 31

# L

LACP (Link Aggregation Control Protocol) 12 LIF failover disabling 40 enabling 40 scenarios causing 37 LIFs about 25 characteristics 28 cluster 26 cluster-management 26 configuring 25 creating 32 data 26 deleting 36 DNS load balancing 52 failover groups, relation 38 load balancing weight, assigning 52 maximum number of 30 migrating 34 modifying 34 node-management 26 reverting to home port 35 roles 26 viewing information about 70 viewing, failover targets displaying, failover-targets 74

limits LIFs 30 Link Aggregation Control Protocol (LACP) 12 load balancing 51 IP address based 13 MAC address based 13 multimode interface groups 13 round-robin 13 types 51 load balancing weight assigning 51 load balancing zone adding a LIF 54 displaying 75 removing a LIF 54

# Μ

Management ports default assignments 8 managing DNS host name entries 48 MIB /etc/mib/iscsi.mib 59 /etc/mib/netapp.mib 59 custom mib 59 iSCSI MIB 59 migrating LIFs 34 monitoring DNS domain configurations 72 failover groups 73 host name entries 72 interface groups 69 load balancing zones 75 network connectivity 81 network information 67 network ports 67 routing groups 71 static routes 71 VLANs 68 multimode interface groups load balancing, IP address based 13 load balancing, MAC address based 13 load balancing, port-based 13 load balancing, round-robin 13

# Ν

network cabling guidelines 2 network configuration 1 cluster setup 3 data Vserver admin Vserver 5 network connectivity discovering 81 network problems commands for diagnosing 81 network traffic optimizing, Cluster-Mode 51 networking components cluster 1 networks ports 7 NFS 56 NIC removing 21

notices 89 Notices 89

# 0

OID 59

# Ρ

port role cluster 8 data 8 node-management 8 ports concepts 7 displaying 67 failover groups 37 ifgrps 9 interface groups 9 interface groups, adding ports 15 interface groups, creating 14 interface groups, displaying information 69 interface groups, removing ports 15 interface groups, restrictions on 13 managing 7 modifying attributes 20 naming conventions 7 roles 8

# R

route creating 44 routes static, deleting 45 static, displaying information about 71 routing managing 43 routing groups 43 static routes 43 routing groups creating 43 deleting 44 displaying information 71

# S

setup network configuration 3 Simple Network Management Protocol See SNMP **SNMP** agent 59 authKey security 62 authNoPriv security 62 authProtocol security 62 cluster Vserver 59 commands 65 configuring traps 64 configuring v3 users 61 example 62 MIBs 59 noAuthNoPriv security 62 security parameters 62 storage system 59

SNMP (continued) traps 59 traps, types 64 SNMP community creating 60 SNMP traps built-in 64 snmpwalk 62 static routes deleting 45 displaying information about 71

# T

trademark information 88 traps configuring 64

# V

virtual LANs creating 19 deleting 20 displaying information 68 managing 16 **VLANs** advantages of 18 creating 19 deleting 20 displaying information 68 managing 16 membership 17 MTU size 20 tagged traffic 19 tagging 16 untagged traffic 19



NA 210-06186\_A0, Printed in USA

SC27-6278-00

