# IBM System Storage N series
## Data ONTAP 8.0 7-Mode Network Management Guide

# Contents

# About this guide

This guide applies to systems, including systems with gateway functionality, running Data ONTAP 8.*x* 7-Mode. In the Data ONTAP 8.*x* 7-Mode product name, the term *7-Mode* signifies that the 8.*x* release has the same features and functionality found in the prior Data ONTAP 7.1, 7.2, and 7.3 release families.

**Note:** 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®, Hitachi Data Systems®, and EMC®. 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, which is accessed and navigated as described in

# 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 services

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 that you have taken these steps to try to solve the problem yourself:

- Check all cables to make sure that they are connected properly.
- Check the power switches to make sure that the system is turned on.
- Use the troubleshooting information in your system documentation and use the diagnostic tools that come with your system.

## Using the documentation

Information about N series hardware products is available in printed documents and a documentation CD that comes with your system. The same documentation is available as PDF files on the IBM NAS support Web site:

*www.ibm.com/storage/support/nas/*

Data ONTAP software publications are available as PDF files on the IBM NAS support Web site:

*www.ibm.com/storage/support/nas/*

## Hardware service and support

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

*www.ibm.com/planetwide/*

# Firmware updates

As with all devices, it is recommended that you run the latest level of firmware, which can be downloaded by visiting the following Web site:

*www.ibm.com/storage/support/nas/*

Verify that the latest level of firmware is installed on your machine before contacting IBM for technical support. See the *Data ONTAP Upgrade Guide* for your version of Data ONTAP for more information on updating firmware.

# 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 e-mail 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

# Network interfaces on your storage system

Your storage system supports physical network interfaces, such as Ethernet and Gigabit Ethernet interfaces, and virtual network interfaces, such as interface group and virtual local area network (VLAN). Each of these network interface types has its own naming convention.

Your storage system supports the following types of physical network interfaces:

- 10/100/1000 Ethernet
- Gigabit Ethernet (GbE)
- 10 Gigabit Ethernet

In addition, some storage system models have a physical network interface named e0M. It is a low-bandwidth interface of 100 Mbps and is used only for Data ONTAP management activities, such as running a Telnet, SSH, or RSH session.

**Related concepts**

*Network interface configuration* on page 37
*How interface groups work in Data ONTAP* on page 101
*How VLANs work* on page 87

## Network interface naming

Network interface names are based on whether the interface is a physical or virtual network interface. Physical interfaces are assigned names based on the slot number of the adapter. Interface group names are user specified. VLANs are named by combining the interface name and VLAN ID.

Physical interfaces are automatically assigned names based on the slot where the network adapter is installed. Because physical interfaces are Ethernet interfaces, they are identified by a name consisting of "e," the slot number of the adapter, and the port on the adapter (if multi-port adapter). A multiport adapter has letters or numbers imprinted next to its ports.

- `e<slot_number>` if the adapter or slot has only one port
- `e<slot_number><port_letter>` if the adapter or slot has multiple ports

Interface group names are user specified. An interface group's name should meet the following criteria:

- It must begin with a letter.
- It must not contain any spaces.
- It must not contain more than 15 characters.
- It must not already be in use by another interface or interface group.

VLAN interface names are in the following format:

- `<physical_interface_name>-<vlan_ID>`
- `<ifgrp_name>-<vlan_ID>`

The following table lists interface types, interface name formats, and example of names that use these identifiers.

| Interface type | Interface name format | Examples of names |
|---|---|---|
| Physical interface on a single-port adapter or slot | `e<slot_number>` | e0 e1 |
| Physical interface on a multiple-port adapter or slot | `e<slot_number><port_letter>` | e0a e0b e0c e0d e1a e1b |
| Interface group | Any user-specified string that meets certain criteria | web_ifgrp ifgrp1 |
| VLAN | `<physical_interface_name>-<vlan-ID>` or `<ifgrp_name>-<vlan_ID>` | e8-2 ifgrp1-3 |

### Host names

When you run the `setup` command on a storage system for the first time, Data ONTAP creates a host name for each installed interface by appending the interface name to the host name of the storage system.

**Note:** The interface host names are not advertised by DDNS, but are available in the /etc/hosts file.

The following table shows examples of host names appended with the interface names.

| Interface type | Host name |
|---|---|
| Single-port Ethernet interface in slot 0 | toaster-e0 |
| Quad-port Ethernet interface in slot 1 | toaster-e1a toaster-e1b toaster-e1c toaster-e1d |

# Maximum number of network interfaces

Beginning with Data ONTAP 7.3, storage systems can accommodate from 256 to 1,024 network interfaces per system, depending on the storage system model, system memory, and whether they are in an HA pair.

You should run the `sysconfig` command and check the Memory size field displayed for the slot 0 system board of the storage system to determine your storage system memory.

The number of physical interfaces depends on the storage system model. Each storage system can support up to 16 interface groups. The maximum number of VLANs that can be supported equals the maximum number of network interfaces shown in the following table minus the total number of physical interfaces, interface groups, vh, and loopback interfaces supported by the storage system.

The maximum number of network interfaces that each system can support is shown in the following table. The total number of interfaces can include physical, interface group, VLAN, vh, and loopback interfaces.

| Storage system memory | Maximum number of network interfaces |
|---|---|
| 2 GB or less | 128 |
| 2 GB or less in an HA pair | 256 |
| 6 GB or less | 256 |
| 6 GB or less in an HA pair | 512 |
| More than 6 GB | 512 |
| More than 6 GB in an HA pair | 1,024 |

**Related references**

# The e0M interface

Some storage system models have an interface named e0M. The e0M interface is dedicated to Data ONTAP management activities. It enables you to separate management traffic from data traffic on your storage system for security and throughput benefits.

On a storage system that has the e0M interface, the Ethernet port (indicated by a wrench icon on the rear of the chassis) is connected to an internal Ethernet switch. The internal Ethernet switch provides connectivity to the e0M interface and the remote management device such as the RLM or the BMC. The following diagram illustrates the connections.

When you set up a system that includes the e0M interface, the Data ONTAP setup script recommends that you use the e0M as the preferred management interface for environments that use dedicated LANs to isolate management traffic from data traffic. The setup script then prompts you to configure e0M. The e0M configuration is separate from the configuration of the remote management device. Both configurations require unique IP addresses to allow the Ethernet switch to direct traffic to either the e0M interface or the remote management device. For information about how to set up the e0M interface, see the *Data ONTAP 7-Mode Software Setup Guide*.

The e0M interface is a low-bandwidth, 100 BT (100 Mbps) interface. It does not support interface groups, VLANs, and jumbo frames.

After you have set up the e0M interface, you can use it to access the storage system with the following management protocols, if they have been enabled:

- Telnet
- RSH
- HTTP or HTTPS
- SSH
- SNMP

**Related concepts**

*Guidelines for configuring e0M* on page 25

**Related tasks**

*Blocking data traffic on e0M* on page 46

## Guidelines for configuring e0M

In storage systems with e0M, you should ensure that e0M is not configured to serve data traffic so that system performance is not affected. Therefore, to achieve optimum system performance, you must consider certain guidelines when configuring e0M.

The guidelines for configuring e0M are as follows:

- e0M and high bandwidth interfaces should not be configured in the same subnet.
- e0M should be configured on an isolated subnet that is not used for data traffic.
  If it is not possible to configure e0M in a dedicated management subnet, then system management activities should be performed through a high-bandwidth interface and e0M should be disabled.
- A default route gateway should not be configured on the same subnet as that of e0M.

# How to use the RLM or BMC to manage Data ONTAP remotely

You can manage your storage system locally from an Ethernet connection by using any network interface. However, to manage your storage system remotely, the system should have a Remote LAN Module (RLM) or Baseboard Management Controller (BMC). These provide remote platform management capabilities, including remote access, monitoring, troubleshooting, and alerting features.

If your data center configuration has management traffic and data traffic on separate networks, you can configure the RLM or the BMC on the management network.

With the RLM, you can remotely access the storage system in the following ways:

- Through the serial console
  The RLM is directly connected to the storage system through the serial console. You use the Data ONTAP CLI to administer the storage system and the RLM.
- Through an Ethernet connection using a secure shell client application
  You use the RLM CLI to monitor and troubleshoot the storage system.

With the BMC, you can access the storage system in the following ways:

- Through the serial console
  You use the Data ONTAP CLI to administer the storage system and the BMC.
- Through an Ethernet connection by using a secure shell client application
  You use the BMC CLI to monitor and troubleshoot the storage system.

For more information about the RLM and the BMC, see the *Data ONTAP 7-Mode System Administration Guide*.

## Ways to configure the RLM

Before using the RLM, you must configure it for your storage system and network. You can configure the RLM when setting up a new storage system with RLM already installed, after setting

up a new storage system with RLM already installed, or when adding an RLM to an existing storage system.

You can configure the RLM by using one of the following methods:

* Initializing a storage system that has the RLM pre-installed
  When the storage system setup process is complete, the `rlm setup` command runs automatically. For more information about the entire setup process, see the *Data ONTAP 7-Mode Software Setup Guide*.

* Running the Data ONTAP `setup` script
  The `setup` script ends by initiating the `rlm setup` command.

* Running the Data ONTAP `rlm setup` command
  For information about using the `rlm setup` command to configure the RLM, see the *Data ONTAP 7-Mode System Administration Guide*.

When the `rlm setup` script is initiated, you are prompted to enter network and mail host information.

## Ways to configure the BMC

Before using the BMC, you must configure it for your storage system and network. You can configure the BMC when setting up a new storage system with BMC already installed or after setting up a new storage system with BMC already installed.

You can configure the BMC by using one of the following methods:

* Initializing a storage system that has the BMC
  When the storage system setup process is complete, the `bmc setup` command runs automatically. For more information about the entire setup process, see the *Data ONTAP 7-Mode Software Setup Guide*.
* Running the Data ONTAP `setup` script
  The `setup` script ends by initiating the `bmc setup` command.
* Running the Data ONTAP `bmc setup` command
  For information about using the `bmc setup` command to configure the BMC, see the *Data ONTAP 7-Mode System Administration Guide*.

When the `bmc setup` script is initiated, you are prompted to enter network and mail host information.

# How TSO increases outbound throughput

TCP segmentation offload (TSO), or Large send offload (LSO), is a technique for reducing the CPU usage of the host system in high-bandwidth outbound network connections. In TSO, data segmentation is offloaded to the NIC that divides the data into the default maximum transmission unit (MTU) size of the outgoing interface.

When a large amount of data is transmitted to a destination system, TCP in the host system segments the data into MTU sized packets. For each segment, a header buffer is allocated and a TCP header is inserted. The packets are sent down through the layers of the TCP stack, and each layer adds its own header to the packet. Processing all the created segments at each layer in the protocol stack can burden the CPU of the controller. This leads to the wastage of CPU cycles in the host system and causes an overhead to the host CPU. However, using TSO, this overhead can be passed down to the NIC to improve CPU efficiency of the host.

## Data ONTAP support for TSO

In Data ONTAP 8.0.1 and later releases, TSO is supported on all Data ONTAP versions. If fast path is enabled on the storage system, TSO is also enabled on all the network interfaces by default. TSO is also enabled on VLANs and any type of interface group if the underlying network interfaces support TSO.

TSO is automatically disabled under the following conditions:

* In an interface group, if any of the constituent network interfaces do not support TSO
* In rate-limited connections such as SnapMirror with the throttle option enabled
  For more information about the `replication.throttle.enable` option, see the options(1) man page.

## Viewing the TSO statistics

Viewing the TSO statistics enables you to find out whether TSO is functioning and gives you an estimate of the saved CPU cycles.

**Step**

1. To view the TSO statistics, enter the following command:

   `netstat -p tcp`

---

**Example**

The example shows a part of the `netstat -p tcp` command output.

```
49759833 segments sent using TSO
1637084830596 bytes sent using TSO
0 TSO segments truncated
0 TSO wrapped sequence space segments
```

The `netstat -p tcp` counters related to TSO are listed in the order of appearance in the output.

* Segments sent using TSO
  This counter indicates the number of segments sent by the application before the NIC performs TSO.

---

For example, when 64 KB of data is transmitted, the TSO segments counter is incremented by 1.

- Bytes sent using TSO

  This counter indicates the number of bytes of segments sent by the application before the NIC performs TSO.

  For example, when 64 KB of data is transmitted, the TSO bytes counter is incremented by 64 KB.

- TSO segments truncated

  This counter indicates the number of times a segment has to be truncated because an application sends a buffer that is larger than the maximum size of the packet transmitted by NIC.

  For example, if the NIC supports a maximum segment size of 48 KB and an application sends 64 KB, then the buffer breaks the larger segment in two TSO segments. The first segment is 48 KB and the second segment is 16 (64 - 48) KB. This ensures that the NIC card does not transmit a large buffer.

- TSO wrapped sequence space segments

  This counter indicates the number of times a TSO segment wraps the 32-bit TCP sequence space.

# How Data ONTAP uses ACP to increase storage availability

ACP, or Alternate Control Path, is a protocol that enables Data ONTAP to manage and control a SAS disk shelf storage subsystem. It uses a separate network (alternate path) from the data path, so management communication is not dependent on the data path being intact and available.

You do not need to actively manage the SAS disk shelf storage subsystem. Data ONTAP automatically monitors and manages the subsystem without operator intervention. However, you must provide the required physical connectivity and configuration parameters to enable the ACP functionality.

**Note:** You can install SAS disk shelves without configuring ACP. However, for maximum storage availability and stability, you should always have ACP configured and enabled.

After you enable ACP, you can use the storage show acp and acpadmin list all commands to display information about your ACP subsystem.

Because ACP communication is on a separate network, it does not affect data access in any way.

## Enabling ACP

ACP can increase your storage availability when you use SAS disk shelves. If your storage system model has a dedicated port for ACP, then ACP is enabled by default, and you do not need to explicitly enable ACP.

**Before you begin**

- Is the ACP subnet cabled on an isolated network, with no switches or hubs?
  For more information, see the *Hardware and Service Guide* for your disk shelf.
- Have you identified a port that is not in use by any other subsystem?
- If you are configuring ACP for disk shelves attached to an HA pair, have you recorded the domain name and network mask to ensure that they are the same for both nodes?

**About this task**

The ACP subnet is a private Ethernet network that enables the ACP processor in the SAS module to communicate both with Data ONTAP and with the SAS IOMs in the disk shelves.

The ACP subnet is separate from the I/O data path that connects the disk shelves to the HBA on the storage controller. When you configure ACP on one of the system's network interfaces, you must supply a private domain name that conforms to the standard for private internet addresses (RFC1918). You can use the system default domain or another network name (that is, an IP address ending in 0) that conforms to the standard.

**Steps**

1. If your system does not have a dedicated port for ACP (e0p), ensure that the port you are assigning to ACP is not in use by any other subsystem by reviewing your `/etc/rc` file and entering the following command:

   **`ifconfig interface_name`**

   The interface you use for ACP should not be part of an Interface Group, and it should have no VLANs or IP addresses configured on it.

2. At the Data ONTAP command line, enter the following command:

   **`acpadmin configure`**

   If you have not previously configured the networking information for ACP, you are prompted for that information. When you select a domain name and network mask for the ACP interface, Data ONTAP automatically assigns IP addresses for the ACP interface on the storage controller and both I/O modules on each disk shelf on the ACP subnet.

3. You can verify your ACP connectivity by entering the following command:

   **`sysconfig -v`**

   The ACP Connectivity Status should show "Full Connectivity".

**Example**

For example, if you select e0P as the interface for ACP traffic, 192.168.0.0 as the ACP
domain, and 255.255.252.0 as the network mask for the ACP subnet, the `storage show acp`
command output looks similar to the following example:

```
my-sys-1> storage show acp

 Alternate Control Path:  enabled
 Ethernet Interface:      e0P
 ACP Status:              Active
 ACP IP address:          192.168.2.61
 ACP domain:              192.168.0.0
 ACP netmask:             255.255.252.0
 ACP Connectivity Status: Full Connectivity

Shelf Module       Reset Cnt    IP address       FW Version    Module
Type    Status
--------------------------------------------------------------------
-----------
7a.001.A           002          192.168.0.145    01.05
IOM6         active
7a.001.B           003          192.168.0.146    01.05
IOM6         active
7c.002.A           000          192.168.0.206    01.05
IOM6         active
7c.002.B           001          192.168.0.204    01.05
IOM6         active
```

**After you finish**

## The ACP subnet

The ACP subnet is a private Ethernet network that enables the ACP processor in the SAS module to
communicate both with Data ONTAP and with the SAS IOMs in the disk shelves.

The ACP subnet is separate from the I/O data path that connects the disk shelves to the HBA on the
storage controller. When you configure ACP on one of the system's network interfaces, you must
supply a private domain name that conforms to the standard for private internet addresses
(RFC1918). You can use the system default domain, 198.15.1.0, or another network name (that is, an
IP address ending in 0) that conforms to the standard.

If you are configuring ACP for disk shelves attached to an HA pair, you must supply the same ACP
domain name and network mask for both systems.

> **Attention:** Do not connect the ACP port to a routed network, and do not configure switches or
> hubs between the ACP port and the designated Ethernet port. Doing so is not supported and causes
> interruptions in service.

After you select a domain name and network mask for the interface, Data ONTAP automatically assigns IP addresses for the ACP interface on the storage controller and both I/O modules on each disk shelf on the ACP subnet.

You can use the `sysconfig -v` command to check whether your ACP subnet is cabled correctly. If ACP is disabled, `sysconfig` shows ACP connectivity as "NA".

# Standards and characteristics of Ethernet frames

Frame size and Maximum Transmission Unit (MTU) size are the two important characteristics of an Ethernet frame. The standard Ethernet (IEEE 802.3) frame size is 1,518 bytes. The MTU size specifies the maximum number of bytes of data that can be encapsulated in an Ethernet frame.

The frame size of a standard Ethernet frame (defined by RFC 894) is the sum of the Ethernet header (14 bytes), the payload (IP packet, usually 1,500 bytes), and the Frame Check Sequence (FCS) field (4 bytes). You can change the default frame size on Gigabit Ethernet network interfaces.

The MTU size specifies the maximum payload that can be encapsulated in an Ethernet frame. For example, the MTU size of a standard Ethernet frame is 1,500 bytes; this is the default for storage systems. However, a jumbo frame, with an MTU size of 9,000 bytes, can also be configured.

## What jumbo frames are

Jumbo frames are larger than standard frames and require fewer frames. Therefore, you can reduce the CPU processing overhead by using jumbo frames with your network interfaces. Particularly, by using jumbo frames with a Gigabit or 10 Gigabit Ethernet infrastructure, you can significantly improve performance, depending on the network traffic.

Jumbo frames are packets that are longer than the standard Ethernet (IEEE 802.3) frame size of 1,518 bytes. The frame size definition for jumbo frames is vendor-specific because jumbo frames are not part of the IEEE standard. The most commonly used jumbo frame size is 9,018 bytes.

Jumbo frames can be used for all Gigabit and 10 Gigabit Ethernet interfaces that are supported on your storage system. The interfaces must be operating at or above 1,000 Mbps.

You can set up jumbo frames on your storage system in the following two ways:

- During initial setup, the `setup` command prompts you to configure jumbo frames if you have an interface that supports jumbo frames on your storage system.
- If your system is already running, you can enable jumbo frames by setting the MTU size on an interface.

## Network interface requirements for jumbo frames

Before you enable jumbo frames on your storage system, jumbo frames must be enabled for the switch ports, client interfaces, and intermediate routers on the network. If your storage system and the client are on different subnets, the next-hop router must be configured for jumbo frames.

## Guidelines to configure clients for jumbo frames

When configuring clients for jumbo frames, you should verify certain configurations, such as the TCP window size of the client and the MTU size of the client, storage system, and any intermediate subnet.

The guidelines for configuring clients for jumbo frames are as follows:

- Configure jumbo frames on the client and on your storage system.
  Find how to configure jumbo frames on your client by checking the network adapter documentation for your client.
- Enlarge the client's TCP window size.
  The minimum value for the client's window size should be two times the MTU size, minus 40, and the maximum value can be the highest value your system allows. Typically, the maximum value you can set for your client's TCP window is 65,535. If your storage system is configured to support jumbo frames and the client is not, the communication between the storage system and the client occurs at the client's frame size.
- Configure the storage system and UDP clients to have the same MTU size.
  User Datagram Protocol (UDP) systems do not negotiate the MTU size. If your storage system and clients do not have the same MTU size, the storage system might send packets that the clients cannot receive.
- Check the MTU size of any intermediate subnets if your storage system and the client are on different subnets.
  If the storage system and the client (both configured to use jumbo frames) are on different subnets and an intermediate subnet does not support jumbo frames, the intermediate router fragments the IP packets and the advantages of using jumbo frames are lost.

**Related tasks**

*Specifying an MTU size for a network interface* on page 40

# Flow control

Flow control enables you to manage the flow of frames between two directly connected link-partners. Flow control can reduce or eliminate dropped packets due to overrun.

To achieve flow control, you can specify a flow control option that causes packets called Pause frames to be used as needed. For example, link-partner A sends a Pause On frame to link-partner B

when its receive buffers are nearly full. Link-partner B suspends transmission until it receives a Pause Off frame from link-partner A or a specified timeout threshold is reached.

# Network interface configuration

Configuring network interfaces involves assigning IP addresses, setting network parameters and hardware-dependent values, specifying network interfaces, and viewing your storage system's network configuration.

When you configure network interfaces, you can do any or all of the following:

- Assign an IP address to a network interface.
- Set parameters such as network mask, broadcast address.
- Set hardware-dependent values such as media type, MTU size, and flow control.
- Specify whether the interface should be attached to a network with firewall security protection.
- Specify whether the network interface must be registered with Windows Internet Name Services (WINS), if CIFS is running and at least one WINS server has been configured.
- Specify the IP address of an interface or specify the interface name on an HA pair partner for takeover mode.
- View the current configuration of a specific interface or all interfaces that exist on your storage system.

**Related concepts**

## Configuring network interfaces

You can configure network interfaces either during system setup or when the storage system is operating. When the storage system is operating, you can use the ifconfig command to assign or modify configuration values of your network interfaces.

During system setup, you can configure the IP addresses for the network interfaces. An ifconfig command is included in the /etc/rc file of the root volume for each network interface that you configured during the system setup. After your storage system has been set up, the ifconfig commands in the /etc/rc file are used to configure the network interfaces on subsequent storage system reboots.

You can use the ifconfig command to change values of parameters for a network interface when your storage system is operating. However, such changes are not automatically included in the /etc/rc file. If you want your configuration modifications to be persistent after a reboot, you must include the ifconfig command values in the /etc/rc file.

## Configuring an IP address for a network interface

You can configure IP addresses for your network interface during system setup. To configure the IP addresses later, you should use the ifconfig command.

### About this task

*   Network configuration changes made by using the ifconfig command are not automatically included in the /etc/rc file. To make the configuration changes persistent after reboots, include the ifconfig command in the /etc/rc file.
*   When you configure an IP address, your storage system creates a network mask based on the class of the address (Class A, B, C, or D) by default.

### Step

1.  To configure an IP address for a network interface, enter the following command:

    **ifconfig *interface_name IP_address***

    *interface_name* is the name of the network interface.

    *IP_address* is the IP address that you want to assign to the network interface.

    ### Example

    To configure a quad-port Ethernet interface e3a to use the IPv4 address 192.0.2.10, enter the following command:

    **ifconfig e3a 192.0.2.10**

### Related tasks

## Specifying a subnet mask for a network interface

You must specify a subnet mask if you have created subnets that do not match the class boundary of the IPv4 address of the network interface. You can specify a subnet mask for a network interface by using the ifconfig command.

### Before you begin

The IP address assigned to e0M must not belong to the same subnet as the IP address configured to serve data traffic.

**Note:** An EMS message is displayed if the management interface and network interfaces belong to the same subnet.

**About this task**

Data ONTAP allows you to configure a 32-bit subnet mask with all bits equal to 1.

**Steps**

**1.** To specify a subnet mask for an interface, enter the following command:

**ifconfig *interface_name* netmask *mask***

*interface_name* is the name of the network interface.

*mask* is the subnet mask.

**Example**
To configure a 24-bit mask for the interface e3a that you have already configured, enter the following command:

**ifconfig e3a netmask 255.255.255.0**

**2.** To change the subnet mask for an interface that has been configured with a primary and an alias address, enter the following command for each IP address:

**ifconfig *interface_nameIP address* netmask *mask***

**Example**
To change the subnet mask of the primary IP address for the interface e3a, enter the following command:

**ifconfig e3a 172.25.206.110 netmask 255.255.255.0**

To change the subnet mask of the alias IP address for e3a, enter the following command:

**ifconfig e3a 120.120.1.1 netmask 255.255.255.0**

**Related concepts**

*Guidelines for configuring e0M* on page 25

**Related tasks**

*Configuring an IP address for a network interface* on page 38

## Specifying a media type for a network interface

You can specify a media type for configuring the speed and duplex of a network interface by using the ifconfig command.

**Step**

**1.** To specify a media type, enter the following command:

**ifconfig *interface_name* mediatype *type***

*interface_name* is the name of the network interface.

*type* specifies the Ethernet media type used. The possible values are `tp`, `tp-fd`, `100tx`, `100tx-fd`, `auto`, or `10g-sr`.

For more information, see the na_ifconfig(1) man page.

**Example**
To configure the interface e2a as a 100Base-TX full-duplex interface, enter the following command:

**`ifconfig e2a mediatype 100tx-fd`**

## Specifying an MTU size for a network interface

The maximum transmission unit (MTU) size is used to specify the jumbo frame size on Gigabit Ethernet interfaces. You can specify the MTU size for transmission between your storage system and its client by using the `ifconfig` command.

**Step**

**1.** To specify an MTU size, enter the following command:

**`ifconfig interface_name mtusize size`**

*interface_name* is the name of the network interface.

*size* is the MTU to be used for the network interface.

**Example**
To specify an MTU size of 9000 for Gigabit Ethernet interface e8, enter the following command:

**`ifconfig e8 mtusize 9000`**

**Related concepts**

*Standards and characteristics of Ethernet frames* on page 33
*What jumbo frames are* on page 33
*Guidelines to configure clients for jumbo frames* on page 34

## Specifying the flow control type for a network interface

You can specify the flow control type for a network interface to manage the flow of frames between two directly connected link-partners by using the `ifconfig` command. By controlling the transmission rate of data packets between the link-partners, you can effectively manage the network traffic.

**About this task**

- Depending on the network infrastructure, you must specify the flow control settings for a network interface.

  For example, if you set the flow control type on the network interface to **full**, the switch (link-partner) should also be able to send and receive flow control frames.

- The configured flow control setting is advertised during autonegotiation. If autonegotiation succeeds, the operational flow control setting is determined based on the negotiated speed and the value advertised by the other device. If autonegotiation fails, the configured flow control setting is used.

  **Note:** Only 1 Gigabit interfaces support autonegotiation.

- The flow control settings of all the physical interfaces that constitute an interface group and VLANs must be the same.

**Step**

1. To specify the flow control type, enter the following command:

   **ifconfig *interface_name* flowcontrol *value***

   *interface_name* is the name of the network interface.

   *value* is the flow control type. You can specify the following values for the flowcontrol option:

   | | |
   |---|---|
   | **none** | No flow control |
   | **receive** | Able to receive flow control frames |
   | **send** | Able to send flow control frames |
   | **full** | Able to send and receive flow control frames |

   The default flow control type is full.

   **Example**
   To turn off flow control on interface e0a, enter the following command:

   **ifconfig e0a flowcontrol none**

**Related concepts**

*Flow control* on page 34

# Specifying whether a network interface is trusted

You can specify whether a network interface is trustworthy or untrustworthy. When you specify an interface as untrusted (untrustworthy), any packets received on the interface are likely to be dropped. For example, if you run a ping command on an untrusted interface, the interface drops any ICMP response packet received.

**About this task**

Applications such as NFS, CIFS, HTTP can choose to accept packets only from trusted interfaces. If the destination interface is set as untrusted, it can receive packets from untrusted interfaces. Otherwise, the packets from untrusted interfaces are dropped. By default, only HTTP allows receiving packets from untrusted interfaces.

**Step**

1. To specify a network interface as trusted or untrusted, enter the following command:

   **`ifconfig interface_name {trusted|untrusted}`**

   `interface_name` is the name of the network interface.

   `trusted` specifies that the network interface is to be trusted.

   `untrusted` specifies that the network interface is not to be trusted.

   **Example**
   To specify that the network attached to interface e8 is not to be trusted for firewall security, enter
   the following command:

   **`ifconfig e8 untrusted`**

# Configuring a partner interface in an HA pair

To prepare for a successful takeover in an HA configuration, you can map a network interface to an
IP address or to another network interface on the partner node. During a takeover, the network
interface on the surviving node assumes the identity of the partner interface.

**Before you begin**

When specifying the partner IP address, you must ensure that both the local network interface and the
partner's network interface are attached to the same network segment or network switch.

**About this task**

- If the network interface is an interface group, the partner interface must be denoted by an
  interface name and not an IP address.
  The partner interface can be an interface group or a physical network interface.
- You cannot specify the underlying physical ports of an interface group in a partner configuration.
- For the partner configuration to be persistent across reboots, you must include the `ifconfig`
  command in the `/etc/rc` file.
  For a successful takeover in both directions, you must repeat the partner configuration in the `/etc/rc` files of each node.
- When specifying the partner interface name, you can configure the interfaces symmetrically, for
  example map interface e1 on one node to interface e1 on the partner node.
  Though symmetrical configuration is not mandatory, it simplifies administration and
  troubleshooting tasks.

**Step**

1. Depending on the partner configuration that you want to specify, enter the following command:

| If you want specify a... | Enter the following command... |
|---|---|
| Partner IP address | **ifconfig *interface_name* partner *address*** |
| | *interface_name* is the name of the network interface. |
| | *address* is the partner IP address. |
| Partner interface name | **ifconfig *interface_name* partner *partner_interface*** |
| | *partner_interface* is the name of the partner network interface. |

---

**Example: Specifying a partner IP address and partner interface name**

Consider node1 and node2 are two storage systems in an HA configuration

If the IP address of the interface e8 on node2 is 198.9.200.38, the following command allows the interface e1 of node1 to take over the IP address of node2 for the duration of the takeover:

```
node1> ifconfig e1 partner 198.9.200.38
```

Instead of specifying the IP address, you can also specify the partner interface name. The following command allows the interface e1 of node1 to assume the identity of e8 of node2 for the duration of the takeover:

```
node1> ifconfig e1 partner e8
```

---

## Enabling or disabling automatic takeover for a network interface

You can enable or disable negotiated failover for a network interface to trigger automatic takeover if the interface experiences a persistent failure. You can use the `nfo` option of the `ifconfig` command to enable or disable negotiated failover.

**About this task**

- You must include the `nfo` option in the `/etc/rc` file for it to persist across reboots.
- You can specify the `nfo` option for an interface group. However, you cannot specify the `nfo` option for any underlying physical interface of the interface group.

**Steps**

1. To enable takeover on interface failures, enter the following command:

   **options cf.takeover.on_network_interface_failure on**

2. To enable or disable negotiated failover, enter the following command:

   **ifconfig *interface_name* {nfo|-nfo}**

   *interface_name* is the name of the network interface.

nfo—Enables negotiated failover

-nfo—Disables negotiated failover

**Example**

To enable negotiated failover on the interface e8 of an HA pair, enter the following command:

**ifconfig e8 nfo**

## Removing a primary IP address from a network interface

You can remove a primary IP address from a network interface to disconnect the network interface from the network or reconfigure the network interface.

**Before you begin**

Ensure that you remove all the manually configured alias addresses for the interface.

**Step**

1. To remove a primary IP address, enter the following command:

   **ifconfig *interface_name* 0**

   *interface_name* is the name of the network interface.

   Alternatively, to remove a primary IPv4 address, you can use the following command:

   **ifconfig *interface_name* 0.0.0.0**

   **Example**

   To remove the primary address of the interface e3, enter the following command:

   **ifconfig e3 0**

**Related tasks**

[Creating or removing aliases](#) on page 44

## Creating or removing aliases

You can create or remove an alias if you are changing the IP address of an interface. You should use the alias command to create an alias IP address, and use the -alias command to remove an alias IP address.

**About this task**

- The alias addresses are lost when the storage system reboots. If you want to make these changes persistent, include the ifconfig commands, which are used for configuring the alias addresses, in the /etc/rc file for the root volume.
- You cannot use FilerView to manage aliases.
- For IPv4 addresses, you can add an alias address only if a primary address for the interface exists.

**Step**

**1.** To create or remove an alias, enter the following command:

`ifconfig interface_name [-]alias address`

**Example**

The following example creates the alias IP address 192.0.2.30 for the interface e0 (already configured with IP address 192.0.2.21):

`ifconfig e0 alias 192.0.2.30`

The following example removes the 192.0.2.30 alias for the interface e0 specified in the previous example:

`ifconfig e0 -alias 192.0.2.30`

## Changing the status of an interface

You must make an interface inactive before performing tasks such as upgrading an interface, disabling a failed interface, or troubleshooting connectivity issues. You must again make the interface active after you complete the task. You can make an interface active or inactive by using the `ifconfig` command.

**Step**

**1.** To change the status of an interface, enter the following command:

`ifconfig interface {up|down}`

`up`—makes the interface active

`down`—makes the interface inactive

# Blocking or unblocking data traffic on the network interfaces

For network security and efficient traffic management, you can block or allow certain types of data traffic, including SnapMirror transfers, SnapVault transfers, and data transfers that use the CIFS, NFS, and NDMP protocols from selected or all network interfaces.

**About this task**

- If you block any data protocol on an interface by using the `interface.blocked.protocol` option, any client requests using that particular data protocol on that interface fails. In such cases, you can either unblock the data protocol on that interface or use a different interface for the client traffic.
- You can block protocols from all interfaces including interface groups and VLANs.

**Step**

**1.** To block or unblock protocols from network interfaces, perform one of the following steps:

| If you want to... | Enter the following command... |
| --- | --- |
| Block a protocol from multiple network interfaces | **options interface.blocked.*protocol_name* *interface_name,interface_name, ...interface_name*** <br><br> *protocol_name* is the protocol that you want to block. <br><br> *interface_name* is the interface on which you want to block the protocol. <br><br> **Note:** To block multiple protocols from a single interface, you must repeat the command for each protocol. |
| Unblock a protocol | **options interface.blocked.*protocol_name* ""** |

**Example**

To block the interface e9 from using the CIFS protocol, enter the following command:

**options interface.blocked.cifs e9**

To block the NDMP protocol from the interfaces e0a and e0b, enter the following command:

**options interface.blocked.ndmp e0a,e0b**

To unblock CIFS from all the network interfaces, enter the following command:

**options interface.blocked.cifs ""**

## Blocking data traffic on e0M

Starting with Data ONTAP 8.0.2, you can block certain types of data traffic on the e0M interface, including SnapMirror transfers, SnapVault transfers, and data transfers that use the CIFS, NFS, and NDMP protocols. This optimizes the system performance by not allowing high-bandwidth data traffic on e0M, which is a low-bandwidth dedicated management interface.

**About this task**

This task needs to be performed for storage systems that are upgraded to Data ONTAP 8.0.2 because the default value of the interface.blocked.mgmt_data_traffic option is off in such systems.

However, in new systems installed with Data ONTAP 8.0.2, the default value of the interface.blocked.mgmt_data_traffic option is on.

**Step**

1. To block all types of data traffic on e0M, enter the following command:

   **`options interface.blocked.mgmt_data_traffic on`**

# Restricting protocol access

If a protocol is enabled for Data ONTAP, you can restrict the protocol's access to the storage system by specifying the host name, IP address, or network interface name.

**Step**

1. At the storage system prompt, enter one of the following commands:

| If you want to restrict a protocol's access to the storage system by using... | Enter... |
|---|---|
| host name or IP address | **`options protocol.access host=[hostname\|IP_address]`** |
| network interface name | **`options protocol.access if=interface_name`** |

- *`protocol`* is the name of the protocol you want to allow access to the storage system. It can be **`rsh`**, **`telnet`**, **`ssh`**, **`httpd`**, **`httpd.admin`**, **`snmp`**, **`ndmpd`**, **`snapmirror`**, or **`snapvault`**.
- *`hostname`* is the name of the host to which you want to allow access by using *`protocol`*.
- *`IP_address`* is the IP address of the host to which you want to allow access by using *`protocol`*.
- *`interface_name`* is the network interface name of the host to which you want to allow access by using *`protocol`*.

  **Note:** If the `telnet.access` option is not set to `legacy`, the `trusted.hosts` option is ignored for Telnet. If the `httpd.admin.access` option is not set to `legacy`, the `trusted.hosts` option is ignored for `httpd.admin`. If the `snapmirror.access` option is not set to `legacy`, the `/etc/snapmirror.allow` file is ignored for SnapMirror destination checking.

For more information about controlling protocol access to a storage system by using multiple host names, IP addresses, and network interfaces, see the na_protocolaccess(8) man page.

For information about NDMP, see the *Data ONTAP 7-Mode Data Protection Tape Backup and Recovery Guide*.

For information about SnapMirror or SnapVault, see the *Data ONTAP 7-Mode Data Protection Online Backup and Recovery Guide*.

**Related tasks**

# Network interface information you can view

You can view the status and performance statistics of your network interfaces, such as packets sent and received, cumulative or continuous packet traffic, collisions and other errors, active sockets, memory buffer, protocol-specific statistics, routing tables.

Data ONTAP provides the following commands that you can use to view network interface information:

| Command | Information displayed |
|---------|----------------------|
| `ifconfig -a` | • Interface status (up or down)<br>• Configuration parameters |
| `ifstat` | • Packets sent and received<br>• Collisions and other errors<br>• Negotiated media type settings between storage system interfaces and link partners |
| `netstat` | • Active sockets for each protocol<br>• Memory buffer (mbuf) pool usage<br>• Protocol-specific statistics for all protocols or a single protocol<br>• Cumulative or continuous packet traffic for all interfaces or a single interface<br>• Routing tables |

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

You can also use FilerView to view interface and routing information.

## Viewing network interface settings

To verify the network connectivity and diagnose any issues, you can view the network interface settings, such as interface status, IP address, and other network parameters. You can view the settings of all interfaces or a specific interface by using the `ifconfig` command.

**Step**

**1.** Depending on the network interface settings that you want to view, complete the following step:

| If you want to view... | Enter the following command... |
|---|---|
| All interfaces | `ifconfig -a` |
| A specific interface | `ifconfig interface_name` |

**Related tasks**

## Viewing or modifying interface settings with FilerView

You can use FilerView to view or modify the settings of all interfaces or a specific interface. The changes made by using FilerView are automatically written to the `/etc/rc` file.

**Steps**

1. From the list on the left pane, click **Network > Manage Interfaces**.

2. Depending on whether you want to view or modify the configuration settings, perform the following step:

| If you want to... | Then... |
|---|---|
| View interface configuration details | Click **Show All Interface Details**. |
| Modify an interface configuration | Select an interface and click **Modify**. |

## Viewing statistics of all active TCP connections

You can view the mapping network context of each TCP connection and the number of bytes of data sent and received over each TCP connection by using the `netstat` command.

**Step**

1. Depending on the statistics that you want to view, perform the following step:

| If you want to view the... | Enter the following command... |
|---|---|
| Mapping context of each TCP connection | `netstat -aM` |
| Number of bytes of data sent and received over each TCP connection | `netstat -aB` |

**Example**

The following example shows the output of the `netstat -aM` command:

```
system1> netstat -aM
Active TCP connections (including servers)
Ctx  Local Address       Remote Address        Swind Send-Q  Rwind Recv-Q State
 lg  *.443               *.*                        0      0      0      0 LISTEN
 lg  *.22                *.*                        0      0      0      0 LISTEN
 lg  *.10568             *.*                        0      0      0      0 LISTEN
```

```
lg  *.10569            *.*                    0      0      0      0 LISTEN
lg  *.10567            *.*                    0      0      0      0 LISTEN
lg  *.10571            *.*                    0      0      0      0 LISTEN
lg  *.8514             *.*                    0      0      0      0 LISTEN
lg  *.514              *.*                    0      0      0      0 LISTEN
lg  *.23               *.*                    0      0      0      0 LISTEN
lg  *.8023             *.*                    0      0      0      0 LISTEN
lg  *.4047             *.*                    0      0      0      0 LISTEN
lg  *.4045             *.*                    0      0      0      0 LISTEN
lg  *.4046             *.*                    0      0      0      0 LISTEN
lg  *.2049             *.*                    0      0      0      0 LISTEN
lg  *.111              *.*                    0      0      0      0 LISTEN
lg  *.28073            *.*                    0      0      0      0 LISTEN
lg  *.32243            *.*                    0      0      0      0 LISTEN
lg  *.22899            *.*                    0      0      0      0 LISTEN
 1  192.168.1.72.2049  192.168.1.36.800   33952    328  26280      0 ESTABLISHED
lg  *.2049             *.*                    0      0      0      0 LISTEN

Active UDP sockets (including servers)
Local Address          Remote Address      Send-Q Recv-Q
*.10570                *.*                      0      0
*.69                   *.*                      0      0
*.161                  *.*                      0      0
*.4049                 *.*                      0      0
*.4047                 *.*                      0      0
*.4045                 *.*                      0      0
*.4046                 *.*                      0      0
*.2049                 *.*                      0      0
*.111                  *.*                      0      0
*.21566                *.*                      0      0
*.520                  *.*                      0      0
```

The following example shows the output of the netstat -aB command:

```
netstat -aB
Active TCP connections (including servers)
Local Address          Remote Address       Swind Send-Q Rwind Recv-Q State
Sent    Rcvd
localhost-20.1023      localhost-10.671     65535      0  8760      0 ESTABLISHED
0      0
localhost-20.8514      localhost-10.626     66608      1  8760      0 ESTABLISHED
1     44
localhost-20.18576     localhost-10.7951    66608      0  8760      0 ESTABLISHED
9284    606K
*.10568                *.*                      0      0     0      0 LISTEN
0      0
*.10569                *.*                      0      0     0      0 LISTEN
0      0
*.10567                *.*                      0      0     0      0 LISTEN
0      0
*.22                   *.*                      0      0     0      0 LISTEN
0      0
*.443                  *.*                      0      0     0      0 LISTEN
0      0
*.8514                 *.*                      0      0     0      0 LISTEN
0      0
*.514                  *.*                      0      0     0      0 LISTEN
0      0
*.23                   *.*                      0      0     0      0 LISTEN
0      0
*.8023                 *.*                      0      0     0      0 LISTEN
0      0
*.32243                *.*                      0      0     0      0 LISTEN
0      0
*.22899                *.*                      0      0     0      0 LISTEN
0      0

Active UDP sockets (including servers)
Local Address          Remote Address       Send-Q Recv-Q
*.10570                *.*                      0      0
```

```
*.69                    *.*                        0        0
*.161                   *.*                        0        0
```

## Viewing or clearing network interface statistics

You can use the `ifstat` command to view the cumulative statistics of packets received and sent on a specified interface or on all interfaces. You can also use the `ifstat` command to clear the statistics.

**About this task**

- The `ifstat` command displays the cumulative network interface statistics that are gathered either from the time of the last reboot or from the last time you cleared them.
- If you use the `ifstat` command on a storage system that is part of an HA pair, the resulting information pertains only to the storage system on which the command was run. The information does not include statistics for the partner node in an HA pair.
- In an HA pair in takeover mode, the `ifstat` command displays the combined statistics of the packets processed by the network interface on the local node and those on the partner node. Because the statistics displayed by the `ifstat` command are cumulative, a giveback does not cause the statistics to zero out.

**Step**

1. Depending on the statistics that you want to view, perform the following step:

| If you want to... | Enter the following command... |
| --- | --- |
| View the network interface statistics of all interfaces | `ifstat -a` |
| View the network interface statistics of a specific interface | `ifstat interface_name`<br><br>`interface_name` is the name of the network interface. |
| Clear the network interface statistics of a network interface | `ifstat -z interface_name` |

The output of the `ifstat` command depends on the type of interface. For example, Ethernet or Gigabit Ethernet interfaces generate different types of statistics.

---

**Example of showing the network interface statistics before and after clearing them**

To view the statistics of the network interface e0a, enter the following command:

`ifstat e0a`

An output similar to the following is displayed.

```
system1> ifstat e0a
```

```
-- interface  e0a  (8 days, 20 hours, 10 minutes, 27 seconds) --

RECEIVE
 Frames/second:        13 | Bytes/second:        800 | Errors/
minute:        0
 Discards/minute:       0 | Total frames:       897k | Total bytes:
62415k
 Total errors:         0 | Total discards:        0 | Multi/
broadcast:   734k
 No buffers:           0 | Non-primary u/c:       0 | Tag
drop:          0
 Vlan tag drop:        0 | Vlan untag drop:       0 | CRC
errors:        0
 Runt frames:          0 | Fragment:              0 | Long
frames:        0
 Jabber:               0 | Alignment errors:      0 | Bus
overruns:        0
 Queue overflows:      0 | Xon:                   0 |
Xoff:              0
 Jumbo:                0 | Reset:                 0 |
Reset1:            0
 Reset2:               0
TRANSMIT
 Frames/second:         2 | Bytes/second:        110 | Errors/
minute:        0
 Discards/minute:       0 | Total frames:       153k | Total bytes:
24129k
 Total errors:         0 | Total discards:        0 | Multi/
broadcast:   9478
 Queue overflows:      0 | No buffers:            0 | Max
collisions:        0
 Single collision:     0 | Multi collisions:      0 | Late
collisions:      0
 Timeout:              0 | Xon:                   0 |
Xoff:              0
 Jumbo:                0
LINK_INFO
 Current state:       up | Up to downs:           0 |
Auto:              on
 Speed:            1000m | Duplex:             full |
Flowcontrol:     none
```

The following command clears and reinitializes the statistics for the network interface e0a:
**ifstat -z e0a**

The following sample output shows the network interface statistics for the network interface
e0a immediately after the statistics are cleared.

```
system1> ifstat e0a

-- interface  e0a  (0 hours, 0 minutes, 8 seconds) --

RECEIVE
 Frames/second:         1 | Bytes/second:         32 | Errors/
minute:        0
 Discards/minute:       0 | Total frames:          7 | Total
bytes:        448
```

```
 Total errors:          0 | Total discards:       0 | Multi/
broadcast:      0
 No buffers:            0 | Non-primary u/c:      0 | Tag
drop:           0
 Vlan tag drop:         0 | Vlan untag drop:      0 | CRC
errors:         0
 Runt frames:           0 | Fragment:             0 | Long
frames:         0
 Jabber:                0 | Alignment errors:     0 | Bus
overruns:       0
 Queue overflows:       0 | Xon:                  0 |
Xoff:              0
 Jumbo:                 0 | Reset:                0 |
Reset1:            0
 Reset2:                0
TRANSMIT
 Frames/second:         1 | Bytes/second:        17 | Errors/
minute:         0
 Discards/minute:       0 | Total frames:         4 | Total
bytes:        361
 Total errors:          0 | Total discards:       0 | Multi/
broadcast:      0
 Queue overflows:       0 | No buffers:           0 | Max
collisions:        0
 Single collision:      0 | Multi collisions:     0 | Late
collisions:        0
 Timeout:               0 | Xon:                  0 |
Xoff:              0
 Jumbo:                 0
LINK_INFO
 Current state:        up | Up to downs:          0 |
Auto:             on
 Speed:             1000m | Duplex:            full |
Flowcontrol:     none
```

**Related references**

# Viewing statistics of dropped data packets

You can view the count of dropped data packets for each data protocol by using the netstat -s command. By default, the statistics for TCP, IP, UDP, ICMP, and IGMP are displayed.

**About this task**

The scenarios that might increase the count of filtered data packets are:

- DNS advertising the management IP address to data clients
- Wrong static route configuration
- Data traffic blocked on the management interface

The number of filtered data packets dropped for IP and UDP protocols are included in the weekly autosupport messages, as a part of the netstat -s command output.

**Step**

1. To view the total number of dropped packets for each protocol, enter the following command:

   **netstat -s**

---

**Example**

The example shows a part of the netstat -s command output.

```
system1> netstat -s

ip:
        4066350 total packets received
        0 bad header checksums
        0 with size smaller than minimum
        0 with size larger than maximum
        0 with data size< data length
        0 with header length< data size
        0 with data length< header length
        0 with bad options
        0 with incorrect version number
        0 packets with spoofed source address
        0 packets arrived on wrong port
        0 fragments received
        the last 2 src addr that send fragments:
        0.0.0.0 0.0.0.0
        0 fragments dropped (dup or out of space)
        0 malformed fragments dropped
        0 overlapping fragments discarded
        0 fragments dropped after timeout
        the last 2 src addr that have fragment time out:
        0.0.0.0 0.0.0.0
        0 packets dropped, too many fragments
        the last 2 src addr that sent too many fragments:
        0.0.0.0 0.0.0.0
        0 packets dropped, reassembly queue overflow
        0 packets reassembled ok
        3958788 packets for this host
        0 packets for unknown/unsupported protocol
        0 packets forwarded
        4986 packets not forwardable
        0 redirects sent
        3057658 packets sent from this host
        0 packets sent with fabricated ip header
        0 output packets dropped due to no bufs, etc.
        0 output packets discarded due to no route
        0 output datagrams fragmented
        0 fragments created
        0 datagrams that can't be fragmented
        498 filtered packets dropped
        0 packets dropped due to vfiler mismatch
```

```
        0 packets forwarded by source interface
```

## Viewing network interface information with FilerView

You can view network interface statistics, such as MTU size, incoming and outgoing packets on each interface, by using the Network Report in FilerView. You can also view routing information including the routing tables by using the Network Report.

### About this task

The Network Report in FilerView provides the same information that you get by running the `netstat -l`, `routed status`, and `netstat -rn` commands.

### Step

1. From the list on the left pane, click **Network > Report**.

   The Network Report displays the interface statistics and routing tables.

# How routing in Data ONTAP works

You can have Data ONTAP route its own outbound packets to network interfaces. Although your storage system can have multiple network interfaces, it does not function as a router. However, it can route its outbound packets.

Data ONTAP uses two routing mechanisms:

**Fast path**      Data ONTAP uses this mechanism to route NFS packets over UDP and to route all TCP traffic.

**Routing table**      To route IP traffic that does not use fast path, Data ONTAP uses the information available in the local routing table. The routing table contains the routes that have been established and are currently in use, as well as the default route specification.

## What fast path is

Fast path is an alternative routing mechanism to the routing table, in which the responses to incoming network traffic are sent back by using the same interface as the incoming traffic. By avoiding the routing table lookup, fast path provides a shorter route for data access. Fast path is used in all TCP and NFS/UDP protocols.

By default, fast path is enabled on the storage system; however, you can disable it.

- Fast path is used for all TCP and NFS/ UDP protocols.

  **Note:** Fast path is used only in the NFS/UDP protocols. However, fast path is not used in other UDP-based NFS services such as portmapper, mountd, and nlm.

- In a TCP connection, fast path is disabled after three consecutive retransmissions of the same data packet.
- Fast path should be disabled in asymmetric networks with lossy data paths.

Using fast path has the following advantages:

- Load balancing between multiple network interfaces on the same subnet.
  Load balancing is achieved by sending responses on the same interface of your storage system that receives the incoming requests.
  If fast path is enabled on an interface group and a physical interface in that group receives an incoming request, it is not necessary that the same interface sends a response to the request. Instead, any other physical interface in an interface group can send the response.
- Improved storage system performance by skipping routing table lookups.

### How fast path works with NFS-over-UDP

NFS-over-UDP traffic uses fast path only when sending a reply to a request. The reply packet is sent out on the same network interface that received the request packet.

For example, a storage system named toaster uses the toaster-e1 interface to send reply packets in response to NFS-over-UDP requests received on the toaster-e1 interface.

Fastpath is used only in the NFS/UDP protocols. However, fast path is not used in other UDP-based NFS services such as portmapper, mountd, and nlm.

### How fast path works with TCP

Data ONTAP can use fast path on every TCP packet transmitted except the first SYN packet (if Data ONTAP initiates a connection). The network interface that is used to transmit a packet is the same interface that received the last packet.

### Fast path not compatible with asymmetric routing

If fast path is enabled on your storage system in an asymmetric network, the destination MAC address of the response packet will be that of the router that forwarded the incoming packet. However, in asymmetric networks, the router that forwards packets to your storage system is not the router that forwards packets sent by the storage system. In such scenarios, you must disable fast path.

**Related concepts**

*Data ONTAP support for TSO* on page 27

**Related tasks**

*Enabling or disabling fast path* on page 61

# How to manage the routing table

You can manage the routing table automatically by using the routed daemon, or manually by using the route command.

## What the routed daemon does

The routed daemon performs several functions automatically and can be configured to perform several additional functions. The routed daemon is enabled by default.

The routed daemon performs the following functions by default:

- Deletes redirected routes after a specified period
- Performs router discovery with ICMP Router Discovery Protocol (IRDP)
  Router discovery is useful only if there is no static default route.
- Listens for Routing Information Protocol (RIP) packets

- Migrates routes to alternate interfaces when multiple interfaces are available on the same subnet
  For example, the network interface e0a is on 10.1.1.2/24 and e0b is on 10.1.1.3/24. The default
  route to 10.1.1.1 is through e0a. The subnet route 10.1.1.0/24 is also through e0a. Now if e0a is
  brought to the down status, the routed daemon moves the default route to e0b and the subnet
  route to e0b. Therefore, this function ensures continued connectivity to the local subnet and
  remote networks.

The routed daemon can also be configured to perform the following functions:

- Control RIP and IRDP behavior
- Generate RIP response messages that update a host route on your storage system
- Recognize distant gateways identified in the /etc/gateways file
- Authenticate RIPv2 advertisements with Message Digest 5 (MD5) algorithm for improved
  network security

For more information about the routed daemon, see the na_routed(1) man page.

## When the routed daemon should be turned off

In some circumstances, you should turn off the routed daemon. For example, you should turn it off
if you have multiple interfaces on the same subnet and you want to direct network traffic to specific
interfaces.

If you want to direct traffic to specific interfaces, you must turn off the routed daemon, because the
daemon sees all interfaces on a subnet as equivalent.

You can safely turn off the routed daemon if the following conditions are true:

- You do not use RIP or router discovery.
- You have a single router per subnet or a network in which redirects are not sent.
- You can manage your routing table directly.

  **Note:** Unless you have specific routing needs and understand network routing configuration, you
  are advised to always keep the routed daemon on. Turning off the routed daemon might cause
  unexpected routing behavior.

### Related tasks

Enabling or disabling the routed daemon from the command-line interface on page 62

## Routing tables in a vFiler unit environment

If you enable the MultiStore license and have more than one IPspace, Data ONTAP disables the
routed daemon. Therefore, routing tables in such vFiler unit environments must be managed
manually with the route command.

The routed daemon runs only in the default IPspace. Data ONTAP does not support running
multiple instances of the routed daemon in multiple IPspaces. Therefore, if you have multiple
IPspaces, the routed daemon is automatically disabled.

All vFiler units in an IPspace share a routing table. Therefore, any commands that display or manipulate the routing table apply to all vFiler units in that IPspace.

## Circumstances that might alter the routing table

Certain events can cause the routing table to be modified. You should check the routing table after these events occur to be sure that it is still configured as required.

The routing table might be modified in the following circumstances:

- A new interface is configured with the `ifconfig` command and there are no existing entries for the new network number in the routing table.
- You use the `route add` command to add an entry to the routing table.
- Your storage system receives an ICMP redirect packet, which notifies the storage system of a better first-hop router for a particular destination.

    **Note:** Your storage system ignores ICMP redirect packets if the `ip.icmp_ignore_redirect.enable` option is turned on.

- Your storage system is rebooted after the default route in the `/etc/rc` file is modified.

# Specifying the default route

A default route is the network route used by the routing table when no other known route exists for a given packet's destination address. You can set the default route in Data ONTAP either during initial setup or by modifying the `/etc/rc` file.

### Before you begin

The IP addresses configured on e0M must not be on the same subnet as the default route gateway.

   **Note:** An EMS message is displayed if the default route uses the dedicated management interface.

### Steps

1. Open the `/etc/rc` file in the root volume by using a text editor.

2. Add the following command to the `/etc/rc` file:

   **`route add default route_IP`**

   `route_IP` is the IP address of the default route

   ### Example
   The following example shows the default route being set in the `/etc/rc` file by using the `route add` command:

```
hostname sys1
ifconfig e0 192.0.2.21 netmask 255.255.255.0 mediatype 100tx-fd
```

```
route add default 192.0.2.1 1
routed on
```

**Related concepts**

*Guidelines for configuring e0M* on page 25

# How to enable or disable routing mechanisms

Both the fast path mechanism and the `routed` daemon are enabled by default in Data ONTAP. You can enable or disable these routing mechanisms using the command-line interface or FilerView.

If you disable both fast path and the `routed` daemon, you must configure routing manually.

## Enabling or disabling fast path

IP fast path is a mechanism that uses the network interface of an inbound request to send the response bypassing a routing table lookup. Fast path is enabled by default for all TCP connections and only for the NFS/UDP protocol.

### About this task

- In TCP connections using an asymmetric network path, if a data packet is not acknowledged even after three consecutive retransmissions, fast path is disabled automatically and all the subsequent packets follow the route specified in the routing table.
- During a takeover, the routing behavior on the failed node is determined by the fast path settings on the surviving node.
- Disabling fast path results in the TCP segmentation offload (TSO) feature being disabled.

### Step

1. To enable or disable fastpath, enter the following command:

   **`options ip.fastpath.enable {on|off}`**

   `on`—Enables fast path

   `off`—Disables fast path

   > **Note:** You can use the `netstat -x` command to check if the fast path mechanism is enabled.

**Related concepts**

*What fast path is* on page 57

## Enabling or disabling the routed daemon from the command-line interface

You can manage the routing table automatically by using the routed daemon. You can turn on or turn off the routed daemon by using the `routed` command.

### About this task

You must add the `routed` command to the `/etc/rc` file for the routed daemon behavior to persist across storage system reboots.

### Step

1. To enable or disable the routed daemon, enter the following command:

   **`routed {on|off}`**

   `on`—Turns on the routed daemon

   `off`—Turns off the routed daemon

### Related concepts

*What the routed daemon does* on page 58
*When the routed daemon should be turned off* on page 59

### Related references

*The routed daemon* on page 188

## Enabling or disabling the routed daemon with FilerView

You can use FilerView to turn on or turn off the `routed` daemon.

### Steps

1. From the list on the left pane, click **Network > Configure**.
2. Select **Yes** (for on) or **No** (for off) from the Routed Enabled drop-down list.
3. Click **Apply**.

### Related concepts

*What the routed daemon does* on page 58
*When the routed daemon should be turned off* on page 59

### Related references

*The routed daemon* on page 188

# How to view the routing table and default route information

You can view the routing table of the storage system and default route information relating to your route's destinations, their gateways, how much each route is used, and the interface used by each route. Flags showing route status information are also displayed.

## Viewing the routing table from the command-line interface

You can view information such as default route and the routes for specific destination addresses.

**Step**

1. To view the Data ONTAP routing table, enter one of the following commands:

   - **netstat -rn**
   - **route -s**

---

**Example for interpreting the routing table**

The output of the netstat -rn command is as follows:

```
netstat -rn
Routing tables

Internet:
Destination     Gateway           Flags      Refs      Use   Interface
default         192.0.2.1         UGS           7   116009   e0a
192.0.2/24      link#3            UC            0        0   e0a
192.0.2.1       0:15:c6:bf:c6:44  UHL           1        0   e0a
192.0.2.68      0:a0:98:7:66:2    UHL           0     1079   lo
192.0.2.216     0:a0:98:7:3e:5    UHL           0        0   e0a
192.0.2.217     0:a0:98:7:3e:5    UHL           0        0   e0a
192.0.2.255     ff:ff:ff:ff:ff:ff UHL           0     9470   e0a
127.0.0.1       127.0.0.1         UH            0       11   lo
```

In this example, the destination can be a host 192.0.2.1, a network 192.0.2/24, or the default route. If the destination is a subnet on a network, the network number is followed by a forward slash (/) and a number that describes the network mask for that network.

---

**Routing table flags**

The following table describes the Flags column in the netstat -rn output.

| Flag | Description |
|------|-------------|
| U | Up—Route is valid |
| G | Gateway—Route is to a gateway router rather than to a directly connected network or host |

| Flag | Description |
|------|-------------|
| H | Host name—Route is to a host rather than to a network, where the destination address is a complete address |
| R | Reject—Set by ARP when an entry expires (for example, the IP address could not be resolved into a MAC address) |
| D | Dynamic—Route added by a route redirect or RIP (if routed is enabled) |
| M | Modified—Route modified by a route redirect |
| C | Cloning—A new route is cloned from this entry when it is used |
| L | Link—Link-level information, such as the Ethernet MAC address, is present |
| S | Static—Route added with the route command |

## Viewing the default route information from the command-line interface

You can view default route information such as whether the routed daemon is turned on or off, default route information, and routing protocols. You can view the default route information by using the routed status command.

**Step**

1. Enter the following command:

   **routed status**

   > **Note:** You can also view the default route by using the netstat -rn or route -s commands.

   **Example**
   The output of the routed status command is as follows:

   ```
   routed status
   RIP snooping is on
   Gateway               Metric  State    Time Last Heard
   example-gateway.com      1     ALIVE    Wed Mar 18 13:58:56 IST 2009
   0 free gateway entries, 1 used
   ```

   In the routed status command output, metric is the route property that is used to determine the preferred route. The route with the lowest metric is the preferred route. You should always use a metric greater than 0 when adding default routes.

### Viewing the routing table and routing information by using FilerView

You can view the routing table, routing information, and routing protocols by using FilerView. You can view information such as default route and the routes for specific destination addresses.

**Step**

1. From the list on the left pane, click **Network > Report**.

   The Routing section of the Network Report shows the default route and protocols in effect, as well as routing tables.

## Modifying the routing table

You might want to add or delete routes in your routing table depending on the changes in your network. You can use the route command to modify the routing table. You cannot modify the routing table using FilerView.

**Step**

1. Depending on whether you want to add or delete a route from the routing table, perform the following step:

| If you want to... | Enter the following command... |
|---|---|
| Add a route | `route add destination [gateway metric]`<br><br>`destination` is the IP address or host name of the destination for which the route is being added or deleted.<br><br>`gateway` is the gateway for the specified `destination`.<br><br>`metric` indicates the number of hops to the `destination`. The value of `metric` should be greater than zero when the route to the destination is through the `gateway`. The value of `metric` is zero when the `destination` is on a directly-attached network. |
| Delete a route | `route delete destination [gateway metric]`<br><br>**Attention:** You must not delete a cloned route (denoted by the C flag) from the routing table; if you do, the network connectivity to that subnet is lost. If you have deleted a cloned route, you must add the route again to the routing table in either of the following ways:<br><br>• Bring the interface that connects to the particular subnet first to the down state and then to the up state.<br>  You can change the state of the interface by using the ifconfig command.<br>• Delete and reconfigure the IP address on the interface that connects to the particular subnet. |

For more information about the `route` command and options, see the na_route(1) man page.

---

**Example**

To add a destination with the IP address 192.0.2.25 to the routing table, enter the following command:

**`route add 192.0.2.25 gateway.com 1`**

You can verify that the route to this destination is added to the routing table by using the `netstat -rn` or `route -sn` command, as shown in the following output:

```
system1> netstat -rn
Routing tables

Internet:
Destination     Gateway             Flags     Refs      Use  Interface
default         192.0.2.1           UGS          4   184855  e0a
127.0.0.1       127.0.0.1           UH           0        0  lo
192.0.2/24      link#11             UC           2     1238  e0a
192.0.2.1       0:d0:d3:0:30:0      UHL          0       40  e0a
192.0.2.23      0:1:30:b8:30:c0     UHL          1        0  e0a
192.0.2.25      192.0.2.1           UHL          0     1285  lo
```

In this example, the subnet route, 192.0.2, is a cloned route.

---

**Related tasks**

[Changing the status of an interface](#) on page 45

**Related references**

[Routing table flags](#) on page 63

# How to maintain host-name information

Data ONTAP relies on correct resolution of host names to provide basic connectivity for storage systems on the network. If you are unable to access the storage system data or establish sessions, there might be problems with host-name resolution on your storage system or on a name server.

Host-name information can be maintained in one or all of the following ways in Data ONTAP:

- In the /etc/hosts file on your storage system's default volume
- On a Domain Name System (DNS) server
- On a Network Information Service (NIS) server

If you use more than one of the resources for host-name resolution, the order in which they are used is determined by the /etc/nsswitch.conf file.

## How the /etc/hosts file works

Data ONTAP uses the /etc/hosts file to resolve host names to IP addresses. You need to keep the /etc/hosts file up-to-date. Changes to the /etc/hosts file take effect immediately.

When Data ONTAP is first installed, the /etc/hosts file is automatically created with default entries for the following interfaces:

- localhost
- All interfaces on your storage system

The /etc/hosts file resolves the host names for the storage system on which it is configured. This file cannot be used by other systems for name resolution.

For more information about file formats, see the na_hosts(5) man page.

You can add IP address and host name entries in the /etc/hosts file in the following two ways:

- Locally—You can add entries by using the command-line interface or FilerView.
- Remotely—If the file has many entries and you have access to an NIS makefile master, you can use the NIS makefile master to create the /etc/hosts file. This method prevents errors that might be caused by editing the file manually.

## Adding a host name in the /etc/hosts file

You can add the host name and aliases of the storage system in the `/etc/hosts` file. You can use the `setup` command to rewrite the `/etc/hosts` file.

**Step**

1. From a workstation that has access to your storage system, edit the `/etc/hosts` file. Add the following line to the `/etc/hosts` file:

   ***IP_address host_name aliases***

   *IP_address* is the IP address of the host.

   *host_name* is the name of the host.

   *aliases* are the alias names for the host.

   **Example**

   To add a host name, myhost, with an IP address 192.0.2.16, add the following line in the `/etc/hosts` file:

   **192.0.2.16 myhost newhost myhost-e0a**

   newhost and myhost-e0a are the alias names for myhost.

   The following is a sample `/etc/hosts` file:

   ```
   #Auto-generated by setup Tue Apr 21 17:41:40 IST 2009
   127.0.0.1 localhost
   192.0.2.16          myhost myhost-e0a
   # 0.0.0.0        myhost-e0b
   # 0.0.0.0        myhost-e0c
   # 0.0.0.0        myhost-e0d
   ```

## Hard limits for the /etc/hosts file

You need to be aware of the hard limits on the line size and number of aliases when you edit the `/etc/hosts` file.

The hard limits are as follows:

- Maximum line size is 1022 characters.
  The line size limit includes the end of line character. You can enter up to 1021 characters per line.
- Maximum number of aliases is 34.

  **Note:** There is no limit on file size.

## Editing the /etc/hosts file with FilerView

You can add entries to the local /etc/hosts file if the number of entries is small.

**Steps**

1. In FilerView, click **Network** in the list on the left pane.

2. In the list under Network, click **Manage Hosts File**.

3. Click in the hosts window, then click **Insert**.

4. Complete the fields in the **Create a New /etc/hosts Line** window for each host you want to add and click **OK**.

5. Click **Apply** in the **Manage Hosts File** window.

## Changing the host name of a storage system

You can change the host name of a storage system by editing the /etc/hosts file, and then using the hostname command.

**Steps**

1. Edit the /etc/hosts file to include the new host name of the storage system.

2. At the command-line interface of the storage system, enter the following command to specify a new name for the host:

   **hostname *new_name***

   *new_name* is the new host name of the storage system. The host name must match the entry made in the /etc/hosts file in Step 1.

   > **Attention:** If you skip this step, any manual or scheduled SnapMirror operations might fail. Use the hostname command to specify the correct name before any SnapMirror operations are initiated.

# How to configure DNS to maintain host information

You can maintain host information centrally using DNS. With DNS, you do not have to update the /etc/hosts file every time you add a new host to the network. You can configure your storage system to use one or more DNS servers either during the setup procedure or later.

If you have several storage systems on your network, maintaining host information centrally saves you from updating the /etc/hosts file on each storage system every time you add or delete a host.

If you configure DNS during the setup procedure, your storage system's DNS domain name and name server addresses are configured in one of the following ways:

- Automatically if you use Dynamic Host Configuration Protocol (DHCP) to configure onboard interfaces. Automatic configuration is possible only if all the DHCP-configured DNS server addresses are IPv4 addresses.
- Manually if you do not use DHCP—you must enter the values when prompted. A maximum of three name server IP addresses can be specified for a DNS server.

If you configure DNS later, you must take the following actions:

- Specify DNS name servers.
- Specify the DNS domain name of your storage system.
- Enable DNS on your storage system.

You can enable DNS and set DNS configuration values in either of the following ways:

- Using FilerView
- Using the command-line interface

If you want to use primarily DNS for host-name resolution, you should specify it ahead of other methods in the hosts section of the `/etc/nsswitch.conf` file.

Correct host-name resolution depends on correctly configuring the DNS server. If you experience problems with host-name resolution or data availability, check the DNS server in addition to local networking.

For more information about storage system DNS resolution of host names, see the na_dns(1) and na_dns(8) man pages.

**Related concepts**

*How the /etc/hosts file works* on page 67

## Configuring DNS from the command-line interface

You can configure your storage system to use one or more DNS servers for host-name resolution. You can configure DNS by first creating or editing the `/etc/resolv.conf` file, then specifying the DNS domain name, and finally enabling DNS through the command-line interface.

**Steps**

1. Depending on whether you want to create or edit the `/etc/resolv.conf` file, perform the following step:

| If you are... | Then... |
| --- | --- |
| Creating the `/etc/resolv.conf` file | By using a text editor, create the `/etc/resolv.conf` file in the root volume. The file can consist of up to three lines, each specifying a name server host in the following format. <br><br> **nameserver *ip_address*** <br><br> *ip_address* is the IP address of the DNS name server. |

| If you are... | Then... |
|---|---|
| Editing the `/etc/resolv.conf` file | From a workstation that has access to the root volume of your storage system, edit the `/etc/resolv.conf` file using a text editor. |

**2.** Enter the following command to specify the DNS domain name:

**`options dns.domainname domain`**

`domain` is the new domain name, which follows the host name of your storage system in the fully qualified domain name.

**3.** Enter the following command to enable DNS:

**`options dns.enable {on|off}`**

`on`—Enables DNS

`off`—Disables DNS

### Hard limits for the /etc/resolv.conf file

You need to be aware of the hard limits for name servers, domain name, and search domains when you create or edit the `/etc/resolv.conf` file.

The hard limits for the `/etc/resolv.conf` file are as follows:

- Maximum line size is 256.
- Maximum number of name servers is 3.
- Maximum domain name length is 256 characters.
- Maximum search domains limit is 6.

    **Note:** You should use only tab or space to separate host names in the search domain list.

- Total number of characters for all search domains is 256.

  **Note:** There is no limit on file size.

## How DNS resolves host names

DNS uses certain records for resolving a domain name to an IP address. To determine a host name based on the IP address, DNS uses the reverse lookup.

For resolving IPv4 addresses, DNS uses the A record. The A record can store a 32-bit address and can resolve IPv4 addresses.

IPv4 reverse DNS lookups use the `in-addr.arpa` domain. An IPv4 address is represented in the in-addr.arpa domain by a sequence of bytes, represented as decimal numbers, in reverse order. The numbers are separated by dots and end with the suffix `.in-addr.arpa`.

The following table shows a sample IPv4 address and its reverse DNS lookup:

| IP address | Reverse lookup domain name |
|---|---|
| 192.0.2.10 | 10.2.0.192.in-addr.arpa |

## DNS name caching

DNS name caching speeds up the process whereby the DNS name resolver converts host names into IP addresses. The DNS name cache stores DNS requests so that they can be easily and quickly found when needed. DNS name caching is enabled by default.

Name caching improves DNS performance during a name server failover and reduces the time needed for an HA pair takeover and giveback.

You can disable DNS name caching by using the `dns.cache.enable` option, but doing so might have an adverse performance impact. The `dns flush` command removes all entries from the DNS name cache. However, the command has no effect if DNS name caching is not enabled.

For more information about the `dns flush` command and the `dns.cache.enable` option, see the na_ dns(1) man page.

## DNS information you can view

You can view information about whether DNS and DNS name caching are enabled, configured name servers, state of these name servers (whether up or down), configured DNS domain name, DNS name cache statistics, and performance statistics for each name server.

The `dns info` command displays the status of the DNS resolver. If DNS is enabled, the command displays the following information:

- Whether DNS is enabled
- Whether DNS name caching is enabled
- Caching statistics

  - Cache hits: Number of DNS requests that were found in the cache
  - Cache misses: Number of DNS requests that were not found in the cache and that required a DNS query to the name server
  - Cache entries: Number of entries currently in the DNS name cache
  - Expired cache entries
  - Number of cache replacements
- Details about each name server that was polled by your storage system:

  - IP address of the DNS server
  - State of the name server, displayed as "UP," "DOWN," or "NO INFO"
  - Date of the last DNS request to that name server
  - Average time in milliseconds for a DNS query
  - Number of DNS queries made
  - Number of DNS queries that resulted in errors
- Default DNS domain name of the storage system

- Search domains of the storage system

The search domains are domain suffixes that are used to convert unqualified domain names into fully qualified domain names (FQDN). The search domains are read from the /etc/resolv.conf file.

For more information about the dns info command and the resulting display, see the na_ dns(1) man page.

# How to use dynamic DNS to update host information

You can use dynamic DNS updates to prevent errors and save time when sending new or changed DNS information to the primary master DNS server for your storage system's zone. Dynamic DNS allows your storage system to automatically send information to the DNS servers as soon as the information changes on the system.

Without dynamic DNS updates, you must manually add DNS information (DNS name and IP address) to the identified DNS servers when a new system is brought online or when existing DNS information changes. This process is slow and error-prone. During disaster recovery, manual configuration can result in a long downtime.

For example, if you want to change the IP address on interface e0 of *storagesystem1*, you can simply configure e0 with the new IP address. The storage system *storagesystem1* automatically sends its updated information to the primary master DNS server.

> **Note:** Data ONTAP supports a maximum of 64 Dynamic Domain Name Server (DDNS) aliases.

## How dynamic DNS updates work in Data ONTAP

If dynamic DNS updates are enabled on your storage system, Data ONTAP periodically sends updates to the primary master DNS server for its zone. Updates are also sent if any DNS information changes on your system.

Your storage system finds the primary master DNS server for its zone by querying the DNS servers configured in your storage system's /etc/resolv.conf file. The primary master DNS server might be different from the ones configured in your storage system's /etc/resolv.conf file.

By default, periodic updates are sent every 12 hours. A time-to-live (TTL) value is assigned to every DNS update sent from your storage system. The TTL value defines the time for which a DNS entry is valid on the DNS server. By default, the TTL value is set to 24 hours, and you can change it.

When your storage system sends an update to the DNS server, it waits up to five minutes to receive an acknowledgement of the update from the server. If it does not receive an acknowledgement, the storage system sends the update again. This time, the storage system doubles the waiting interval (to 10 minutes), before sending the update. The storage system continues to double the waiting interval with each retry until a waiting interval of 160 minutes or TTL/2, whichever is less, is reached.

## Support for dynamic DNS updates in Data ONTAP

When you use dynamic DNS updates in Data ONTAP, you must be aware of certain conditions, such as the types of systems and network interfaces that support dynamic DNS, secure updates, and behavior of vFiler units with dynamic DNS.

The following conditions apply to dynamic DNS updates:

- By default, dynamic DNS updates are disabled in Data ONTAP.
- Dynamic DNS updates are supported on UNIX and Windows systems.
- On Windows DNS servers, secure dynamic DNS updates can be used to prevent malicious updates on the DNS servers. Kerberos is used to authenticate updates.
  Even if secure dynamic DNS updates are enabled, your storage system initially tries sending updates in clear text. If the DNS server is configured to accept only secure updates, the updates sent in clear text are rejected. Upon rejection, the storage system sends secure DNS updates.
- For secure dynamic DNS updates, your storage system must have CIFS running and must be using Windows Domain authentication.
- Dynamic DNS updates can be sent for the following:

  - Physical interfaces
  - Interface group and VLAN interfaces
  - vFiler units
- You cannot set TTL values for individual vFiler units. All vFiler units inherit the TTL value that is set for vfiler0, which is the default vFiler unit and is the same as the physical storage system.
- DHCP addresses cannot be dynamically updated.
- In a takeover situation, the hosting storage system is responsible for sending DNS updates for IP addresses for which it is responding.

## Enabling or disabling dynamic DNS updates

Dynamic DNS enables your storage system to automatically send information to the DNS servers as soon as the information changes on the system. By default, dynamic DNS is disabled on the storage system. You can enable dynamic DNS on your storage system by using the `options dns.update.enable` command.

**Step**

1. Enter the following command:

   **`options dns.update.enable {on|off|secure}`**

   `on`—Enables dynamic DNS updates

   `off`—Disables dynamic DNS updates

   `secure`—Enables secure dynamic DNS updates

   **Note:** Secure dynamic DNS updates are supported for Windows DNS servers only.

## Disabling the transmission of DNS updates for an IP address

You can disable the transmission of dynamic DNS updates for an IP address by using the `ifconfig` command.

### About this task

You should not disable dynamic DNS updates for an interface that is part of an interface group.

### Step

1. Enter the following command:

   **ifconfig *interface_name* no_ddns *IP_address***

   *interface_name* is the name of the interface.

   *IP_address* is the IP address of the interface.

   ### Example
   Use the following command to ensure that dynamic DNS updates are not sent from the interface e0a:

   **ifconfig e0a no_ddns 192.0.2.30**

   The following output shows the output of the `ifconfig` command after the dynamic DNS is disabled for the interface:

   ```
   ifconfig e0a
   e0a: flags=0x2d48867<UP,BROADCAST,RUNNING,MULTICAST,TCPCKSUM> mtu 1500
   inet 192.0.2.30 netmask 0xff000000 broadcast 10.255.255.255 noddns
   ether 00:a0:98:07:66:02 (auto-1000t-fd-up) flowcontrol full
   ```

   The `ifconfig` command output shows the "noddns" keyword that indicates that dynamic DNS updates are disabled for this IP address.

## Changing the time-to-live setting for DNS entries

You can change the time-to-live setting for DNS entries by using the `options dns.update.ttl` command.

### Step

1. Enter the following command:

   **options dns.update.ttl *time***

   *time* can be set in seconds (**s**), minutes (**m**), or hours (**h**), with a minimum value of 600 seconds and a maximum value of 24 hours.

   ### Example
   To set the TTL to two hours, enter the following command:

```
options dns.update.ttl 2h
```

**Related concepts**

*How dynamic DNS updates work in Data ONTAP* on page 73

# How to use NIS to maintain host information

NIS enables you to centrally maintain host information. In addition, NIS enables you to maintain user information.

NIS provides the following methods for resolving the storage system's host name:

- Using the `/etc/hosts` file on the NIS server
  You can download the `/etc/hosts` file on the NIS server to your storage system's default volume for local host-name lookup.
- Using a hosts map that is maintained as a database on the NIS server
  The storage system uses the hosts map to query during a host lookup request across the network.

## How using NIS slaves can improve performance

Host-name resolution by using a hosts map can have a performance impact because each query for the hosts map is sent across the network to the NIS server. You can improve the performance of your storage system by downloading the maps and listening for updates from the NIS master server.

The NIS slave improves performance by establishing contact with an NIS master server and performing the following two tasks:

- Downloading the maps from the NIS master server
  You can download the maps from the NIS master server to the NIS slave by running the `yppush` command from the NIS server. You can also download the maps by disabling and then enabling the NIS slave from your storage system. After the maps are downloaded, they are stored in the `/etc/yp/nis_domain_name` directory. The NIS slave then services all the NIS requests from your storage system by using these maps. The NIS slave checks the NIS master every 45 minutes for any changes to the maps. If there are changes, they are downloaded.
- Listening for updates from the NIS master
  When the maps on the NIS master are changed, the NIS master administrator can optionally notify all slaves. Therefore, in addition to periodically checking for updates from the NIS master, the NIS slave also listens for updates from the master.

You cannot configure the NIS slave during the setup procedure. To configure the NIS slave after the setup procedure is complete, you need to enable NIS slave by setting `options nis.slave.enable` to `on`.

**Note:** The NIS slave does not respond to remote NIS client requests and therefore cannot be used by other NIS clients for name lookups.

**Related concepts**

**Related tasks**

# How an NIS master is selected

If you enable the NIS slave on your storage system, the NIS servers listed with the nis.servers option are contacted to determine the master NIS server.

The NIS master can be different from the servers that are listed with the nis.servers option. In such a case, the servers listed with the nis.servers option inform the slave about the master server.

The NIS slave on your storage system can contact the master only if any one of the following conditions is true:

- The NIS server has an entry in the hosts map for the master.
- The /etc/hosts file on your storage system is able to resolve the IP address of the master.

# Creating /etc/hosts from the NIS master

You can create a host file remotely and modify the NIS master to install the host file in the /etc directory. This method is useful if you have many entries in your host file.

**Steps**

1. On the NIS server, open the NIS Makefile with a text editor.

2. Locate the section for hosts.time.

3. Add the following lines at the end of the hosts.time section, replacing *dirname* with a directory name of your choice, and *toaster 1*, *toaster2*, and so on with names of the storage systems:

    ```
    @mntdir=/tmp/dirname_etc_mnt_$$$$;\ if [ ! -d $$mntdir ]; then rm -f $
    $mntdir; \ mkdir $$mntdir; fi;\ for s_system in toaster1 toaster2
    toaster3 ; do \ mount $$s_system:/etc $$mntdir;\ mv $$mntdir/hosts $
    $mntdir/hosts.bak;\ cp /etc/hosts $$mntdir/hosts;\ umount $$mntdir;\
    done;\ rmdir $$mntdir
    ```

4. Save the NIS Makefile.

   The /etc/hosts file on your storage system is updated whenever the NIS Makefile is run.

**Related concepts**

## Guidelines for using NIS slaves

When using an NIS slave, you should follow certain guidelines, such as the available space in the storage system, conditions for enabling DNS, and supported configurations.

The following guidelines apply when using the NIS slave:

* The root volume of your storage system must have sufficient space to download maps for the NIS slave. Typically, the space required in the root volume is same as the size of the maps on the NIS server.
* If the root volume does not have enough space to download maps, the following occurs:

  * An error message is displayed informing you that the space on the disk is not sufficient to download or update the maps from the NIS master.
  * If the maps cannot be downloaded, the NIS slave is disabled. Your storage system switches to using hosts map on the NIS server for name resolution.
  * If the maps cannot be updated, your storage system continues to use the old maps.
* If the NIS master server was started with the `-d` option or if the `hosts.byname` and `hosts.byaddr` maps are generated with the `-b` option, your storage system must have DNS enabled, DNS servers must be configured, and the hosts entry in the `/etc/nsswitch.conf` file must contain DNS as an option to use for host name lookup.

If you have your NIS server configured to perform host name lookups using DNS, or if you use DNS to resolve names that cannot be first resolved using the `hosts.by*` maps, using the NIS slave causes those lookups to fail. This is because when the NIS slave is used, all lookups are performed locally using the downloaded maps. However, if you configure DNS on your storage system, the lookups succeed.

You can use the NIS slave for the following:

* Interface groups and VLAN interfaces
* vFiler units
* HA pairs

  > **Note:** In an HA pair, you should ensure that the `nis.servers` options value is the same on both nodes and that the `/etc/hosts` file on both nodes can resolve the name of the NIS master server.

**Related concepts**

## NIS administrative commands

You can use the NIS administrative commands to view the NIS server information.

Data ONTAP supports the standard NIS administrative commands listed in the following table. For more information, see each command's man page.

| Command | Function |
|---------|----------|
| ypcat | Prints an entire NIS map. |
| ypgroup | Displays the NIS group cache entries. |
| ypmatch | Looks up specific entries in an NIS map. |
| ypwhich | Returns the name of the current NIS server. |

# How to configure NIS with Data ONTAP interfaces

You can configure your storage system to use one or more NIS servers either during the setup procedure or later using the Data ONTAP command-line interface or FilerView.

If you want to use NIS primarily for host-name resolution, specify it ahead of other methods in the hosts map in the /etc/nsswitch.conf file.

To configure NIS, you need to do all of the following:

- Specify the NIS server to which your storage system should bind.
- Specify the NIS domain name of your storage system.
- Enable NIS on your storage system.

Correct host-name resolution depends on correctly configuring the NIS server. If you experience problems with host-name resolution or data availability, check the NIS server in addition to local networking.

For more information about your NIS client, see the na_nis(1) and na_nis(8) man pages.

## Enabling NIS using the command-line interface

You can enable NIS on your storage system for host-name resolution. Whenever you enable NIS for group lookup services, you must also enable local caching of the NIS group files.

### About this task

If you use NIS for group lookup services, disabling NIS group caching can severely degrade performance. Failure to enable NIS lookups and NIS caching together could lead to timeouts as CIFS clients attempt authentication.

### Steps

**1.** To enable NIS lookups, enter the following command:

**options nis.enable on**

**2.** To enable NIS caching, enter the following command:

**options nis.group_update.enable on**

## Specifying the NIS domain name

You can specify the NIS domain name to which your storage system belongs.

**Step**

1. Enter the following command:

   **options nis.domainname *domain***

   *domain* is the NIS domain name to which your storage system belongs. For example, typical NIS domain names might be sales or marketing. The NIS domain name is usually not the same as the DNS domain name.

## Specifying NIS servers to bind to your storage system

You can specify an ordered list of NIS servers to which you want your storage system to bind. The list should begin with the closest NIS server (closest in network terms) and end with the farthest one.

**About this task**

Keep the following in mind before performing the binding procedure:

- Using the NIS broadcast feature can incur security risks.
- You can specify NIS servers by IP address or host name. If host names are used, ensure that each host name and its IP address are listed in the /etc/hosts file of your storage system. Otherwise, the binding with the host name fails.
- You can only specify IPv4 addresses or server names that resolve to IPv4 addresses by using the /etc/hosts file on your storage system.

**Step**

1. Enter the following command to specify the NIS servers and their order:

   **options nis.servers *ip_address, server_name,*[*]**

   The asterisk (*) specifies that broadcast is used to bind to NIS servers if the servers in the list are not responding. The '*' is the default value. If you do not specify the broadcast option, and none of the listed servers is responding, NIS services are disrupted until one of the preferred servers responds.

   > **Example**
   >
   > The following command lists two servers and uses the default broadcast option:
   >
   > **options nis.servers 192.0.2.1,nisserver-1,***
   >
   > Your storage system first tries to bind to 192.0.2.1. If the binding fails, the storage system tries to bind to nisserver-1. If this binding also fails, the storage system binds to any server that

responds to the broadcast. However, the storage system continues to poll the preferred servers. When one of the preferred server is found, the storage system binds to the preferred server.

**Related concepts**

*How an NIS master is selected* on page 77

## Enabling an NIS slave on your storage system

You can enable an NIS slave on your storage system to reduce traffic over your network.

**Step**

1. To enable or disable an NIS slave on your storage system, enter the following command:

   **options nis.slave.enable {on|off}**

**Related concepts**

*How using NIS slaves can improve performance* on page 76
*Guidelines for using NIS slaves* on page 78

# What NIS information you can view

You can view information about NIS master and slave servers, netgroup caches, and performance statistics.

The nis info command displays the following types of NIS information:

- NIS domain name
- Last time the local group cache was updated
- Information about each NIS server that was polled by your storage system:

  - IP address of the NIS server
  - Type of NIS server
  - State of the NIS server
  - Whether your storage system is bound to the NIS server
  - Time of polling
- Information about the NIS netgroup cache:

  - Status of the cache
  - Status of the "*.*" entry in the cache
  - Status of the "*.nisdomain" entry in the cache
- Whether an NIS slave is enabled
- NIS master server
- Last time the NIS map was checked by the NIS slave

- NIS performance statistics:

  - Number of YP lookup network retransmission
  - Total time spent in YP lookups
  - Number of network retransmission
  - Minimum time spent in a YP lookup
  - Maximum time spent in a YP lookup
  - Average time spent in a YP lookup
- Response statistics for the three most recent YP lookups

For more information about the `nis info` command and resulting display, see the na_nis(1) man page.

## Viewing NIS performance statistics

You can use the `nis info` command to view NIS performance statistics for your storage system.

**Step**

1. Enter the following command to view NIS information:

   **nis info**

   ---

   **Example**

   The following example shows the statistics provided by the `nis info` command.

   ```
   system1*> nis info
   NIS domain is lab.example.com
       NIS group cache has been disabled

        IP Address    Type State Bound   Last Polled
   Client
   calls Became Active
   ------------------------------------------------------------------
   ---------------------------
       192.0.2.12 PREF ALIVE YES     Mon Jan 23 23:11:14
   GMT 2008 0 Fri Jan 20 22:25:47 GMT 2008

                   NIS Performance Statistics:
                 Number of YP Lookups: 153
                 Total time spent in YP Lookups: 684 ms, 656 us
                 Number of network re-transmissions: 0
                 Minimum time spent in a YP Lookup: 0 ms, 1 us
                 Maximum time spent in a YP Lookup: 469 ms, 991 us
                 Average time spent in YP Lookups: 4 ms, 474 us

                   3 Most Recent Lookups:
               [0] Lookup time: 0 ms, 1 us Number of network re-
   transmissions: 0
               [1] Lookup time: 5 ms, 993 us Number of network re-
   transmissions: 0
               [2] Lookup time: 0 ms, 1 us Number of network re-
   ```

```
transmissions: 0
NIS netgroup (*.* and *.nisdomain) cache status:
uninitialized
         *.* eCode: 0
         *.nisdomain eCode: 0
   NIS Slave disabled
```

# Configuring DNS and NIS with FilerView

You can configure DNS and NIS for host-name resolution by using FilerView. You can also configure the host-name service configuration file (/etc/nsswitch.conf) with FilerView.

**Steps**

1. Click **Network** in the list on the left pane.

2. In the list under Network, click **Configure Host Name Resolution (DNS & NIS)**.

   The Host Name Resolution Policy Wizard is displayed.

3. Click **Next** and complete the steps in the **Host Name Resolution Policy Wizard** to set or modify the DNS and NIS configuration values.

   You can perform the following tasks by using the Host Name Resolution Policy Wizard:

   - Enable DNS and NIS
   - Enter a DNS domain name
   - Specify the dynamic DNS update interval
   - Enable dynamic DNS
   - Use DNS cache
   - Specify IP addresses of DNS servers (maximum of three IP addresses)
   - Specify the domain search list
   - Specify a NIS domain name
   - Specify NIS servers
   - Enable NIS domain search
   - Enable NIS slave
   - Enable local caching of NIS group files
   - Specify the schedule to update the local cache of NIS group files
   - Define the search order for hosts, password, shadow, group, and netgroup information

**Related concepts**

*How to configure DNS to maintain host information* on page 69

*DNS name caching* on page 72

*How to use dynamic DNS to update host information* on page 73

# How to change the host-name search order

If you use more than one method for host-name resolution, you must specify the order in which each name resolution service is used. This order is specified in the /etc/nsswitch.conf file in your storage system's root volume. You can change this order at any time.

Data ONTAP creates a default /etc/nsswitch.conf file when you run the setup command on your storage system. The contents of the default file are as follows:

```
hosts: files nis dns

passwd: files nis ldap

netgroup: files nis ldap

group: files nis ldap

shadow: files nis
```

**Note:** Only the hosts entry in the /etc/nsswitch.conf file pertains to host-name resolution. For information about other entries, see the *Data ONTAP 7-Mode System Administration Guide* and the na_nsswitch.conf(5) man page.

By default, the host information is searched in the following order:

- /etc/hosts file
- NIS
- DNS

You can change the host-name resolution order in either of the following ways:

- By using FilerView
- By editing the /etc/nsswitch.conf file

## Changing the host-name search order with FilerView

You can use FilerView to change the order in which Data ONTAP searches for host information.

**Steps**

1. From the list on the left pane, click **Network**.

2. In the list under Network, click **Manage DNS** and **NIS Name Service**.

**3.** In the Name Service section, select the desired values from the **Hosts** drop-down list.

## Changing the host-name search order

You can change the order in which Data ONTAP searches for host information by editing the /etc/nsswitch.conf file.

**Steps**

**1.** If the /etc/nsswitch.conf file does not exist in the root volume of the storage system, create it.

**2.** Edit the file, entering each line in the following format:

**hosts: *service***

*service* is one or more of the following: files, dns, nis.

**3.** Save the file.

---

**Example**

To change the resolution order to use NIS exclusively, change the hosts line to read as follows:

**hosts: nis**

---

# How VLANs work

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.

## 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. The switch ports can be configured to belong to one or more VLANs (static registration), or end-stations can register their VLAN membership dynamically, with VLAN-aware switches.

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.

### How VLAN membership affects communication

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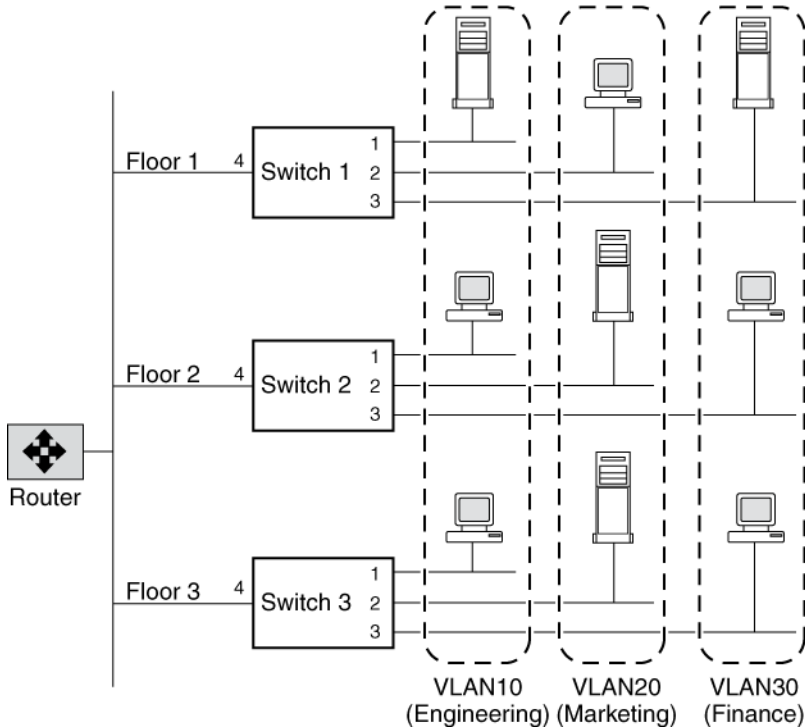
In this figure, VLAN 10 (Engineering), VLAN 20 (Marketing), and VLAN 30 (Finance) span three floors of a building. If a member of VLAN 10 on Floor 1 wants to communicate with a member of VLAN 10 on Floor 3, the communication occurs without going through the router, and packet flooding is limited to port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

**Related concepts**

*VLAN membership* on page 87

# GARP VLAN Registration Protocol

GARP VLAN Registration Protocol (GVRP) uses Generic Attribute Registration Protocol (GARP) to allow end-stations on a network to *dynamically* register their VLAN membership with GVRP-aware switches. Similarly, these switches dynamically register with other GVRP-aware switches on the network, thus creating a VLAN topology across the network.

GVRP provides dynamic registration of VLAN membership; therefore, members can be added or removed from a VLAN at any time, saving the overhead of maintaining static VLAN configuration on switch ports. Additionally, VLAN membership information stays current, limiting the broadcast domain of a VLAN only to the active members of that VLAN.

For more information about GVRP and GARP, see IEEE 802.1Q and IEEE 802.1p (incorporated in the 802.1D:1998 edition).

## GVRP configuration for VLAN interfaces

By default, GVRP is disabled on all VLAN interfaces in Data ONTAP; however, you can enable it.

After you enable GVRP on an interface, the VLAN interface informs the connecting switch about the VLANs it supports. This information (dynamic registration) is updated periodically. This information is also sent every time an interface comes up after being in the down state or whenever there is a change in the VLAN configuration of the interface.

### Related tasks

# VLAN tags

A VLAN tag is a unique identifier that indicates the VLAN to which a frame belongs. Generally, 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.
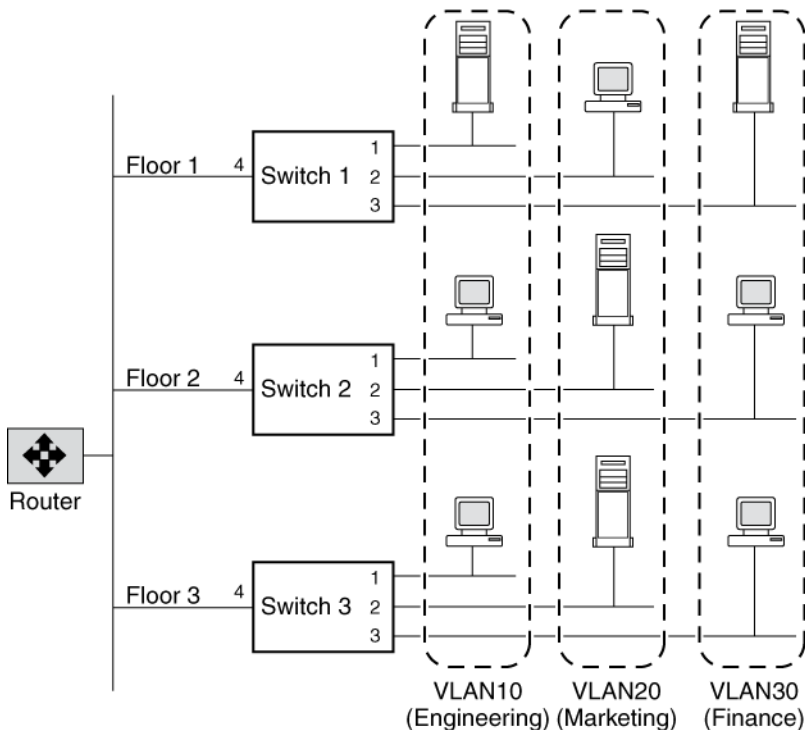
For example, in this figure, if a member of VLAN 10 on Floor 1 sends a frame for a member of VLAN 10 on Floor 2, Switch 1 inspects the frame header for the VLAN tag (to determine the VLAN) and the destination MAC address. The destination MAC address is not known to Switch 1. Therefore, the switch forwards 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.

# Advantages of VLANs

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

VLANs provide the following advantages:

*   Ease of administration

    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 function, 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.
*   Confinement of broadcast domains

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.
- Reduction in network traffic

  Confinement of broadcast domains on a network significantly reduces traffic.
- Enforcement of security policies

  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.

# Prerequisites for setting up VLANs

You must meet certain prerequisites for switches and end-stations before you can set up VLANs in a network.

The following are the prerequisites for setting up VLANs:

- The switches deployed in the network either must comply with IEEE 802.1Q standards or must have a vendor-specific implementation of VLANs.
- For an end-station to support multiple VLANs, it must be able to dynamically register (using GVRP) or must be statically configured to belong to one or more VLANs.

# Guidelines for setting up VLANs in Data ONTAP

VLANs in Data ONTAP are implemented in compliance with the IEEE 802.1Q standard. Before creating and configuring LANs, you must consider certain guidelines about the type, MTU size, speed and media of the underlying interfaces.

You should consider the guidelines when setting up VLANs in Data ONTAP:

- You cannot set up VLANs using the setup procedure.

  You must use the command-line interface FilerView to create, change, or delete VLANs.
- You must add the commands to create VLANs on the storage system to the /etc/rc file to make the VLANs persistent across reboots.
- You can create any number of VLANs on a NIC (supporting IEEE 802.1Q) on the storage system.

  However, Data ONTAP imposes a limit on the number of interfaces (including physical, interface group, VLAN, vh, and loopback interfaces) per storage system.
- The maximum number of VLANs that you can configure is determined by the system memory

  - For systems with 2 GB memory, you can create a maximum of 128 VLANs
  - For systems with 6 GB memory, you can create a maximum of 256 VLANs
  - For systems with greater than 6 GB memory, you can create a maximum of 512 VLANs

> **Note:** VLAN tags 0 and 1 are reserved and not supported.

- You can create VLANs on physical interfaces and interface groups.
- You can use VLANs to support packets of different Maximum Transmission Unit (MTU) sizes on the same network interface.

  If a network interface is a member of multiple VLANs, you can specify different MTU sizes for individual VLANs.
- You can assign an identification number ranging from 1 to 4094 to a VLAN.
- You must ensure that the interface on your storage system is also a member of its partner's VLANs in an HA pair.
- You cannot configure any parameters except `mediatype` for the physical network interface configured to handle VLANs.
- You must set the same flow control settings for all the underlying physical network interfaces that constitute a VLAN.

  You should set the flow control settings of all the network interfaces to "full", which is the default flow control type.

**Related concepts**

*Maximum number of network interfaces* on page 23

**Related tasks**

*Specifying the flow control type for a network interface* on page 40

# The vlan command syntax

You can use the `vlan` command to create, add interfaces to, delete, modify, and view the statistics of a VLAN.

The following table gives the syntax of the `vlan` command:

| Command | Description |
|---------|-------------|
| `vlan create [-g {on|off}]` *ifname vlanid_list* | Create a VLAN |
| `vlan add` *ifname vlanid_list* | Add an interface to a VLAN |
| `vlan delete -q` *ifname [vlanid_list]* | Delete an interface from a VLAN |
| `vlan modify -g {on|off}` *ifname* | Enable or disable GVRP on VLAN interfaces |
| `vlan stat ifname` *[vlanid_list]* | View the statistics of the network interfaces of a VLAN |

For more information about the `vlan` command, see the na_vlan(1) man page.

**Note:** The VLANs created or changed using the `vlan` command are not persistent across reboots unless the `vlan` commands are added to the `/etc/rc` file.

# Creating a VLAN

You can create a VLAN for ease of administration, confinement of broadcast domains, reduced network traffic, and enforcement of security policies. You can use the `vlan create` command to include an interface in one or more VLANs as specified by the VLAN identifier, enable VLAN tagging, and optionally enable GVRP.

### About this task

- By default, GVRP is disabled on VLAN interfaces created by using the `vlan create` command; however, you can enable it with the `-g` option of the `vlan create` command.
- VLANs created by using the `vlan create` command are not persistent across reboots unless the `vlan` commands are added to the `/etc/rc` file.
- A VLAN name should not exceed 15 characters. A VLAN is named by combining the base interface name (physical or interface group) and the VLAN identifier. If the resulting VLAN name exceeds 15 characters, the base interface name is truncated and appended to the VLAN identifier with a hyphen (-) in between.
- You should be aware of the limit on the interface name when making an entry in the `/etc/rc` file.

### Step

1. Enter the following command:

   **`vlan create [-g {on|off}] ifname vlanid`**

   `-g` enables (`on`) or disables (`off`) GVRP on an interface. By default, GVRP is disabled on the interface.

   `ifname` is the name of the network interface.

   `vlanid` is the VLAN identifier to which the `ifname` interface belongs. You can include a list of VLAN identifiers.

   > **Example: Creating and naming of VLAN interfaces**
   >
   > Create VLANs with identifiers 10, 20, and 30 on the interface e4 of a storage system by using the following command:
   >
   > **`vlan create e4 10 20 30`**
   >
   > As a result, VLAN interfaces e4-10, e4-20, and e4-30 are created. The `ifconfig` command output displays e4 as a VLAN interface as follows:

```
ifconfig -a
e0a: flags=0x80e08866<BROADCAST,RUNNING,MULTICAST,VLAN> mtu 1500
ether 00:0c:29:56:54:7e (auto-1000t-fd-up) flowcontrol full
```

The following example displays the truncation of the base interface name when creating a VLAN. To create a VLAN on the interface group "reallylongname," enter the following command:

**vlan create reallylongname 100**

The resulting VLAN name is "reallylongn-100". The base interface name is truncated and the VLAN name is restricted to 15 characters. When you edit the /etc/rc file, ensure that you enter the truncated VLAN name.

**After you finish**

You must configure the VLAN interface by using the ifconfig command.

**Related concepts**

*Prerequisites for setting up VLANs* on page 91
*Guidelines for setting up VLANs in Data ONTAP* on page 91

# Configuring a VLAN

After you create a VLAN, you must configure it with an IP address. By using the ifconfig command, you can configure all the parameters for a VLAN interface in the same way that you configure the parameters for a physical interface.

**About this task**

You can configure the following parameters for a VLAN:

- IP address
- Network mask
- Interface status
- Partner

**Step**

1. Enter the following command:

   **ifconfig *ifname-vlanid IP_address* netmask *mask***

   *ifname-vlanid* is the VLAN interface name.

   *IP_address* is the IP address for this interface.

   *mask* is the network mask for this interface.

**Example**

Create VLANs with identifiers 1760 on the interface e5a of a storage system by using the
following command:

**vlan create e5a 1760**

Configure the VLAN interface e5a-1760 by using the following command:

**ifconfig e5a-1760 192.0.2.11 netmask 255.255.255.0**

**Related concepts**

*Configuring network interfaces* on page 37

## 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 use the ifconfig command to configure an IP address for the underlying interface of the
VLAN. Any tagged frame is received by the matching VLAN interface. Untagged traffic is received
by the native VLAN on the underlying interface.

> **Note:** You should not create a VLAN on a physical 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 a network interface that is configured to receive tagged and untagged traffic
(IP configured for the underlying interface and VLANs). You must bring down all VLANs on the
underlying interface before you bring down the interface. However, you can delete the IP address of
the underlying interface.

When you configure a VLAN using GVRP, the GVRP packets are sent as both untagged and tagged
from the native VLAN.

For information about reverting with a configuration for receiving tagged and untagged frames on the
same network interface, see the *Data ONTAP 7-Mode Upgrade Guide*.

# Adding an interface to a VLAN

If a physical interface does not belong to any VLAN, you can use the vlan create command to
make the interface a member of one or more VLANs. However, if the interface is already a member
of a VLAN, you should use the vlan add command to add the interface to subsequent VLANs.

**About this task**

VLANs created using the vlan add commands are not persistent across reboots unless the vlan
commands are added to the /etc/rc file.

**Step**

1. Enter the following command:

   **vlan add *interface_name vlanid***

   *interface_name* is the name of the network interface.

   *vlanid* is the VLAN identifier to which the interface belongs. You can include a list of VLAN identifiers.

   **Example**
   Add VLANs with identifiers 40 and 50 on the interface e4 of a storage system by using the following command:

   **vlan add e4 40 50**

   VLAN interfaces e4-40 and e4-50 are created.

**After you finish**

You must configure the VLAN interface by using the ifconfig command.

**Related tasks**

*Configuring a VLAN* on page 94
*Creating a VLAN* on page 93

# Deleting VLANs

You can delete a specific VLAN or all VLANs that are configured on a network interface. When you delete all VLANs on an interface, the interface is then available to be configured as a regular physical interface.

**Step**

1. Enter the following command:

   **vlan delete [-q] *interface_name***

| If you want to... | Enter the following command: |
| --- | --- |
| Delete one or more specific VLANs | **vlan delete [-q] *interface_name vlanid***<br><br>**Note:** If you want to delete more than one specific VLAN, you can include a list of VLAN identifiers.<br><br>For example, to delete the VLAN e4-30, enter the following command:<br><br>**vlan delete e4 30** |

| If you want to... | Enter the following command: |
|---|---|
| Delete all VLANs configured on a network interface | `vlan delete [-q] interface_name`<br><br>For example, to delete all VLANs configured on the interface e4, enter the following command:<br><br>`vlan delete e4` |

`interface_name` is the name of the network interface.

`vlanid` is the VLAN identifier to which the `interface_name` interface belongs. You can include a list of VLAN identifiers.

`-q` option invokes the quiet mode.

**Result**

By default, the `vlan delete` command prompts you to confirm the deletion.

**Note:** If you do not want to receive this prompt, use the `-q` option with the `vlan delete` command. This action invokes the quiet mode, which causes the operation to complete without prompting.

# Enabling or disabling GVRP on your VLAN interface

GVRP dynamically registers the VLAN memberships of stations on your network. This reduces the overhead of maintaining static VLAN configuration on switch ports every time a change occurs in your network. To enable or disable GVRP on all interfaces of a network adapter, you should use the `vlan modify` command.

**About this task**

- When you enable GVRP on a network interface, it is enabled on all the associated VLANs.
  For example, you can enable GVRP only on the network interface e8 of a storage system. However, you cannot enable or disable GVRP for the VLAN e8-2.
- If you enable GVRP on an interface that is configured to the `down` status, the state of the interface and all associated VLAN interfaces is automatically configured to the `up` status.
  This state change occurs so that the interface can start sending VLAN registration frames to register its VLAN membership with the switch.
- VLANs modified using the `vlan modify` command are not persistent across reboots unless the `vlan` commands are added to the `/etc/rc` file.

**Step**

1. Enter the following command:

```
vlan modify -g {on|off} adap_name
```

`-g on` enables GVRP.

`-g off` disables GVRP.

`adap_name` is the name of the network adapter.

**Related concepts**

*GARP VLAN Registration Protocol* on page 88
*GVRP configuration for VLAN interfaces* on page 89

# Viewing VLAN statistics

You can use the `vlan stat` command to view the statistics of all VLANs configured on a network interface. You can view the frames received and transmitted on an interface and the number of frames that were rejected because the frames did not belong to any of the VLAN groups.

**Step**

1. Enter the following command:

   ```
   vlan stat interface_name
   ```

   `interface_name` is the name of the network interface.

   **Example**
   The following example displays the statistics of all VLANs on a storage system:

   ```
   vlan stat e4
   Vlan Physical Interface: e4 (5 hours, 50 minutes, 38 seconds) --
   Vlan IDs: 3,5
   GVRP: enabled

   RECEIVE STATISTICS
   Total frames: 0 | Total bytes: 0 |Multi/broadcast: 0

   Untag drops:0 | Vlan tag drops: 0

   TRANSMIT STATISTICS
   Total frames: 8 | Total bytes: 368

   Vlan Interface: e4-3 (0 hours, 20 minutes, 45 seconds) --

   ID: 3 MAC Address: 00:90:27:5c:58:14
   ```

# Viewing statistics for a specific VLAN

You can use the `vlan stat` command to view the statistics for a specific VLAN configured on a network interface. You can view the frames received and transmitted on an interface and the number of frames that were rejected because the frames did not belong to any of the VLAN groups.

**Step**

1. Enter the following command:

   **`vlan stat interface_name vlanid`**

   `interface_name` is the name of the network interface.

   `vlanid` is the VLAN identifier to which the `interface_name` interface belongs. You can include a list of VLAN identifiers.

# How interface groups work in Data ONTAP

An interface group is a feature in Data ONTAP that implements link aggregation on your storage system. Interface groups provide a mechanism to group together multiple network interfaces (links) into one logical interface (aggregate). After an interface group is created, it is indistinguishable from a physical network interface.

The following figure shows four separate network interfaces, e3a, e3b, e3c, and e3d, before they are grouped into an interface group.



The following figure shows the four network interfaces grouped into a single interface group called Trunk1.



Different vendors refer to interface groups by the following terms:

- Virtual aggregations
- Link aggregations
- Trunks
- EtherChannel

Interface groups provide several advantages over individual network interfaces:

- Higher throughput

  Multiple interfaces work as one interface.
- Fault tolerance

  If one interface in an interface group goes down, your storage system stays connected to the network by using the other interfaces.
- No single point of failure

  If the physical interfaces in an interface group are connected to multiple switches and a switch goes down, your storage system stays connected to the network through the other switches.

# Types of interface groups

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

The following figure is an example of a single-mode interface group. In the figure, e0 and e1 are part of the SingleTrunk1 single-mode interface group. If the active interface, e0, fails, the standby e1 interface takes over and maintains the connection to the switch.

## Static multimode interface group

The static multimode interface group implementation in Data ONTAP is in compliance 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 support IEEE 802.3ad (dynamic), also known as Link Aggregation Control Protocol (LACP). Port Aggregation Protocol (PAgP), the proprietary link aggregation protocol from Cisco, too is not supported.

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 provides higher throughput than a single-mode interface group. Static multimode interface groups can recover from a failure of up to "n-1" interfaces, where n is the total number of interfaces that form the interface group.

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 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 an administrator wants to manage or control the transmission of inbound traffic on an aggregated link, it must be modified on the directly connected network device.

The following figure is an example of a static multimode interface group. Interfaces e0, e1, e2, and e3 are part of the MultiTrunk1 multimode interface group. All four interfaces in the MultiTrunk1 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 stated to interoperate or conform to the IEEE 802.3 standards, it should operate with Data ONTAP.

## Dynamic multimode interface group

Dynamic multimode interface groups can detect not only the loss of link status (as do static multimode interface groups), but also a loss of data flow. This feature makes dynamic multimode interface groups compatible with high-availability environments. The dynamic multimode interface group implementation in Data ONTAP is in compliance with IEEE 802.3ad (dynamic), also known as Link Aggregation Control Protocol (LACP).

Dynamic multimode interface groups have some special requirements. They include the following:

- Dynamic multimode interface groups must be connected to a switch that supports LACP.
- Dynamic multimode interface groups must be configured as first-level interface groups.
- Dynamic multimode interface groups should be configured to use the port-based and IP-based load-balancing methods.

In a dynamic multimode interface group, all interfaces in the interface group are active and share a single MAC address. This logical aggregation of interfaces provides higher throughput than a single-mode interface group.

A dynamic 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. For information about configuring the switch, see your switch vendor's documentation. Some switches might not support link aggregation of ports configured for jumbo frames. For more information, see your switch vendor's documentation.

> **Attention:** Data ONTAP supports only the active and passive modes of LACP.

Several load-balancing options are available to distribute traffic among the interfaces of a dynamic multimode interface group.

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 an administrator wants to manage or control the transmission of inbound traffic on an aggregated link, it must be modified on the directly connected network device.

The following figure is an example of a dynamic multimode interface group. Interfaces e0, e1, e2, and e3 are part of the MultiTrunk1 multimode interface group. All four interfaces in the MultiTrunk1 dynamic 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. Dynamic multimode interface groups in Data ONTAP conform to the IEEE 802.3 standards. If a particular multiple switch link aggregation technology is stated to interoperate or conform to the IEEE 802.3 standards, it should operate with Data ONTAP.

# 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. If no method is specified, the IP address based load-balancing method is used.

## 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.

## 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 can lead to 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.

# Guidelines for configuring interface groups

Before creating and configuring interface groups, you must follow certain guidelines about the type, MTU size, speed, and media of the underlying interfaces.

The following guidelines apply when you create and configure interface groups on your storage system:

- It is best that the network interfaces that are part of an interface group are on the same network adapter.
- In a single interface group, you can group up to 16 physical Ethernet interfaces.
- You cannot include a VLAN interface in an interface group.
- The interfaces that form an interface group must have the same Maximum Transmission Unit (MTU) size.
  If you attempt to create or add to an interface group and the member interfaces have different MTU sizes, Data ONTAP automatically modifies the MTU size to be the same. To ensure that the desired MTU size is configured, you can use the ifconfig command to configure the MTU size of the interface group after it is created. You need to configure the MTU size only if you are enabling jumbo frames on the interfaces.
- You can include any interface, except the e0M management interface that is present on some storage systems.
- Do not mix interfaces of different speeds or media in the same multimode interface group.
- You should set the same flow control settings for all the underlying physical network interfaces that constitute an interface group.
  It is best to set the flow control settings of all the network interfaces to **full**, the default control type.

Some switches might not support multimode link aggregation of ports configured for jumbo frames. For more information, see your switch vendor's documentation.

**Related tasks**

# The ifgrp command

You can manage interface groups on your storage system by using the `ifgrp` command. This command enables you to create interface groups, add interfaces to interface groups, delete interfaces from interface groups, view status and statistics of interface groups, and destroy interface groups.

The following table gives the `ifgrp` command syntax:

| Command | Description |
|---|---|
| **ifgrp create [single\|multi\|lacp]** ***ifgrp_name* -b [rr\|mac\|ip]** **[*interface_list*]** | Create a single-mode or multimode interface group |
| ifgrp {favor\|nofavor} *interface_name* | Designate a favored or nonfavored interface in a single-mode interface group |
| ifgrp add *ifgrp_name interface_list* | Add network interfaces to an interface group |
| ifgrp delete *ifgrp_name* *interface_name* | Delete a network interface from an interface group |
| ifgrp destroy *ifgrp_name* | Destroy an interface group |
| ifgrp status [*ifgrp_name*] | View the status of an interface group |
| ifgrp stat *ifgrp_name* [*interval*] | View the statistics of data packets on the network interfaces of an interface group |

The following `ifgrp` commands are not persistent if used from the command-line interface; however, you can put any of these commands in the `/etc/rc` file to make it persistent across reboots:

- ifgrp create
- ifgrp add
- ifgrp delete
- ifgrp destroy
- ifgrp favor
- ifgrp nofavor

# Creating a single-mode interface group

You can create a single-mode interface group in which only one interface is active at a time and the others are ready to take over if the active interface fails. A single-mode interface group increases the redundancy for providing more availability.

**Before you begin**

- Decide on a case-sensitive name for the interface group that meets the following criteria:

  - It must begin with a letter.
  - It must not contain any spaces.
  - It must not contain more than 15 characters.
  - It must not already be in use for an interface group.
- Decide on a list of the interfaces you want to combine into the interface group.
- To make a specific interface active, you must specify that interface as preferred by using the `ifgrp favor` command; otherwise, an interface is randomly selected to be the active interface.

**Steps**

1. Configure all interfaces that are to be included in the interface group to the `down` status by entering the following command:

   **`ifconfig interface_list down`**

   *interface_list* is a list of the interfaces you want as part of the interface group.

   **Example**

   **`ifconfig e0a e0b down`**

2. To create an interface group, enter the following command:

   **`ifgrp create single ifgrp_name [interface_list]`**

   *ifgrp_name* is the name of the interface group.

   *interface_list* is a list of the interfaces you want as part of the interface group.

   **Note:** The operation performed using the `ifgrp create` command is not persistent across reboots unless you add the command to the `/etc/rc` file.

3. To configure the interface group, enter the following command:

   **`ifconfig ifgrp_name IP_address`**

   *ifgrp_name* is the name of the interface group.

   *IP_address* is the IP address for this interface.

---

**Example: Creating a single-mode interface group**

The following command creates a single-mode interface group SingleTrunk1:

**`ifgrp create single SingleTrunk1 e0 e1`**

The following command configures an IP address of 192.0.2.4 and a netmask of 255.255.255.0 on the single-mode interface group SingleTrunk1:

**`ifconfig SingleTrunk1 192.0.2.4 netmask 255.255.255.0`**

---

**Related concepts**

*Single-mode interface group* on page 102

**Related tasks**

*Changing the status of an interface* on page 45

## Selecting an active interface in a single-mode interface group

When you create a single-mode interface group, an interface is randomly selected to be the active interface. You can specify another interface as active—for example, when you add a higher speed or higher bandwidth interface—by using the ifgrp favor command to override the random selection.

**Step**

**1.** Enter the following command:

**`ifgrp favor interface_name`**

*interface_name* is the name of the interface that you want to specify as active.

**Example**
To specify the interface e1 as preferred, enter the following command:

**`ifgrp favor e1`**

> **Note:** The operation performed using the ifgrp favor command is not persistent across reboots unless the command is added to the /etc/rc file.

**Related concepts**

*Single-mode interface group* on page 102

**Related tasks**

*Designating a nonfavored interface in a single-mode interface group* on page 110

## Designating a nonfavored interface in a single-mode interface group

When you create a single-mode interface group, an interface is randomly selected to be the active interface. You can designate an interface as nonfavored so that it is not considered during the random selection of an active interface in a single-mode interface group.

### About this task

The interface marked as nonfavored can become the active interface when all other interfaces in a single-mode interface group fail. Even after other interfaces come to the up state, a nonfavored interface continues to remain the active interface until it fails or until you, the system administrator, change the active interface by using the ifgrp favor command.

### Step

1. Enter the following command:

   **ifgrp nofavor *interface_name***

   *interface_name* is the name of the interface you do not want to be considered during the random selection of an active interface.

   **Note:** The operation performed using the ifgrp nofavor command is not persistent across reboots unless the command is added to the /etc/rc file.

   ### Example
   Specify the interface e2 to be nonfavored with the following command:

   **ifgrp nofavor e2**

### Related concepts

*Single-mode interface group* on page 102

### Related tasks

*Selecting an active interface in a single-mode interface group* on page 109

## Failure scenarios for a single-mode interface group

A single-mode interface group fails when the link status of the interface group is down. Failure can also occur if link-monitoring Address Resolution Protocol (ARP) packets do not reach any of the interfaces that form the interface group.

When the link status of a single-mode interface group is configured to the down status, it signals that the interfaces that are part of the interface group have lost connection with the switch.

Link-monitoring ARP packets are sent over the ports of a single-mode interface group to detect whether the ports are in the same broadcast domain. If these ARP packets do not reach any of the interfaces in the interface group, the interface group is configured to the down status. To avoid this

problem, you must ensure that all the interfaces of a single-mode interface group are in the same broadcast domain (for example, a LAN or a VLAN).

**Related concepts**

*Single-mode interface group* on page 102

**Related tasks**

*Viewing interface group status* on page 116

# Creating a static multimode interface group

You can use the ifgrp create command to create a static multimode interface group. If you do not specify the type of interface group in the ifgrp create command, a static multimode interface group is created by default.

**Before you begin**

You must meet the following prerequisites to create a multimode interface group:

- Identify or install a switch that supports link aggregation over multiple port connections in your network, configured according to your switch vendor's instructions.
- Decide on a case-sensitive name for the interface group that meets the following criteria:

  - It must begin with a letter.
  - It must not contain a space.
  - It must not contain more than 15 characters.
  - It must not already be in use for an interface group.
- Decide on the interfaces that you want to select as part of the interface group.
- Configure all interfaces that are to be included in the interface group to the down status, by using the ifconfig command.

**About this task**

You can improve throughput by creating a static multimode interface group. With a multimode interface group, all interfaces in the interface group are active and share a single MAC address. This logical aggregation of interfaces provides higher throughput than a single-mode interface group.

**Steps**

1. To create the interface group, enter the following command:

   **ifgrp create multi *ifgrp_name* -b {rr|mac|ip|port} [*interface_list*]**

   -b describes the load-balancing method.

rr specifies round-robin load balancing.

mac specifies MAC address load balancing.

> **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 indicates IP address load balancing (default).

port indicates port-based load balancing.

*ifgrp_name* is the name of a previously created interface group.

*interface_list* is a list of the interfaces you want to add to the interface group.

**Example**

To create a static multimode interface group, comprising interfaces e0, e1, e2, and e3 and using
MAC address load balancing, enter the following command:

**ifgrp create multi MultiTrunk1 -b mac e0 e1 e2 e3**

**2.** To configure the interface group, enter the following command:

**ifconfig *ifgrp_name IP_address* netmask *mask***

**Related concepts**

**Related tasks**

# Creating a dynamic multimode interface group

With a dynamic multimode interface group, all interfaces in the interface group are active and share a
single MAC address. This logical aggregation of interfaces provides higher throughput than a single-
mode interface group. Dynamic multimode interface groups can detect both loss of link and loss of
data flow.

**Before you begin**

You must meet the following prerequisites to create a multimode interface group:

- Identify or install a switch that supports LACP over multiple port connections in your network,
  configured according to your switch vendor's instructions.
- Decide on a case-sensitive name for the interface group that meets the following criteria:

- It must begin with a letter.
- It must not contain a space.
- It must not contain more than 15 characters.
- It must not already be in use for an interface group.
- Decide on the interfaces you want to select as part of the interface group.
- Configure all interfaces that are to be included in the interface group to the `down` status, by using the `ifconfig` command.

**About this task**

Data ONTAP logs information about the LACP negotiation for dynamic multimode interface groups in the `/vol0/etc/log/lacp_log` file.

**Steps**

**1.** To create a dynamic multimode interface group, enter the following command:

**ifgrp create lacp *ifgrp_name* -b {rr|mac|ip|port} [*interface_list*]**

`-b` specifies the load-balancing method.

`rr` specifies round-robin load balancing.

`mac` specifies MAC address load balancing.

> **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` specifies IP address based load balancing (default).

`port` indicates port-based load balancing.

*ifgrp_name* is the name of a previously created interface group.

*interface_list* is a list of the interfaces that form the interface group.

**Example**
To create a dynamic multimode interface group, comprising interfaces e0, e1, e2, and e3 and using IP address based load balancing, enter the following command:

**ifgrp create lacp MultiTrunk1 -b ip e0 e1 e2 e3**

**2.** To configure the dynamic multimode interface group, enter the following command:

**ifconfig *ifgrp_name IP_address* netmask *mask***

**Related concepts**

**Related tasks**

# Adding interfaces to an interface group

You can add one or more interfaces to an interface group to expand your network. You can add
physical interfaces to an interface group any time after you create it by using the ifgrp add
command.

**Before you begin**

- You must configure additional ports on the switch where the new interfaces will be connected.
  For information about configuring the switch, see your switch vendor's documentation.
- The interface to be added to the interface group must be configured to the down status by using
  the ifconfig command.

**Step**

1. Enter the following command:

   **ifgrp add *ifgrp_name interface_list***

   *ifgrp_name* is the name of a previously configured interface group.

   *interface_list* is a list of the interfaces you want to add to the interface group.

   **Note:** The operation performed using the ifgrp add command is not persistent across reboots
   unless the command is added to the /etc/rc file.

   **Example**
   To add the interface e4 to the multimode interface group MultiTrunk1, enter with the following
   command:

   **ifgrp add MultiTrunk1 e4**

**Related tasks**

# Deleting interfaces from an interface group

You might have to delete a physical interface from an interface group—for example, when the interface needs maintenance or when you want to use the interface for a different purpose. You can use the ifgrp delete command to delete one or more interfaces from an interface group.

**Before you begin**

You must configure the interface group to the down state before you can delete a network interface from the interface group. You can configure the interface group to the down status by using the following command:

**ifconfig *ifgrp_name* down**

*ifgrp_name* is the name of the interface group that you want to configure to the down status.

**About this task**

The operation performed using the ifgrp delete command is not persistent across reboots unless the command is added to the /etc/rc file.

**Step**

1. Enter the following command:

   **ifgrp delete *ifgrp_name interface***

   *ifgrp_name* is the name of the interface group.

   *interface* is the interface of the interface group you want to delete.

   **Example**
   To delete the interface e4 from a multimode interface group MultiTrunk1, enter the following commands:

   **ifconfig MultiTrunk1 down**

   **ifgrp delete MultiTrunk1 e4**

**Related tasks**

# Viewing interface group status

You can view the current status of a specified interface group or all single-mode and multimode interface groups on your storage system.

**Step**

1. Enter the following command:

   **ifgrp status [*ifgrp_name*]**

   *ifgrp_name* is the name of the interface group whose status you want to display.

   If you do not specify the interface group name, the status of all interface groups is displayed.

---

**Example**

The following example displays the status of the interface group ifgrp1:

```
ifgrp status ifgrp1
default: transmit 'IP Load balancing', Ifgrp Type 'multi_mode', fail 'log'
ifgrp1: 1 link, transmit 'none', Ifgrp Type 'single_mode' fail 'default'
         Ifgrp Status   Up      Addr_set
        up:
        e0b: state up, since 23Jun2009 08:18:12 (00:01:16)
                mediatype: auto-1000t-fd-up
                flags: enabled
                input packets 54, input bytes 4858
                output packets 35, output bytes 2830
                output probe packets 0, input probe packets 0
                strike count: 0 of 10
                up indications 1, broken indications 0
                drops (if) 0, drops (link) 0
                indication: up at 23Jun2009 08:18:12
                        consecutive 75, transitions 1
        broken:
        e0c: state broken, since 23Jun2009 08:18:22 (00:01:06)
                mediatype: auto-unknown-down
                flags:
                input packets 0, input bytes 0
                output packets 0, output bytes 0
                output probe packets 0, input probe packets 0
                strike count: 0 of 10
                up indications 0, broken indications 0
                drops (if) 0, drops (link) 0
                indication: broken at 23Jun2009 08:18:22
                        consecutive 0, transitions 1
```

---

## What the interface group status information table contains

You can view the status information of an interface group by using the ifgrp status command.

The following table describes the information that is shown in each field and subfield of the ifgrp status command output.

| Field | Subfield | Description |
|---|---|---|
| `default` | | Indicates the default values for fields such as transmit, Ifgrp Type, and fail. These values apply if you do not specify any values for these fields when creating an interface group. |
| | `transmit` | Indicates the default load-balancing method. |
| | `Ifgrp Type` | Indicates the default interface group type. |
| | `fail` | Indicates the default location where the errors are logged. |
| `ifgrp_name` | | Indicates that the data that follows this field pertains to the interface group, `ifgrp_name`. |
| | `transmit` | Indicates the load-balancing method used. |
| | `Ifgrp Type` | Indicates the type of interface group. Possible values are single-mode, multi_mode, or lacp. |
| | `fail` | Indicates the location where errors are logged for the interface group. |
| | `Ifgrp Status` | Indicates the current status of the interface group, `ifgrp_name`. |
| | `Addr_set` | Indicates that a MAC address has been configured for the interface group, `ifgrp_name`, and all its interfaces. |
| | `state` | Indicates the current link-state of the interface. Possible values are up or down. |
| | `since` | Indicates the date, time, and number of hours since the interface has been up. |
| | `mediatype` | Indicates the media type that defines the speed and duplex for that interface. |
| | `flags` | Indicates whether the interface is enabled or disabled to send and receive data. |
| | `strike count` | Indicates the number of attempts for link-monitoring. When an underlying link of an interface group does not receive any packets (including ARP packets that are used for link-monitoring), the strike count gets incremented once in 5 seconds. If this strike count reaches 10, the underlying link is brought "down." |

| Field | Subfield | Description |
|-------|----------|-------------|
| | consecutive | Indicates the number of consecutively received "up" or "broken" indications from the switch and link interaction. |
| | transitions | Indicates the number of indications received that caused a state transition from "up" to "broken" or "down" to "up". |

# Viewing interface group statistics

You can view the statistics for a specific interface group or for all interface groups. The statistics include the number of packets received and sent by each interface group.

**Step**

1. Enter the following command:

   **ifgrp stat [*ifgrp_name*] [*interval*]**

   *ifgrp_name* is the name of the interface group. If you do not specify an interface group, the status of all interface groups is displayed.

   *interval* is the interval, in seconds. The default is one second.

---

**Example**

The following example displays the output of the ifgrp stat command for a multimode interface group created with the round-robin load-balancing method:

```
ifgrp stat ifgrp0
ifgrp (trunk) ifgrp0
         e3a                       e3b
  Pkts In    Pkts Out    Pkts In    Pkts Out
  8637076    47801540    158        159
  1617       9588        0          0
  1009       5928        0          0
  1269       7506        0          0
  1293       7632        0          0
  920        5388        0          0
  1098       6462        0          0
  2212       13176       0          0
  1315       7776        0          0
```

The first row of the output shows the total number of packets received and sent until the time the ifgrp stat command was run. The following rows show the total number of packets received and sent per second thereafter.

For interface groups created with the round-robin load-balancing option, the outgoing packets are balanced among the network interfaces of the interface group.

```
ifgrp stat ifgrp1
Interface group (trunk) ifgrp1
        e0c                       e0b
  Pkts In    Pkts Out   Pkts In    Pkts Out
  82         208k       796k       208k
  1          27342      104774     27326
  2          26522      102088     26560
  8          20332      77275      20335
  5          27198      103529     27186
```

# Destroying an interface group

You destroy an interface group when you no longer need it or when you want to use the interfaces that form the interface group for other purposes. After you destroy the interface group, the interfaces in the interface group act individually rather than as an aggregate.

**Steps**

1. Configure the interface group to the down status by entering the following command:

   **ifconfig *ifgrp_name* down**

   *ifgrp_name* is the name of the interface group you want to configure to the down status.

2. Enter the following command:

   **ifgrp destroy *ifgrp_name***

   *ifgrp_name* is the name of the interface group you want to destroy.

# Second-level interface groups

If you have more than one multimode interface group, you can use the ifgrp create command to group them by creating a second layer of interface group called the *second-level interface group*. Second-level interface groups enable you to provide a standby multimode interface group in case the primary multimode interface group fails.

You can use second-level interface groups on a single storage system or in an HA pair.

**Note:** You cannot use LACP interface groups as second-level interface groups.

## Guidelines for creating a second-level interface group

You can create a single-mode second-level interface group over two multimode interface groups. The ports of the underlying multimode interface groups should be connected to the same switch. If you

create a second-level interface group over two multimode interface groups that are connected to two different switches, you must connect the two 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. If the ports are not in the same broadcast domain and if no data traffic is detected over the ports, the interface group is configured to the down status.

When the ports of a single-mode interface group are connected to different broadcast domains, it is called a *split-network condition*, which is difficult to diagnose for problems. To avoid such a situation, if a second-level interface group over two multimode interface groups are connected to two different switches, you must connect the switches by using an ISL.

## Creating a second-level interface group

You can create a second-level interface group by using two multimode interface groups. Second-level interface groups enable you to provide a standby multimode interface group in case the primary multimode interface group fails.

### Before you begin

You must meet the following prerequisites to create a second-level interface group:

- Identify or install a switch that supports link aggregation over multiple port connections in your network, configured according to your switch vendor's instructions.
- Decide on a name for the second-level interface group:

  - It must begin with a letter.
  - It must not contain a space.
  - It must not contain more than 15 characters.
  - It must not already be in use for an interface group.
- Decide on a list of the interfaces you want to select as part of the interface group.
- Configure all interfaces that are to be included in the interface group to the down status, by using the ifconfig command.

### Steps

1. Enter the following command to create the first of two multimode interfaces:

   **ifgrp create multi -b {rr|mac|ip|port} *ifgrp_name1 if1 if2***

   The *ifgrp_name1* interface group is composed of two physical interfaces, *if1* and *if2*.

   –b—specifies the type of load-balancing method.

   rr—specifies the round-robin load-balancing option.

   mac—specifies the MAC address load-balancing option.

ip—indicates the IP address load-balancing option (default option).

port—indicates the port-based load-balancing option.

2. Enter the following command to create the second of two multimode interfaces:

   **ifgrp create multi -b {rr|mac|ip|port} *ifgrp_name2 if3 if4***

   The *ifgrp_name2* interface group is composed of two physical interfaces, *if3* and *if4*.

3. Enter the following command to create a single-mode interface from the multimode interfaces:

   **ifgrp create single *ifgrp_name ifgrp_name1 ifgrp_name2***

   *ifgrp_name* is the second-level interface group created with two multimode interface groups, *ifgrp_name1* and *ifgrp_name2*.

---

**Example**

Use the following commands to create two interface groups and a second-level interface group. In this example, IP address load balancing is used for the multimode interface groups.

**ifgrp create multi Firstlev1 e0 e1**

**ifgrp create multi Firstlev2 e2 e3**

**ifgrp create single Secondlev Firstlev1 Firstlev2**

---

**Related tasks**

## Enabling failover in a second-level interface group

If the active interface group participating in a second-level interface group experiences a link failure and becomes degraded, failover to the other interface group occurs if the surviving interface group has a higher aggregate bandwidth. However, if the active (now degraded) interface group has the favor attribute set, failover is allowed only if the option ifgrp.failover.link_degraded is on.

**About this task**

- By default, the ifgrp.failover.link_degraded option is off.
- If the active interface group (now degraded) has the favor attribute set and if the ifgrp.failover.link_degraded option is set to off, failover does not occur.

**Step**

1. If a favored interface group becomes degraded, enable failover to an interface group with a higher aggregate bandwidth by entering the following command:

**options ifgrp.failover.link_degraded on**

**Example**

In the following scenario, consider a second-level interface group over two multimode interface groups with the following configuration settings:

- ifgrp1: Multimode interface group over two 1GbE interfaces
- ifgrp2: Multimode interface group over two 1GbE interfaces
- ifgrp3: Single-mode interface group over ifgrp1 and ifgrp2

The following are the possible situations and failover behavior depending on the options set:

- Situation 1: ifgrp1 is the active and neither interface group is favored
  If one of the links in ifgrp1 fails, ifgrp2 becomes the active interface group. This happens regardless of whether the option `ifgrp.failover.link_degraded` is `on` or `off`.
- Situation 2: ifgrp1 is the favored and active interface, and the option `ifgrp.failover.link_degraded` is `off`
  If one of the links in ifgrp1 fails, ifgrp1 continues to be the active interface group.
- Situation 3: ifgrp1 is the favored and active interface, and the option `ifgrp.failover.link_degraded` is `on`
  If one of the links in ifgrp1 fails, ifgrp2 becomes the active interface group.

**Related tasks**

[Selecting an active interface in a single-mode interface group](#) on page 109

# Second-level interface groups in an HA pair

In an HA pair, you can access data from both storage systems even if one of the storage system in the configuration fails.

With a second-level interface group connected in a single-mode configuration, you can maintain connectivity to your storage system even if one of the switches fails. Therefore, by using the two configurations together, you can achieve a fully redundant storage system connectivity architecture.

The following figure shows second-level interface groups in an HA pair.
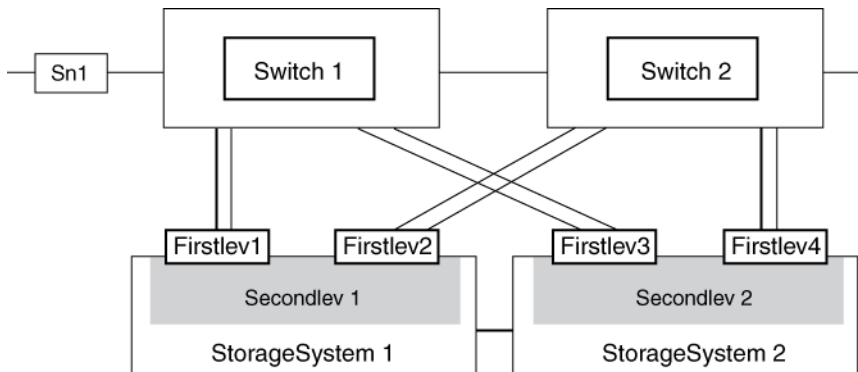
When both storage systems are in operation, the following connections exist:

- Firstlev1 in Secondlev 1 connects StorageSystem 1 to the network through Switch 1.
- Firstlev2 in Secondlev 1 connects StorageSystem 1 to Switch 2.
- Firstlev4 in Secondlev 2 connects StorageSystem 2 to the network through Switch 2.
- Firstlev3 in Secondlev 2 connects StorageSystem 2 to Switch 1.

Firstlev2 and Firstlev3 are in standby mode.

If one of the switches fails, the following happens:

- If Switch 1 fails, Firstlev2 and Firstlev4 maintain the connection for their storage systems through Switch 2.
- If Switch 2 fails, Firstlev1 and Firstlev3 maintain the connection for their storage systems through Switch 1.

In the following figure, Switch 1 fails in an HA pair. Firstlev2 takes over the MAC address of Firstlev1 and maintains the connectivity through Switch 2.
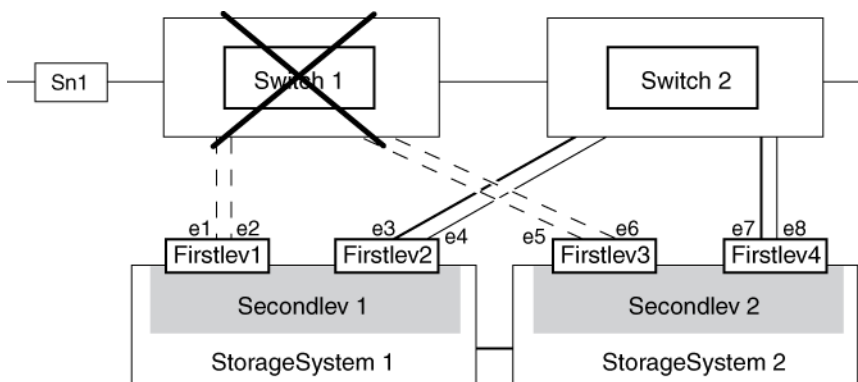
## Creating a second-level interface group in an HA pair

You can create two second-level interface groups in an HA pair so that you can access data from both storage systems even if one of the storage system in the configuration fails.

### Before you begin

You must ensure that all interfaces to be included in the interface group are configured to the `down` status. You can use the `ifconfig` command to configure an interface to the `down` status.

### About this task

The operation performed using the `ifgrp create` command is not persistent across reboots unless the command is added to the `/etc/rc` file.

Assume *StorageSystem1* and *StorageSystem2* are the storage systems that are configured in an HA pair.

### Steps

1. Enter the following commands on *StorageSystem1* to create two multimode interface groups:

   **ifgrp create multi -b {rr|mac|ip|port} *ifgrp_name1 if1 if2***

   **ifgrp create multi -b {rr|mac|ip|port} *ifgrp_name2 if3 if4***

   `-b` specifies the type of load-balancing method.

   `rr` specifies the round-robin load-balancing option.

   `mac` specifies the MAC address load-balancing option.

   `ip` specifies the IP address load-balancing option (default option).

   `port` specifies the port-based load-balancing option.

   *if1*, *if2*, *if3*, *if4* are the network interfaces.

   *ifgrp_name1* and *ifgrp_name2* are the names of the multimode interface groups.

2. Enter the following command on *StorageSystem1* to create a second-level interface from the multimode interface groups:

   **ifgrp create single *secondlev1 ifgrp_name1 ifgrp_name2***

   *secondlev1* is the name of the second-level interface group.

3. Enter the following commands on *StorageSystem2* to create two multimode interface groups:

   **ifgrp create multi -b {rr|mac|ip|port} *ifgrp_name3 if5 if6***

   **ifgrp create multi -b {rr|mac|ip|port} *ifgrp_name4 if7 if8***

**4.** Enter the following command on *StorageSystem2* to create a second-level interface from the multimode interface groups:

**`ifgrp create single secondlev2 ifgrp_name3 ifgrp_name4`**

**5.** Enter the following command on *StorageSystem1* to configure the second-level interface groups for takeover:

**`ifconfig secondlev1 partner secondlev2`**

**6.** Enter the following command on *StorageSystem2* to configure the second-level interface groups for takeover:

**`ifconfig secondlev2 partner secondlev1`**

In steps 5 and 6, secondlev1 and secondlev2 (arguments to the partner option) must be interface names and not interface IP addresses. If secondlev1 is an interface group, secondlev2 can be a an interface group or a physical network interface.

---

**Example**

Use the following commands to create a second-level interface group in an HA pair. In this example, IP-based load balancing is used for the multimode interface groups.

On *StorageSystem1*:

**`ifgrp create multi Firstlev1 e1 e2`**

**`ifgrp create multi Firstlev2 e3 e4`**

**`ifgrp create single Secondlev1 Firstlev1 Firstlev2`**

On *StorageSystem2* :

**`ifgrp create multi Firstlev3 e5 e6`**

**`ifgrp create multi Firstlev4 e7 e8`**

**`ifgrp create single Secondlev2 Firstlev3 Firstlev4`**

On *StorageSystem1*:

**`ifconfig Secondlev1 partner Secondlev2`**

On *StorageSystem2* :

**`ifconfig Secondlev2 partner Secondlev1`**

---

**Related tasks**

# How CDP works with Data ONTAP

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 is a protocol that 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 or 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 controller sends CDPv1 advertisements, the immediately connected CDP-compliant devices send back CDPv1 advertisements.

You should consider the following information before enabling CDP on your storage system:

- CDP advertisements are sent only by the ports that are in the up state.
  In Data ONTAP 8.0.1, the network interface need not be configured with an IP address to be able to send the CDP advertisement.
- CDP must be enabled on both the transmitting and receiving devices for sending and receiving CDP advertisements.
- CDP advertisements are sent at regular intervals. You can configure the time interval.
- 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 IP addresses are changed at the storage system side, 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 storage system 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 IP addresses 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 IP addresses 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 1500 bytes, only the number of IP addresses that can fit into a 1500 MTU-sized packet are advertised.

# Enabling or disabling CDP on your storage system

To discover and send advertisements to CDP-compliant neighboring devices, CDP must be enabled on the storage system. You should use the `cdpd.enable` option to enable or disable CDP on your storage system.

**About this task**

When the `cdpd.enable` option is set to `on`, CDPv1 is enabled on all physical ports of the storage system.

**Step**

1. To enable or disable CDP, enter the following command:

   **`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 by the storage controller in each CDPv1 packet. You can use the `cdpd.holdtime` option to configure hold time.

**About this task**

The value of the `cdpd.holdtime` option applies to both members of an HA pair.

The default value of hold time is 180 seconds.

**Step**

1. To configure the hold time, enter the following command:

   **`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 members of an HA pair.

**Step**

1. To configure the interval for sending CDP advertisements, enter the following command:

   **`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 cleared the last time. To reinitialize the CDP statistics, you can use the `cdpd zero-stats` command.

**Step**

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

| If you want to... | Enter the following command... |
|---|---|
| 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:            9116  | Csum Errors:         0  | Unsupported
Vers:    4561
```

```
 Invalid length:       0  | Malformed:          0  | Mem alloc
fails:        0
 Missing TLVs:         0  | Cache overflow:     0  | Other
errors:          0

TRANSMIT
 Packets:           4557  | Xmit fails:         0  | No
hostname:          0
 Packet truncated:    0  | Mem alloc fails:    0  | Other
errors:          0
```

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

Enter the following command to clear the statistics:

**cdpd zero-stats**

The following output shows the statistics after they are cleared:

```
system1> cdpd show-stats

RECEIVE
 Packets:              0  | Csum Errors:        0  | Unsupported
Vers:     0
 Invalid length:      0  | Malformed:          0  | Mem alloc
fails:      0
 Missing TLVs:        0  | Cache overflow:     0  | Other
errors:         0


TRANSMIT
 Packets:             0  | Xmit fails:         0  | No
hostname:          0
 Packet truncated:    0  | Mem alloc fails:    0  | Other
errors:         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

To verify your network connectivity, you can view information about the neighboring devices connected to each port of your storage system, provided 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 on your storage system.

**About this task**

If a network interface does not have an IP address configured, you can view the CDP information of the storage system from the switch. However, you cannot view the CDP information of the switch from the storage system.

**Step**

1. To view information about all CDP-compliant devices connected to your storage system, enter the following command:

   `cdpd show-neighbors`

---

**Example**

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

```
system1> cdpd show-neighbors
Local  Remote           Remote                Remote           Hold
Remote
Port   Device           Interface             Platform         Time
Capability
------ --------------   --------------------- ---------------- -----
----------
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

# How to monitor your storage system with SNMP

If you enable SNMP in Data ONTAP, the SNMP managers can query your storage system's SNMP agent for information. The SNMP agent gathers information and forwards it to the managers by using SNMP. The SNMP agent also generates trap notifications whenever specific events occur.

For diagnostic and other network management services, Data ONTAP provides an SNMP agent compatible with SNMP versions 1 and 3. SNMPv3 offers advanced security by using pass phrases and encryption. SNMPv3 supports the MIB-II specification and the MIBs of your storage system. The following MIB-II groups are supported:

- System
- Interfaces
- Address translation
- IP
- ICMP
- TCP
- UDP
- SNMP

   **Note:** Transmission and EGP MIB-II groups are not supported.

## Types of SNMP traps in Data ONTAP

SNMP traps capture system monitoring information in Data ONTAP. There are two types of traps in Data ONTAP: built-in and user-defined.

- 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 are based on one of the following:
  - RFC 1213, which defines traps such as coldStart, linkDown, linkUp, and authenticationFailure.
  - Specific traps defined in the custom MIB, such as diskFailedShutdown, cpuTooBusy, and volumeNearlyFull.
- User-defined traps are defined by `snmp traps` commands or the FilerView SNMP Traps windows. These traps are sent using proxy trap ID numbers 11 through 18, which correspond to a trap's MIB priority.

# What a MIB is

A MIB file is a text file that describes SNMP objects and traps. MIBs are not configuration files. Data ONTAP does not read these files and their contents do not affect SNMP functionality.

Data ONTAP provides two MIB files:

- A custom MIB (`/etc/mib/netapp.mib`)
- An Internet SCSI (iSCSI) MIB (`/etc/mib/iscsi.mib`)

Data ONTAP also provides a short cross-reference between object identifiers (OIDs) and object short names in the `/etc/mib/traps.dat` file. This file is useful for creating user-defined traps.

> **Note:** The latest versions of the Data ONTAP MIBs and `traps.dat` files are available online on the IBM NAS support site. 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.

# What the SNMP agent does

The storage system includes an SNMP agent that responds to queries and sends traps to network management stations.

The SNMP agent on the storage system has only read privileges—that is, it cannot be used to take corrective action in response to a trap.

# How to configure the SNMP agent

You need to configure the SNMP agent on your storage system to set SNMP values and parameters. You can configure your SNMP agent through the command-line interface or with FilerView.

To configure the SNMP agent on your storage system, you need to perform the following tasks:

- Verify that SNMP is enabled.

    > **Note:** SNMP is enabled by default in Data ONTAP.

- If you are running SNMPv3, configure SNMPv3 for read-only access.
- Enable traps.
  Although SNMP is enabled by default, traps are disabled by default.
- Specify host names of one or more network management stations.
  Traps can only be sent when at least one SNMP management station is specified as a traphost.
  Trap notifications can be sent to a maximum of eight network management stations.

You can perform the following tasks after configuring SNMP:

- Provide courtesy information about storage system location and contact personnel.
- Set SNMP access privileges. You can restrict SNMP access on a host or interface basis.
- Specify SNMP communities.

  Community strings function as group names to establish trust between SNMP managers and clients. Data ONTAP imposes the following limitations on SNMP communities:

  - No more than eight communities are allowed.
  - Only read-only communities are supported.
- Enable query authentication.

  You can enable authentication failure traps, which are generated when the agent receives queries with the wrong community string, for the SNMP agent. The traps are sent to all hosts specified as traphosts.
- Create and load user-defined traps.

  **Note:** Storage systems in an HA pair can have different SNMP configurations. For more information, see the na_snmp(1) man page.

**Related concepts**

[User-defined SNMP traps](#) on page 141

## Enabling or disabling SNMP using the command-line interface

You can enable or disable SNMP from the command-line interface by entering the options snmp.enable command.

**Step**

1. Enter the following command:

   **options snmp.enable {on|off}**

   on—Enables SNMP

   off—Disables SNMP

## Configuring SNMPv3 users

To access MIB objects by using SNMPv3, you should create users with login-snmp capability.

**Steps**

1. Enter the following command to create a role with login-snmp capability:

   **useradmin role add *role_name* -a login-snmp**

   *role_name* is the role name with login-snmp capability.

**Example**

`useradmin role add myrole1 -a login-snmp`

2. Enter the following command to create a group and add the created role to that group:

`useradmin group add group_name -r role_name`

*group_name* is the group name to which you want to add the created role, *role_name*.

**Example**

`useradmin group add mygroup1 -r myrole1`

3. Enter the following command to create a user and add the user to the created group:

`useradmin user add user_name -g group_name`

*user_name* is the user name belonging to the group, *group_name*.

**Example**

`useradmin user add myuser1 -g mygroup1`

4. Create a password for the new user.

Ensure that the password has a minimum of eight and a maximum of sixteen characters.

5. Enter the snmpwalk command through the system MIB:

`snmpwalk -v 3 -u user_name -l authNoPriv -A password storage_system system`

*password* is the user's password that you entered in Step 3.

*storage_system* is the storage system that contains the MIBs.

**Example**

`snmpwalk -v 3 -u myuser1 -l authNoPriv -A johndoe123 host1 system`

## Setting SNMP access privileges

You can set SNMP access privileges on a host or an interface by using the command-line interface. The snmp.access option defines a method to restrict access to the storage system on a protocol-by-protocol basis.

**About this task**

You cannot set access privileges from FilerView.

**Step**

1. Enter the following command:

`options snmp.access access_spec`

*access_spec* consists of keywords and their values. Access can be allowed or restricted by host name, IP address, and network interface name.

**Example**

To allow access to SNMP for network interfaces e0, e1, and e2, enter the following command:

**options snmp.access if=e0,e1,e2**

**Related tasks**

[Restricting protocol access](#) on page 47

# Viewing or modifying your SNMP configuration from the command-line interface

You can use the snmp  command to view or modify your SNMP configuration values.

**Step**

1. Enter the following command:

   **snmp {*options values*}**

   *options* are the available options for the snmp command, such as authtrap, community, contact, init, location, traphost, and traps.

   *values* are the values that you want to set for a particular option.

**Related references**

[SNMP command syntax](#) on page 138

# Modifying your SNMP configuration from FilerView

You can use FilerView to modify your SNMP configuration.

**Steps**

1. From the list on the left pane, click **SNMP > Configure**.

   The current SNMP configuration is displayed.

2. To set or modify SNMP configuration values, enter configuration values in the drop-down lists or text fields.

3. Click **Apply**.

## SNMP command syntax

If you specify one or more values for an option of the SNMP commands, the value of that option is set or changed. However, if no values are specified, the current value of that option is returned.

The following table describes the syntax and parameters of SNMP commands.

| Command | Description |
|---|---|
| `snmp` | Displays the current values of all SNMP options, such as `init`, `community`, `contact`, and `traphost`. |
| `snmp authtrap [0|1]` | **With a value**: Enables (with value 1) or disables (with value 0) authentication failure traps on the SNMP agent.<br>**Without a value**: Displays the current value of `authtrap` set in Data ONTAP. |
| `snmp community` | Displays the current list of communities. |
| `snmp community add` `rocommunity` | Adds a community.<br>**Default value**: The default community for the SNMP agent in Data ONTAP is public. The only access mode available on storage systems is the default `ro` (read-only). |
| `snmp community delete {all | rocommunity }` | Deletes one or all communities. |
| `snmp contact [contact]` | **With a value**: Sets the contact name for your storage system. You must enclose the contact string in single quotes (' ') if the string contains spaces.<br>You can enter a maximum of 255 characters for the contact information.<br>**Without a value**: Displays the current contact name set in Data ONTAP. |
| `snmp init [0|1]` | **With a value**: Enables (with value 1) or disables (with value 0) built-in traps and the traps defined using the `snmp traps` command.<br>**Without a value**: Displays the current value of `snmp init` in Data ONTAP.<br>**Default value**: By default, SNMP traps are disabled in Data ONTAP; the system uses the equivalent of `snmp init 0`. |

| Command | Description |
|---------|-------------|
| `snmp location [location]` | **With the option**: Sets the location of your storage system. You must enclose the `location` string in single quotes (' ') if the string contains spaces. |
| | **Without the option**: Displays the current location set in Data ONTAP. |
| `snmp traphost [{add\| delete} { hostname/ ipaddress}]` | **With the option**: Adds or deletes SNMP hosts that receive traps from Data ONTAP. |
| | **Without the option**: Displays the current traphosts set in Data ONTAP. |
| `snmp traps [options]` | Displays the list of user-defined traps set in Data ONTAP |

## SNMP security parameters

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

If the security level is set to authNoPriv, authentication is performed by using a user's authKey to sign the message being sent. The authProtocol parameter must be MD5. The authKey parameters are generated from a passphrase that must have a minimum of eight characters.

If the security level is set to authNoPriv, you must enter the following parameters:

| Parameter | Command-line flag | Description |
|-----------|-------------------|-------------|
| securityName | -u *Name* | User name must not exceed 31 characters. |
| authProtocol | -a (MD5) | Authentication type must be *MD5*. **Note:** Data ONTAP does not support SHA authentication protocol. |
| authKey | -A *PASSPHRASE* | Passphrase with a minimum of eight characters. |
| securityLevel | -l (authNoPriv) | Security level: must be Authentication, No Privacy. **Note:** Data ONTAP does not support retrieving MIB values using the AuthPriv and noAuthNoPriv security levels. |
| context | -n *CONTEXTNAME* | Sets the context name used for SNMPv3 messages. |

## Example: SNMP commands

You can use the snmpget, snmpwalk, snmpbulkget, and snmpbulkwalk commands to retrieve information from network elements with SNMP agents.

**snmpwalk**

The following command retrieves all the variables under the system sys1:

```
snmpwalk -Os -c public -v 1 sys1 system
sysDescr.0 = STRING: Data ONTAP Release 7.3.1
sysObjectID.0 = OID: enterprises.789.2.3
sysUpTimeInstance = Timeticks: (121596665) 14 days, 1:46:06.65
sysContact.0 = STRING:
sysName.0 = STRING: sys1.lab.example.com
sysLocation.0 = STRING:
sysServices.0 = INTEGER: 72
```

The following command is an example of an SNMPv3 request to retrieve all variables under the system sys1:

```
snmpwalk -v 3 -u joe -l authNoPriv -A joe12 sys1 system
SNMPv2-MIB::sysDescr.0 = STRING: Data ONTAP Release 7.3.1
SNMPv2-MIB::sysObjectID.0 = OID:
SNMPv2-SMI::enterprises.789.2.3
DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (121622059) 14
days, 1:50:20.59
SNMPv2-MIB::sysContact.0 = STRING:
SNMPv2-MIB::sysName.0 = STRING: sys1.lab.example.com
SNMPv2-MIB::sysLocation.0 = STRING:
SNMPv2-MIB::sysServices.0 = INTEGER: 72
```

**Note:** You need to enter authentication information for using SNMPv3.

**snmpget**

The following command retrieves the system.sysDescr.0 object from the host sys1 by using the public community string:

```
snmpget -c public sys1 system.sysDescr.0
system.sysDescr.0 = Data ONTAP Release 8.0  Mon Mar 16 16:56:43 IST 2010
```

The following command retrieves the value of an ICMP object (OID=56.1.1.1.1) from the host sys1:

```
snmpget -c public -v 2c sys1 .1.3.6.1.2.1.56.1.1.1.1
    56.1.1.1.1.1 = Counter32: 0
```

**snmpbulkget**

The following command retrieves the system object sysDescr.0 and the first three objects in the ifTable:

```
snmpbulkget -v2c -Cn1 -Cr3 -Os -c public sys1 system ifTable
sysDescr.0 = STRING: Data ONTAP Release 7.3.1
ifIndex.1 = INTEGER: 1
```

```
ifIndex.2 = INTEGER: 2
ifDescr.1 = STRING: "lo0"
```

**snmpbulkwalk**

The following command retrieves all the variables under the system sys1:

```
snmpbulkwalk -v2c -Os -c public sys1 system
sysDescr.0 = STRING: Data ONTAP Release 7.3.1
sysObjectID.0 = OID: enterprises.789.2.3
sysUpTimeInstance = Timeticks: (121603434) 14 days, 1:47:14.34
sysContact.0 = STRING:
sysName.0 = STRING: sys1.lab.example.com
sysLocation.0 = STRING:
sysServices.0 = INTEGER: 72
```

The following example shows a part of the output from retrieving all the variables for the UDP object:

```
snmpbulkwalk -c public -v 2c 192.0.2.19 udp
    udp.udpInDatagrams.0 = Counter32: 347
    udp.udpNoPorts.0 = Counter32: 4
    udp.udpInErrors.0 = Counter32: 0
    udp.udpOutDatagrams.0 = Counter32: 138
    udp.udpTable.udpEntry.udpLocalAddress.0.0.0.0.69 = IpAddress: 00 00 00 00
    udp.udpTable.udpEntry.udpLocalAddress.0.0.0.0.111 = IpAddress: 00 00 00 00
```

# User-defined SNMP traps

If the predefined built-in traps are not sufficient to create alerts for conditions you want to monitor, you can create user-defined traps in Data ONTAP.

Before you define a new trap, you should consult the Data ONTAP MIBs to see if any existing traps serve your purpose.

## How SNMP traps work

You can set SNMP traps to inspect the value of MIB variables periodically. Whenever the value of a MIB variable meets the conditions you specify, a trap is sent to the network management stations on the traphost list. The traphost list specifies the network management stations that receive the trap information.

You can set traps on any numeric variable in the MIB. For example, you can set a trap to monitor the fans on your storage system. You can then set the SNMP application on your network management station to display a flashing message on your console when a fan stops working.

Traps are persistent across reboots until you remove it or modify it.

Follow these guidelines when creating traps:

- Use the /etc/mib/traps.dat file to find Object Identifiers (OIDs) for objects in the MIB files of your storage system.
- Ensure that the trap can be generated in the storage system environment.

- Do not set traps on tabular data.

  It is possible to set traps on row entries in a sequence—for example, an entry in a table. However, if the order in the table is changed by adding or removing rows, you will no longer be trapping the same numeric variables.

## How to define or modify a trap

You can define traps or modify traps you have already defined by entering values from the command-line interface, in FilerView, or in a configuration file.

You must supply the following elements when you create or modify traps:

- Trap name

  Trap name is the name of the user-defined trap you want to create or change. A trap name must not have any embedded periods.
- Trap parameters
- Parameter values

**Note:** When you create a user-defined trap, it is initially disabled by default. You must enable a trap before it can be triggered. You enable traps by using the snmp traps command or FilerView.

## Viewing or modifying trap values from the command-line interface

You can view or modify your trap values by using the snmp traps command.

**Step**

1. Enter the following command:

   **snmp traps {*options variables*}**

   *options* are the options for SNMP traps such as walk, load, trapname, and so on.

   *variables* are values for the specified option.

## Viewing or modifying trap values from FilerView

You can use FilerView to view or modify a trap value.

**Steps**

1. From the list on the left pane, click **SNMP > Traps**.
2. Depending on whether you want to create, modify, or view a trap, perform the following step:

| If you want to... | Then... |
| --- | --- |
| Create a new trap | **a.** Click **Add**. |
| | **b.** In the Add an SNMP Trap window, enter the requested information. |
| | **c.** Click **Add** again. |
| View or modify an existing trap | **a.** Click **Manage** for the trap you want to view or modify. |
| | **b.** To modify the trap, click **Modify** in the Manage SNMP Traps window. |

## Defining traps in a configuration file

You can define SNMP traps in a configuration file and then load the file with the snmp traps load command. Data ONTAP automatically backs up your SNMP configuration as Snapshot copies, making it easy to transfer user-defined traps to other storage systems and to recover SNMP configurations in case of a disaster.

**Steps**

1. Create a traps configuration file on your storage system.

   You can choose the name and location of the file.

   **Example**
   /etc/mib/mytraps

2. Enter the traps in the configuration file in the following format:

   **trapname.parmvalue**

   The parameters are the same as those used with the snmp traps command.

3. Test each line of the file by entering the snmp traps command using the command-line interface or by specifying the trap using FilerView. Make corrections as needed.

4. Load the configuration file with the following command:

   **snmp traps load *file_name***

   *file_name* is the name of the configuration file.

   **Example**
   **snmp traps load /etc/mib/mytraps**

**Related references**

## Example: Trap definitions

You can define a group of traps by using the command-line interface or FilerView.

The following example sets a group of traps. The trap descriptions are numbered in brackets.

```
snmp traps cifstotalops.var snmp.1.3.6.1.4.1.789.1.7.3.1.1.1.0
```
[1]

```
snmp traps cifstotalops.trigger level-trigger
```

```
snmp traps cifstotalops.edge-1 1000000
```
[4]

```
snmp traps cifstotalops.interval 10
```
[2]

```
snmp traps cifstotalops.backoff-calculator step-backoff
```
[5]

```
snmp traps cifstotalops.backoff-step 3590
```
[5]

```
snmp traps cifstotalops.rate-interval 3600
```
[3]

```
snmp traps cifstotalops.priority alert
```

```
snmp traps cifstotalops.message snmp.1.3.6.1.4.1.789.1.7.3.1.1.1.0
```

A cifstotalops trap [1] is evaluated every 10 seconds [2]. The value received from the previous evaluation and the current value are used to calculate the number of CIFS operations per hour [3]. If the number exceeds one million [4], the trap triggers and continues to be triggered every hour [5] until the total number of CIFS operations drops below one million.

## Command syntax for SNMP trap parameters

If you specify one or more values for an option of the SNMP commands, the value of that option is set or changed. However, if no values are specified, the current value of that option is returned.

The following table describes the syntax and parameters for the `snmp traps` command.

| Command | Description |
|---|---|
| `snmp traps` | Displays the list of user-defined traps set in Data ONTAP. |
| `snmp traps [enable\|disable\| reset\|delete] trapname` | Enables, disables, resets, or deletes the trap *trapname*. If you do not specify a trap name, all traps defined so far are acted on. |
| `snmp traps walk prefix` | Walks (traverses in order) the trap list by prefix; that is, lists all traps that have names beginning with *prefix*. |

| Command | Description |
|---|---|
| `snmp traps`<br>`load trap_list_filename` | Loads a set of traps from a configuration file. The file contains a list of traps and parameters without the `snmp traps` command preceding each trap. If the specified file name is defaults, traps are read from the `/etc/defaults/` traps file. |
| `snmp traps trapname.parm value` | Defines or changes a user-defined trap parameter. |

## SNMP trap parameters

You must specify certain parameters to create SNMP traps.

The following table lists SNMP trap parameters that you enter with the `snmp traps` command in the command-line interface and the equivalent parameters that you select in FilerView.

| Parameter in command-line interface | Equivalent in FilerView |
|---|---|
| `var` | OID |
| `trigger` | Trigger |
| `edge-1`<br>`edge-2` | Edge 1<br>Edge 2 |
| `edge-1-direction`<br>`edge-2-direction` | Edge 1 Direction<br>Edge 2 Direction |
| `interval` | Interval |
| `interval-offset` | Interval Offset |
| `rate-interval` | Rate Interval |
| `backoff-calculator` | Backoff Style |
| `backoff-step` | Backoff Step |
| `backoff-multiplier` | Backoff Multiplier |
| `priority` | Priority |
| `message` | Not available |

### The var parameter

The `var` parameter associates a user-defined trap name (specified by the `trapname` variable in the `snmp traps` command or Trap Name in FilerView) with a specific MIB object. The MIB object is

specified in the value field of the `snmp traps` command. It must be in the format `snmp.oid`, where `oid` is an Object Identifier (OID).

The `traps.dat` file, which is located in the `/etc/mib` directory on your storage system, can help you determine OIDs. This file maps MIB objects' short names in the Data ONTAP MIB files to their numeric OIDs. For more information about a particular OID, see the MIB.

In FilerView, it is necessary to enter only the numerical OID, and not the "snmp" prefix.

## The trigger parameter

The trigger parameter specifies the type of triggers that you can set for a trap. If a trap is triggered, data about the event that caused the trigger is sent to the network management stations.

You can specify the following values for the `trigger` parameter:

| | |
|---|---|
| **single-edge-trigger** | Triggers a trap and sends data when the value of the trap's MIB variable crosses an edge (a value that you specify) for the first time. |
| **double-edge-trigger** | Triggers a trap and sends data when either of two edges is crossed. A double-edge-trigger enables you to set two edges, each with its own direction. |
| **level-trigger** | Triggers a trap and sends data whenever the trap's value crosses a specified edge value. |
| **change-trigger** | Keeps track of the last value received from the trap. If the current value differs from the previously received value, the trap is triggered. |
| **always-trigger** | Enables a trap to always trigger at the specified evaluation interval (specified by the `interval` parameter). For example, a trap can trigger every 24 hours for the agent to send the total number of CIFS operations to an SNMP manager. |

## The edge-1 and edge-2 parameters

The edge-1 and edge-2 parameters of a trap specify the threshold values that are compared during trap evaluation to determine whether to fire a trap and send data.

The edge-1 parameter specifies the value for the edge in a single-edge-triggered trap or the first edge in a double-edge-triggered trap. The default value for the edge-1 parameter is MAXINT.

The edge-2 parameter specifies the value for the second edge in a double-edge-triggered trap. The default value for the edge-2 parameter is 0.

**Note:** The edge-2 parameter is not displayed in FilerView during trap creation unless double-edge-trigger is selected in the trigger parameter.

### The edge-1-direction and edge-2-direction parameters

The edge-1-direction and edge-2-direction parameters enable you to set or change the direction that is used to evaluate a trap. The edge-triggered traps only send data when the edge is crossed in either the up or down direction.

The default value for the edge-1-direction parameter is `up` and for the edge-2-direction parameter is `down`.

**Note:** You enter the direction values on the same line as the edge value when you run the `snmp traps` command. The edge-2-direction parameter is not displayed in FilerView during trap creation unless double-edge-trigger is selected in the trigger parameter.

### The interval parameter

The `interval` parameter is the time, in seconds, between evaluations of a trap.

A trap can only send data as often as it is evaluated, even if the edge values are exceeded sooner. The default value for the `interval` parameter is 3600.

**Note:** The maximum value that can be specified for the `interval` parameter in Data ONTAP is 2147482 seconds, which is approximately equal to 24 days.

### The interval-offset parameter

The `interval-offset` parameter is the amount of time, in seconds, until the first trap evaluation.

The default value for the `interval-offset` parameter is 0. You can set it to a nonzero value to prevent too many traps from being evaluated at once (for example, at system startup).

### The rate-interval parameter

The `rate-interval` parameter specifies the time, in seconds, in which the change in value of a trap's variable (rate of change) is expressed.

If the `rate-interval` value is set for a trap, the samples of data obtained at the interval points (set using the `interval` parameter) for a trap variable are used to calculate the rate of change. If the calculated value exceeds the value set for the `edge-1` or `edge-2` parameter, the trap is fired.

For example, to obtain the number of CIFS operations per hour, you specify a `rate-interval` of 3600. If `rate-interval` is set to 0, no sampling at interval points occurs and trap evaluation proceeds as with any other kind of trap. The default value for the `rate-interval` parameter is 0.

## The backoff-calculator parameter

The `backoff-calculator` parameter enables you to change the trap evaluation interval for a trap after a trap fires.

After a trap fires and sends data, you might not want it to be evaluated so often. For instance, you might want to know within a minute of when a file system is full, but only want to be notified every hour thereafter that it is still full.

The `backoff-calculator` parameter can take the following values in the value variable field:

- `step-backoff`
- `exponential-backoff`
- `no-backoff`

The default value for the `backoff-calculator` parameter is `no-backoff`.

## The backoff-step parameter

The `backoff-step` parameter specifies the number of seconds by which the trap evaluation interval is increased.

If a trap interval is 10 and its backoff-step is 3590, the trap is evaluated every 10 seconds until it fires the first time and sends data, and once an hour thereafter. The default value for the `backoff-step` parameter is 0.

**Note:** The `backoff step` parameter is not displayed in FilerView during trap creation unless "step" is selected in the Backoff Style field.

## The backoff-multiplier parameter

The `backoff-multiplier` parameter specifies the value by which to multiply a trap's evaluation interval each time it fires.

If you set `backoff-calculator` to `exponential-backoff` and `backoff-multiplier` to 2, the interval doubles each time the trap fires. The default value of the `backoff-multiplier` parameter is 1.

**Note:** The `backoff multiplier` parameter is not displayed in FilerView during trap creation unless "exponential" is selected in the Backoff Style field.

## The priority parameter

The priority parameter sets the priority of a trap. If several traps are scheduled to be triggered at the same time, you can use the priority parameter to decide which trap is serviced first.

The possible values for the priority parameter, from highest to lowest, are as follows:

- `emergency`
- `alert`

- `critical`
- `error`
- `warning`
- `notification`
- `informational`
- `debug`

The default value for the priority parameter is `notification`.

### The message parameter

The `message` parameter specifies a message that goes out with a trap.

The message can be a string of text or simply the SNMP OID, in the form `snmp.oid`. If you specify the OID as your message, Data ONTAP sends the information that was trapped concerning the OID. If you do not specify a message parameter for a trap, when the trap fires you see a string with the numerical OID value and its priority level.

For example, the following string is sent to the network management stations for the trap `cpuUpTime` if the message parameter is not set:

**`cpuUpTime == 10562288.priority == notification`**

**Note:** If the message is a string that includes spaces, you must enclose the string in quotation marks (" ").

You cannot set the `message` parameter in FilerView.

# How to diagnose network problems

You can diagnose problems on your network by using commands such as cdpd, netdiag, ping, and pktt.

**cdpd**    You can use the cdpd command to view information about the devices that advertise themselves by using the CDPv1 protocol. You can also view information about the CDP neighbors of the storage system and thus detect network connectivity.

**netdiag**  You can use the netdiag command to continuously gather and analyze statistics, and perform diagnostic tests. These diagnostic tests identify and report problems with your physical network or transport layers and suggest remedial action.

For more information about the netdiag command along with all available options, see the na_netdiag(1) man page.

**ping**    You can use the ping command to test whether your storage system can reach other hosts on your network.

For more information about the ping command, see the na_ping(1) man page.

**pktt**    You can use the pktt command to trace the packets sent and received in the storage system's network.

For more information about the pktt command, see the na_pktt(1) man page.

**Related references**

[Error codes for the netdiag command](#) on page 191

## Diagnosing transport layer problems

You can use the netdiag -t command to diagnose problems with the transport layer in your storage system.

**Step**

1. Enter the following command:

   **netdiag -t**

   > **Example**
   >
   > A storage system whose TCP window size is smaller than the recommended value displays the following output:
   >
   > ```
   > Performing transport layer diagnostics.....
   > The TCP receive window advertised by CIFS client
   > ```

```
192.0.2.13 is 8760. This is less than the recommended
value of 32768 bytes. You should increase the TCP receive
buffer size for CIFS on the client. Press enter to continue.
```

# Viewing diagnostic results

You can use the netdiag -s command to view a summary of the various diagnostic checks and tests performed on your storage system.

**Step**

1. Enter the following command:

   **netdiag -s**

# How to diagnose ping problems

You can use the Data ONTAP ping throttling mechanism and its ip.ping_throttle.drop_level option to help avoid denial-of-service attacks that can occur when using ICMP.

The ping throttling mechanism is active in intervals of 1 second. If the number of ICMP echo and reply packets that the storage system receives in a one-second interval exceeds the ping throttling threshold, the storage system drops all subsequent packets that are received within that one-second interval.

**Note:** Regardless of whether the ping throttling threshold has been reached, clients that send more than 16 packets per second to a storage system might experience packet loss. To allow clients to send more than 16 packets per second, you must disable ping throttling.

If your storage system supports a very large number of CIFS clients that use ICMP pings to determine CIFS shares accessibility, you can increase the ping throttling threshold value in the ip.ping_throttle.drop_level option.

If a large number of CIFS clients are experiencing temporary or persistent unavailability of the storage system, you should check to see if the ping throttling threshold has been exceeded for the storage system. If so, you can increase the ping throttling threshold value.

## Increasing the ping throttling threshold value

If your storage system supports a very large number of CIFS clients that use ICMP pings to determine CIFS shares accessibility, you might need to increase the ping throttling threshold value.

**Step**

1. Enter the following command:

```
options ip.ping_throttle.drop_level packets_per_second
```

*packets_per_second* specifies the maximum number of ICMP echo or echo reply packets (ping packets) that the storage system accepts per second. Any further packets within 1 second are dropped. The default value is 150.

## Checking the ping throttling threshold status

If a large number of CIFS clients are experiencing temporary or persistent unavailability of the storage system, you should check if the ping throttling threshold has been exceeded for the storage system.

### Step

1. Enter the following command:

   ```
   netstat -p icmp
   ```

### Result

The resulting report lists the number of pings and ping replies that have been dropped, if any.

If the number of pings dropped, the number of ping replies dropped, or the number of both pings and ping replies dropped is greater than zero, you should increase the ip.ping_throttle.drop_value option to a number that is higher than the current value.

## Disabling ping throttling

To allow clients to send more than 16 packets per second, you need to disable ping throttling.

### Step

1. Enter the following command:

   ```
   options ip.ping_throttle.drop_level 0
   ```

# Protecting your storage system from forged ICMP redirect attacks

You can disable ICMP redirect messages to protect your storage system against forged ICMP redirect attacks.

### About this task

To efficiently route a series of datagrams to the same destination, your storage system maintains a route cache of mappings to next-hop gateways. If a gateway is not the best next-hop for a datagram with a specific destination, the gateway forwards the datagram to the best next-hop gateway and sends an ICMP redirect message to the storage system. By forging ICMP redirect messages, an attacker can modify the route cache on your storage system, causing it to send all of its

communications through the attacker. The attacker can then hijack a session at the network level, easily monitoring, modifying, and injecting data into the session.

**Step**

1. Enter the following command:

   **`options ip.icmp_ignore_redirect.enable on`**

   Your storage system now ignores ICMP redirect messages.

   For more information about the `ip.icmp_ignore_redirect.enable` option, see the na_options(1) man page.

   > **Note:** By default, the `ip.icmp_ignore_redirect.enable` option is `off`.

# Network interface statistics

You can use the `ifstat` command to view statistics for the network interfaces supported by Data ONTAP. To determine the Ethernet controllers in your system, you can use the `sysconfig` command.

## Statistics for Gigabit Ethernet controller VI, VII, and G20 interfaces

The `ifstat` command output displays several statistics when you use the command for the Gigabit Ethernet controllers and G20 interfaces.

The statistics in this section are for the following controllers:

- 10/100/1000 Ethernet controller VI and VII
- Gigabit Ethernet controller VI
- 10/100/1000 Ethernet controller G20
- Gigabit Ethernet controller G20

### RECEIVE section statistics

The following table describes the statistics in the RECEIVE section of the `ifstat` command output.

| Statistic | Definition |
|---|---|
| `Frames/second` | Rate of received frames per second. |
| `Bytes/second` | Rate of received bytes per second. |
| `Errors/minute` | Rate of errors (which led to frames being lost) per minute. |
| `Discards/minute` | Rate per minute of packets discarded due to unavailable resources. |
| `Total frames` | Total frames that are received on the interface. |
| `Total bytes` | Total bytes that are received on the interface. |
| `Total errors` | Total errors that occur on the interface. |
| `Total discards` | Total number of packets that were discarded even though no errors were detected. This number is a sum of the "No buffers", "Bus overruns", and "Queue overflows" statistics. |
| `Multi/broadcast` | Total number of multicast or broadcast packets received. |

| Statistic | Definition |
|---|---|
| Alignment errors | Number of frames that are both misaligned and contain CRC errors. |
| Non-primary u/c | Number of Ethernet frames received for the partner's MAC address after a failover in an HA pair. |
| Tag drop | Number of tagged frames dropped on an interface that is not configured to support VLAN tagging. |
| Vlan tag drop | Number of tagged frames dropped that do not match the VLAN tags configured on the interface. |
| Vlan untag drop | Number of untagged frames dropped on an interface that is configured to be part of a VLAN. |
| CRC errors | Number of packets received with bad CRC. |
| Bad length | Total number of received packets with a bad length. These are frames counted as undersize, fragment, oversize, or jabber. |
| Runt frames | Number of received frames that were less than the minimum size (64 bytes) and had a valid CRC. |
| Fragment | Number of received frames that were less than the minimum size and had a bad CRC. |
| Long frames | Number of received frames that were greater than the maximum size and had a valid CRC. |
| Jabber | Number of received frames that were greater than the maximum size and had a bad CRC. |
| Bus overruns | Number of times the adapter's receive FIFO overflowed and a packet was dropped. This occurs when the bus is very busy and the adapter cannot transfer data into host memory. This might also occur when your storage system CPU is very busy and cannot process the received packets fast enough. |
| Queue overflows | Number of frames dropped on receive due to the driver receive queue overflowing. |
| No buffer | Number of times the driver could not allocate a buffer and a packet was dropped. This might happen when your storage system is very busy. If the count increases continually, it might indicate that a software component is not returning buffers. |
| Xon | Number of XON frames received when receive or full flow control is enabled. |

| Statistic | Definition |
|---|---|
| Xoff | Number of XOFF frames received when receive or full flow control is enabled. |
| Jumbo | Number of good packets received that were larger than the standard Ethernet packet size when jumbo frames are enabled. |
| Reset | Number of times the driver reset the NIC because the NIC was in a bad state. |
| Reset1 | Number of times the driver reset the NIC because the NIC was in a bad state. |
| Reset2 | Number of times the driver reset the NIC because the NIC was in a bad state. |

## TRANSMIT section statistics

The following table describes the statistics in the TRANSMIT section of the `ifstat` command output.

| Statistic | Meaning |
|---|---|
| Frames/second | Rate of transmitted frames per second. |
| Bytes/second | Rate of transmitted bytes per second. |
| Errors/minute | Rate of errors (which led to frames being lost) per minute. |
| Discards/minute | Rate per minute of packets discarded due to unavailable resources. |
| Total frames | Total frames that are transmitted on the interface. |
| Total bytes | Total bytes that are transmitted on the interface. |
| Total errors | Total errors that occur on the interface. |
| Total discards | Total number of packets that were discarded even though no errors were detected. This number is a sum of the "No buffers" and "Queue overflows" statistics. |
| Multi/broadcast | Total number of multicast or broadcast packets transmitted. |
| No buffers | Number of times the driver failed to allocate a buffer for the transmit packet. |
| Queue overflow | Number of outgoing packets dropped because the driver's queue was full. It might indicate a system problem. |

| Statistic | Meaning |
|---|---|
| Max collisions | Number of frames that were not transmitted because they encountered the maximum number of allowed collisions. Only valid in half-duplex mode. |
| Single collision | Number of frames that encountered exactly one collision. Only valid in half-duplex mode. |
| Multi collisions | Number of frames that encountered more than one collision, but less than the maximum allowed. Only valid in half-duplex mode. |
| Late collisions | Number of collisions that occurred outside the collision window. Only valid in half-duplex mode. |
| Xon | Number of XON frames transmitted when send or full flow control is enabled. |
| Xoff | Number of XOFF frames transmitted when send or full flow control is enabled. |
| Timeout | Number of times the adapter's transmitter hung and the adapter had to be reset. This can happen when the cable is pulled and the transmitter cannot transmit a packet. The adapter is reset to reclaim packet buffers. |
| Jumbo | Number of packets transmitted that were larger than the standard Ethernet frame size (1,518 bytes). |

## LINK INFO section statistics

The following table describes the statistics in the LINK INFO section of the ifstat command output.

| Statistic | Meaning |
|---|---|
| Current state | Current state of the interface: <br><br> • up or down—The state of the link. <br> • cfg_down—The interface is configured to the down status. <br> • enabling—The interface is moving to the up status. |
| Up to downs | Number of times the link switched between the up status and the down status. |

| Statistic | Meaning |
|---|---|
| Auto | Operational state of autonegotiation:<br><br>• on—Autonegotiation is enabled and succeeded.<br>• off—Autonegotiation failed. This happens when the device to which the interface is connected has disabled autonegotiation or is incompatible with the interface. This might also indicate that the interface is in the down status. |
| Speed | Speed of link negotiated or set. |
| Duplex | Duplex of the link negotiated or set. |
| Flow control | The operational (autonegotiated) flow control setting. |

**Related tasks**

# Statistics for Gigabit and 10 Gigabit Ethernet controllers T204, T210, and T320 interfaces

The ifstat command output displays several statistics when you use the command for the 10/100/1000 Ethernet controllers T204V and T204E, and the 10 Gigabit Ethernet controllers T210 and T320.

**RECEIVE section statistics**

The following table describes the statistics in the RECEIVE section of the ifstat command output when you use the command on the 10/100/1000 Ethernet controllers T204V and T204E, and the 10 Gigabit Ethernet controllers T210 and T320.

| Statistic | Meaning |
|---|---|
| Frames/second | Rate of received frames per second. |
| Bytes/second | Rate of received bytes per second. |
| Errors/minute | Rate of errors (which led to frames being lost) per minute. |
| Discards/minute | Rate per minute of packets discarded due to unavailable resources. |
| Total frames | Total frames that are received on the interface. |
| Total bytes | Total bytes that are received on the interface. |

| Statistic | Meaning |
|---|---|
| Total errors | Total errors that occur on the interface. |
| Total discards | Total number of packets that were discarded even though no errors were detected. This number is a sum of the "No buffers", "Bus overruns", and "Queue overflows" statistics. |
| Multi/broadcast | Total number of multicast or broadcast packets received. |
| Alignment errors | Number of frames that are both misaligned and contain CRC errors. |
| Non-primary u/c | Number of Ethernet frames received for the partner's MAC address after a failover in an HA pair. |
| Tag drop | Number of tagged frames dropped on an interface that is not configured to support VLAN tagging. |
| Vlan tag drop | Number of tagged frames dropped that do not match the VLAN tags configured on the interface. |
| Vlan untag drop | Number of untagged frames dropped on an interface that is configured to be part of a VLAN. |
| CRC errors | Number of packets received with bad CRC. |
| Runt frames | Number of received frames that were less than the minimum size (64 bytes) and had a valid CRC. |
| Long frames | Number of received frames that were greater than the maximum size and had a valid CRC. |
| Jabber | Number of received frames that were greater than the maximum size and had a bad CRC. |
| No buffer | Number of times the driver could not allocate a buffer and a packet was dropped. This might happen when your storage system is very busy. If the count increases continually, it might indicate that a software component is not returning buffers. |
| Jumbo | Number of good packets received that were larger than the standard Ethernet packet size when jumbo frames are enabled. |
| Pause Frames | Number of MAC Control PAUSE frames sent to the link partner to inhibit transmission of data frames for a specified period of time. This can help the partner from overrunning the controller's receive buffers. |

## TRANSMIT section statistics

The following table describes the statistics in the TRANSMIT section of the `ifstat` command output when you use the command on the 10/100/1000 Ethernet controllers T204V and T204E, and the 10 Gigabit Ethernet controllers T210 and T320.

| Statistic | Meaning |
| --- | --- |
| Frames/second | Rate of transmitted frames per second. |
| Bytes/second | Rate of transmitted bytes per second. |
| Errors/minute | Rate of errors (which led to frames being lost) per minute. |
| Discards/minute | Rate per minute of packets discarded due to unavailable resources. |
| Total frames | Total frames that are transmitted on the interface. |
| Total bytes | Total bytes that are transmitted on the interface. |
| Total errors | Total errors that occur on the interface. |
| Total discards | Total number of packets that were discarded even though no errors were detected. This number is a sum of the "No buffers" and "Queue overflows" statistics. |
| Multi/broadcast | Total number of multicast or broadcast packets transmitted. |
| No buffers | Number of times the driver failed to allocate a buffer for the transmit packet. |
| Queue overflow | Number of outgoing packets dropped because the driver's queue was full. It might indicate a system problem. |
| Bus Underruns | FIFO goes empty before an internal End-Of-Packet indicator is read. |
| Pause Frames | Number of MAC Control PAUSE frames sent to the link partner to inhibit transmission of data frames for a specified period of time. This can help the partner from overrunning the controller's receive buffers. |

## LINK INFO section statistics

The following table describes the statistics in the LINK INFO section of the `ifstat` command output when you use the command on the 10/100/1000 Ethernet controllers T204V and T204E, and the 10 Gigabit Ethernet controllers T210 and T320.

| Statistic | Meaning |
|---|---|
| Current state | Current state of the interface:<br><br>• up or down—The state of the link.<br>• cfg_down—The interface is configured to the down status.<br>• enabling—The interface is coming to the up status. |
| Up to downs | Number of times the link switched between the up status and the down status. |
| Speed | Speed of link negotiated or set. |
| Duplex | Duplex of the link negotiated or set. |
| Flow control | The operational (autonegotiated) flow control setting for T204 Gigabit Ethernet controller.<br><br>The configured flow control setting for T210 and T320 10 Gigabit Ethernet controller.<br><br>**Note:** Flow control cannot be autonegotiated on 10 Gigabit Ethernet controllers. |

# Statistics for the BGE 10/100/1000 Ethernet interface

The ifstat command output displays several statistics when you use the command on the BGE 10/100/1000 Ethernet interface.

### RECEIVE section statistics

The following table describes the statistics in the RECEIVE section of the ifstat command output when you use the command on the BGE 10/100/1000 Ethernet interface.

| Statistic | Meaning |
|---|---|
| Frames/second | Rate of received frames per second. |
| Bytes/second | Rate of received bytes per second. |
| Errors/minute | Rate of errors (which led to frames being lost) per minute. |
| Discards/minute | Rate per minute of packets discarded due to unavailable resources. |
| Total frames | Total frames that are received on the interface. |
| Total bytes | Total bytes that are received on the interface. |

| Statistic | Meaning |
|---|---|
| `Total errors` | Total errors that occur on the interface. |
| `Total discards` | Total number of packets that were discarded even though no errors were detected. This number is a sum of the "No buffers", "Bus overruns", and "Queue overflows" statistics. |
| `Multi/broadcast` | Total number of multicast or broadcast packets received. |
| `Alignment errors` | Number of frames that are both misaligned and contain CRC errors. |
| `Non-primary u/c` | Number of Ethernet frames received for the partner's MAC address after a failover in an HA pair. |
| `Tag drop` | Number of tagged frames dropped on an interface that is not configured to support VLAN tagging. |
| `Vlan tag drop` | Number of tagged frames dropped that do not match the VLAN tags configured on the interface. |
| `Vlan untag drop` | Number of untagged frames dropped on an interface that is configured to be part of a VLAN. |
| `CRC errors` | Number of packets received with bad CRC. |
| `Runt frames` | Number of received frames that were less than the minimum size (64 bytes) and had a valid CRC. |
| `Fragment` | Number of received frames that were less than the minimum size and had a bad CRC. |
| `Long frames` | Number of received frames that were greater than the maximum size and had a valid CRC. |
| `Jabber` | Number of received frames that were greater than the maximum size and had a bad CRC. |
| `No buffer` | Number of times the driver could not allocate a buffer and a packet was dropped. This might happen when your storage system is very busy. If the count increases continually, it might indicate that a software component is not returning buffers. |
| `Xon` | Number of XON frames received when receive or full flow control is enabled. |
| `Xoff` | Number of XOFF frames received when receive or full flow control is enabled. |

| Statistic | Meaning |
|-----------|---------|
| Jumbo | Number of good packets received that were larger than the standard Ethernet packet size when jumbo frames are enabled. |
| Ring full | Not used. Ignore. |
| Jumbo error | Error detected while processing a jumbo packet. Packet is discarded. |

### TRANSMIT section statistics

The following table describes the statistics in the TRANSMIT section of the ifstat command output when you use the command on the BGE 10/100/1000 Ethernet interface.

| Statistic | Meaning |
|-----------|---------|
| Frames/second | Rate of transmitted frames per second. |
| Bytes/second | Rate of transmitted bytes per second. |
| Errors/minute | Rate of errors (which led to frames being lost) per minute. |
| Discards/minute | Rate per minute of packets discarded due to unavailable resources. |
| Total frames | Total frames that are transmitted on the interface. |
| Total bytes | Total bytes that are transmitted on the interface. |
| Total errors | Total errors that occur on the interface. |
| Total discards | Total number of packets that were discarded even though no errors were detected. This number is a sum of the "No buffers" and "Queue overflows" statistics. |
| Multi/broadcast | Total number of multicast or broadcast packets transmitted. |
| No buffers | Number of times the driver failed to allocate a buffer for the transmit packet. |
| Queue overflow | Number of outgoing packets dropped because the driver's queue was full. It might indicate a system problem. |
| Max collisions | Number of frames that were not transmitted because they encountered the maximum number of allowed collisions. Only valid in half-duplex mode. |
| Single collision | Number of frames that encountered exactly one collision. Only valid in half-duplex mode. |

| Statistic | Meaning |
|---|---|
| Multi collisions | Number of frames that encountered more than one collision, but less than the maximum allowed. Only valid in half-duplex mode. |
| Late collisions | Number of collisions that occurred outside the collision window. Only valid in half-duplex mode. |
| Xon | Number of XON frames transmitted when send or full flow control is enabled. |
| Xoff | Number of XOFF frames transmitted when send or full flow control is enabled. |
| Jumbo | Number of packets transmitted that were larger than the standard Ethernet packet size when jumbo frames are enabled. |
| Deferred | Number of frames for which the first transmission was delayed because the medium was busy. |
| MAC Internal | Number of frames not transmitted due to an internal MAC sublayer error. |

### LINK INFO section statistics

The following table describes the statistics in the LINK INFO section of the ifstat command output when you use the command on the BGE 10/100/1000 Ethernet interface.

| Statistic | Meaning |
|---|---|
| Current state | Current state of the interface:<br><br>• up or down—The state of the link.<br>• cfg_down—The interface is configured to the down status.<br>• enabling—The interface is coming to the up status. |
| Up to downs | Number of times the link switched between the up status and the down status. |
| Speed | Speed of link negotiated or set. |
| Duplex | Duplex of the link negotiated or set. |
| Flow control | The operational (autonegotiated) flow control setting. |

**Related tasks**

# Statistics for 10 Gigabit Ethernet Controller IX1 - SFP+

The `ifstat` command output displays several statistics when you use the command for the 10 Gigabit Ethernet controllers.

### RECEIVE section statistics

The following table describes the statistics in the RECEIVE section of the `ifstat` command output.

| Statistic | Definition |
|---|---|
| Frames/second | Rate of received frames per second. |
| Bytes/second | Rate of received bytes per second. |
| Errors/minute | Rate of errors (which led to frames being lost) per minute. |
| Discards/minute | Rate per minute of packets discarded due to unavailable resources. |
| Total frames | Total frames that are received on the interface. |
| Total bytes | Total bytes that are received on the interface. |
| Total errors | Total errors that occur on the interface. |
| Total discards | Total number of packets that are discarded even though no errors were detected. This number is a sum of the "No buffers", "Bus overruns", and "Queue overflows" statistics. |
| Multi/broadcast | Total number of multicast or broadcast packets received. |
| No buffers | Number of times the driver could not allocate a buffer and a packet was dropped. This might occur when your storage system is very busy. If the count increases continually, it might indicate that a software component is not returning buffers. |
| Non-primary u/c | Number of Ethernet frames received for the partner's MAC address after a failover in an HA pair. |
| Tag drop | Number of tagged frames dropped on an interface that is not configured to support VLAN tagging. |
| Vlan tag drop | Number of tagged frames dropped that do not match the VLAN tags configured on the interface. |
| Vlan untag drop | Number of untagged frames dropped on an interface that is configured to be part of a VLAN. |
| CRC errors | Number of packets received with bad CRC. |

| Statistic | Definition |
|---|---|
| Runt frames | Number of received frames that were less than the minimum size (64 bytes) and had a valid CRC. |
| Fragment | Number of received frames that were less than the minimum size and had a bad CRC. |
| Long frames | Number of received frames that were greater than the maximum size and had a valid CRC. |
| Jabber | Number of received frames that were greater than the maximum size and had a bad CRC. |
| Bus overruns | Number of times the adapter's receive FIFO overflowed and a packet was dropped. This occurs when the bus is very busy and the adapter cannot transfer data into host memory. This might also occur when your storage system CPU is very busy and cannot process the received packets quickly. |
| Xon | Number of XON frames received when receive or full flow control is enabled. |
| Xoff | Number of XOFF frames received when receive or full flow control is enabled. |
| Jumbo | Number of good packets received that are larger than the standard Ethernet packet size when jumbo frames are enabled. |

### TRANSMIT section statistics

The following table describes the statistics in the TRANSMIT section of the `ifstat` command output.

| Statistic | Meaning |
|---|---|
| Frames/second | Rate of transmitted frames per second. |
| Bytes/second | Rate of transmitted bytes per second. |
| Errors/minute | Rate of errors (which led to frames being lost) per minute. |
| Discards/minute | Rate per minute of packets discarded due to unavailable resources. |
| Total frames | Total frames that are transmitted on the interface. |
| Total bytes | Total bytes that are transmitted on the interface. |
| Total errors | Total errors that occur on the interface. |

| Statistic | Meaning |
|---|---|
| Total discards | Total number of discarded packets even though no errors are detected. This number is a sum of the "No buffers" and "Queue overflows" statistics. |
| Multi/broadcast | Total number of multicast or broadcast packets transmitted. |
| No buffers | Number of times the driver failed to allocate a buffer for the transmit packet. |
| Queue overflows | Number of outgoing packets dropped because the driver's queue was full. It might indicate a system problem. |
| Xon | Number of XON frames transmitted when send or full flow control is enabled. |
| Xoff | Number of XOFF frames transmitted when send or full flow control is enabled. |
| Timeout | Number of times the adapter's transmitter was hung and the adapter had to be reset. This might occur when the cable is pulled and the transmitter cannot transmit a packet. The adapter is reset to reclaim packet buffers. |
| Jumbo | Number of packets transmitted that are larger than the standard Ethernet frame size (1,518 bytes). |

## LINK INFO section statistics

The following table describes the statistics in the LINK INFO section of the `ifstat` command output.

| Statistic | Meaning |
|---|---|
| Current state | Current state of the interface:<br><br>• `up` or `down`—The state of the link.<br>• `cfg_down`—The interface is configured to the `down` status.<br>• `enabling`—The interface is moving to the `up` status. |
| Up to downs | Number of times the link is switched between the `up` status and the `down` status. |
| Speed | Speed of link negotiated or set. |
| Duplex | Duplex of the link negotiated or set. |
| Flow control | The configured flow control setting, as flow control is not autonegotiated on this interface. |

# Statistics for Dual 10G Ethernet Controller CNA - SFP+

The `ifstat` command output displays several statistics when you use the command for the Dual 10G Ethernet controllers.

### RECEIVE section statistics

The following table describes the statistics in the RECEIVE section of the `ifstat` command output.

| Statistic | Definition |
|---|---|
| `Frames/second` | Rate of received frames per second. |
| `Bytes/second` | Rate of received bytes per second. |
| `Errors/minute` | Rate of errors (which led to frames being lost) per minute. |
| `Discards/minute` | Rate of packets discarded per minute due to unavailable resources. |
| `Total frames` | Total frames received by the interface. |
| `Total bytes` | Total bytes received by the interface. |
| `Total errors` | Total errors that occur on the interface. |
| `Total discards` | Total number of packets that were discarded inspite of undetected errors. This number is a sum of the "No buffers", "Bus overruns", and "Queue overflows" statistics. |
| `Multi/broadcast` | Total number of multicast or broadcast packets received. |
| `No buffers` | Number of times the driver could not allocate a buffer and a packet was dropped. This might occur when your storage system is busy. If the count increases continually, it might indicate that a software component is not returning buffers. |
| `Non-primary u/c` | Number of Ethernet frames received for the partner's MAC address after a failover in an HA pair. |
| `Tag drop` | Number of tagged frames dropped on an interface that is not configured to support VLAN tagging. |
| `Vlan tag drop` | Number of tagged frames dropped that do not match the VLAN tags configured on the interface. |
| `Vlan untag drop` | Number of untagged frames dropped on an interface that is configured to be part of a VLAN. |
| `CRC errors` | Number of packets received with bad CRC. |

| Statistic | Definition |
|---|---|
| Alignment errors | Number of frames that are both misaligned and contain CRC errors. |
| Runt frames | Number of received frames that were less than the minimum size (64 bytes) and had a valid CRC. |
| Long frames | Number of received frames that were greater than the maximum size and had a valid CRC. |
| Code error | Number of received packets/ frames with code errors. |
| Preamble error | Number of received packets/ frames with preamble errors. |
| Frame length error | Number of received frames with frame length errors. |
| Received VLAN frames | Number of received VLAN frames. |

### TRANSMIT section statistics

The following table describes the statistics in the TRANSMIT section of the ifstat command output.

| Statistic | Meaning |
|---|---|
| Frames/second | Rate of transmitted frames per second. |
| Bytes/second | Rate of transmitted bytes per second. |
| Errors/minute | Rate of errors (which led to frames being lost) per minute. |
| Discards/minute | Rate of packets discarded per minute due to unavailable resources. |
| Total frames | Total frames that were transmitted on the interface. |
| Total bytes | Total bytes that were transmitted on the interface. |
| Total errors | Total errors that occur on the interface. |
| Total discards | Total number of discarded packets even though no errors were detected. This number is a sum of the "No buffers" and "Queue overflows" statistics. |
| Multi/broadcast | Total number of multicast or broadcast packets transmitted. |
| No buffers | Number of times the driver failed to allocate a buffer for the transmit packet. |
| Queue overflows | Number of outgoing packets dropped because the driver's queue was full. It might indicate a system problem. |

| Statistic | Meaning |
|---|---|
| `Backplane error` | Number of PCI backplane errors detected. |
| `Padded sent` | Number of 16-60 byte frames, padded and sent. |
| `Too long frames` | Number of frames that were too long but were sent anyway. |
| `Too short frames` | Number of frames that were too short and illegal and were not sent. |
| `Unmatched length` | Total descriptor length that does not match transfer length. |

### LINK INFO section statistics

The following table describes the statistics in the LINK INFO section of the `ifstat` command output.

| Statistic | Meaning |
|---|---|
| `Current state` | Current state of the interface:<br><br>• `up` or `down`—The state of the link.<br>• `cfg_down`—The interface is configured to the `down` status.<br>• `enabling`—The interface is moving to the `up` status. |
| `Up to downs` | Number of times the link is switched between the `up` status and the `down` status. |
| `Speed` | Speed of link negotiated or set. |
| `Duplex` | Duplex of the link negotiated or set. |
| `Flow control` | The configured flow control setting, as flow control is not autonegotiated on this interface. |

# Statistics for Quad Gigabit Ethernet Controller 82850

The `ifstat` command output displays several statistics when you use the command for the Quad Gigabit Ethernet Controllers.

### RECEIVE section statistics

The following table describes the statistics in the RECEIVE section of the `ifstat` command output.

| Statistic | Definition |
|---|---|
| `Frames/second` | Rate of received frames per second. |

| Statistic | Definition |
|---|---|
| Bytes/second | Rate of received bytes per second. |
| Errors/minute | Rate of errors (which led to frames being lost) per minute. |
| Discards/minute | Rate of packets discarded per minute due to unavailable resources. |
| Total frames | Total frames received by the interface. |
| Total bytes | Total bytes received by the interface. |
| Total errors | Total errors that occur on the interface. |
| Total discards | Total number of packets that were discarded inspite of undetected errors. This number is a sum of the "No buffers", "Bus overruns", and "Queue overflows" statistics. |
| Multi/broadcast | Total number of multicast or broadcast packets received. |
| No buffers | Number of times the driver could not allocate a buffer and a packet was dropped.<br><br>This might occur when your storage system is busy. If the count increases continually, it is a indication that a software component is not returning buffers. |
| Non-primary u/c | Number of Ethernet frames received for the partner's MAC address after a failover in an HA pair. |
| Tag drop | Number of tagged frames dropped on an interface that is not configured to support VLAN tagging. |
| Vlan tag drop | Number of tagged frames dropped that do not match the VLAN tags configured on the interface. |
| Vlan untag drop | Number of untagged frames dropped on an interface that is configured to be part of a VLAN. |
| CRC errors | Number of packets received with bad CRC. |
| Runt frames | Number of received frames that were less than the minimum size (64 bytes) and had a valid CRC. |
| Fragment | Number of received frames that were less than the minimum size and had a bad CRC. |
| Long frames | Number of received frames that were greater than the maximum size and had a valid CRC. |
| Jabber | Number of received frames that were greater than the maximum size and had a bad CRC. |

| Statistic | Definition |
|-----------|------------|
| `Alignment errors` | Number of frames that are misaligned and contain CRC errors. |
| `Bus overruns` | Number of times the adapter's receive FIFO overflowed and a packet was dropped. This occurs when the bus is too busy and the adapter cannot transfer data into host's memory. This might also occur when your storage system CPU is too busy and cannot process the received packets fast enough. |
| `Xon` | Number of XON frames received when receive or full flow control is enabled. |
| `Xoff` | Number of XOFF frames received when receive or full flow control is enabled. |
| `Jumbo` | Number of good packets received that were larger than the standard Ethernet packet size when jumbo frames are enabled. |

### TRANSMIT section statistics

The following table describes the statistics in the TRANSMIT section of the `ifstat` command output.

| Statistic | Meaning |
|-----------|---------|
| `Frames/second` | Rate of transmitted frames per second. |
| `Bytes/second` | Rate of transmitted bytes per second. |
| `Errors/minute` | Rate of errors (which led to frames being lost) per minute. |
| `Discards/minute` | Rate of packets discarded per minute due to unavailable resources. |
| `Total frames` | Total frames that were transmitted on the interface. |
| `Total bytes` | Total bytes that were transmitted on the interface. |
| `Total errors` | Total errors that occur on the interface. |
| `Total discards` | Total number of discarded packets even though no errors were detected. This number is a sum of the "No buffers" and "Queue overflows" statistics. |
| `Multi/broadcast` | Total number of multicast or broadcast packets transmitted. |
| `Queue overflows` | Number of outgoing packets dropped because the driver's queue was full and might indicate a system problem. |
| `No buffers` | Number of times the driver failed to allocate a buffer for the transmit packet. |

| Statistic | Meaning |
|-----------|---------|
| Max collisions | Number of frames that encountered more than one collision, but less than the maximum allowed. Only valid in half-duplex mode. |
| Single collision | Number of frames that encountered exactly one collision. Only valid in half-duplex mode. |
| Multi collisions | Number of frames that encountered more than one collision, but less than the maximum allowed. Only valid in half-duplex mode. |
| Late collisions | Number of collisions that occurred outside the collision window. Only valid in half-duplex mode. |
| Xon | Number of XON frames transmitted when send or full flow control is enabled. |
| Xoff | Number of XOFF frames transmitted when send or full flow control is enabled. |
| Jumbo | Number of packets transmitted that were larger than the standard Ethernet frame size (1,518 bytes). |

### LINK INFO section statistics

The following table describes the statistics in the LINK INFO section of the ifstat command output.

| Statistic | Meaning |
|-----------|---------|
| Current state | Current state of the interface:<br><br>• up or down—The state of the link.<br>• cfg_down—The interface is configured to the down status.<br>• enabling—The interface is moving to the up status. |
| Up to downs | Number of times the link is switched between the up status and the down status. |
| Speed | Speed of link negotiated or set. |
| Duplex | Duplex of the link negotiated or set. |
| Flow control | Operational (autonegotiated) flow control setting. |

# Ways to improve your storage system's performance

You can improve your storage system's performance by performing certain configuration procedures, such as using interface groups, correcting duplex mismatches, and upgrading to Ethernet interfaces.

The following configuration procedures might improve the performance of your storage system:

- Using static or dynamic multimode interface groups to aggregate the bandwidth of multiple interfaces
- Using jumbo frames with your network interfaces to reduce CPU processing overhead
- Upgrading to a faster network interface

  You can increase the storage system's performance by upgrading to a faster network interface (10 Gigabit Ethernet interfaces).

- Correcting duplex mismatches on 10Base-T or 100Base-T Ethernet networks

  On 10Base-T or 100Base-T Ethernet networks, the speed and duplex settings for the interfaces at both ends of a link must match exactly. You can use the `ifconfig interface` command to check the duplex setting of your storage system's interface.

  If the setting is to autonegotiate, the `ifconfig` command displays a setting that begins with `auto` (for example, `auto-100tx-fd-up`). Otherwise, the `ifconfig` command displays the configured media type setting (for example, `100tx-fd-up`).

  > **Note:** If one end of the link is set to autonegotiate, the other end must also be set to autonegotiate; otherwise, a mismatch might occur. You can determine the negotiated setting with the `ifstat` command.

- Using iSCSI multiconnection sessions to balance the load across interfaces

  For each iSCSI session, multiple connections are created. The number of allowed connections is negotiated during login and session creation. While it is possible to create multiple connections over a single physical interface, it is best to use multiple physical interfaces for bandwidth enhancement.

- Enabling fast path on your storage system

  Fast path provides load balancing by sending responses on the same network interface that receives the incoming requests and improved performance by skipping routing table lookups.

- Blocking data traffic on the dedicated management interface

  In storage systems with dedicated management interface, e0M, you should block certain types of data traffic on e0M, such as SnapMirror transfers and SnapVault transfers, and other data transfers that use the file access protocols such as CIFS, NFS, NDMP. If the low-bandwidth management interface is configured to serve data traffic, system performance might be affected. Therefore, e0M should be used only for Data ONTAP management activities such as running a Telnet, RSH, or SSH session.

**Related concepts**

**Related tasks**

# IP port usage on a storage system

The Data ONTAP services file is available in the `/etc` directory. The `/etc/services` file is in the same format as its corresponding UNIX system's `/etc/services` file.

## Host identification

Although some port scanners are able to identify storage systems as storage systems, other port scanners report storage systems as UNIX systems if they support NFS or as Windows systems if they support CIFS. There are several services that are not currently listed in the `/etc/services` file.

The following table gives a sample content of the `/etc/services` file.

| Service | Port/ Protocol | Description |
|---------|----------------|-------------|
| ftp-data | 20/tcp | # File transfer protocol |
| ftp | 21/tcp | # File transfer protocol |
| ssh | 22/tcp | # SecureAdmin rsh replacement |
| telnet | 23/tcp | # Remote login (insecure) |
| smtp | 25/tcp | # outbound connections for autosupport |
| time | 37/tcp | # Time Service |
| time | 37/udp | # Time Service |
| domain | 53/udp | # DNS - outbound only |
| domain | 53/tcp | # DNS zone transfers - unused |
| dhcps | 67/udp | # DHCP server - outbound only |
| dhcp | 68/udp | # DHCP client - only first-time setup |
| tftp | 69/udp | # Trivial FTP - for netboot support |

| Service | Port/ Protocol | Description |
|---|---|---|
| http | 80/tcp | # HTTP license, FilerView, SecureAdmin |
| kerberos | 88/udp | # Kerberos 5 - outbound only |
| kerberos | 88/tcp | # Kerberos 5 - outbound only |
| portmap | 111/udp | # aka rpcbind, used for NFS |
| portmap | 111/tcp | # aka rpcbind, used for NFS |
| nntp | 119/tcp | # unused, shouldn't be listed here. |
| ntp | 123/tcp | # Network Time Protocol |
| ntp | 123/udp | # Network Time Protocol |
| netbios-name | 137/udp | # NetBIOS nameserver - for CIFS |
| netbios-dg | 138/udp | # NetBIOS datagram service - for CIFS |
| netbios-ssn | 139/tcp | # NetBIOS service session - for CIFS |
| ssl | 443/tcp | # Secure FilerView (SecureAdmin) |
| cifs-tcp | 445/tcp | # CIFS over TCP with NetBIOS framing |
| snmp | 161/udp | # For Data Fabric Manager or other such tools |
| ldap | 389/tcp | # LDAP session |
| https | 443/tcp | # SecureAdmin/SSL |
| shell | 514/tcp | # rsh, insecure remote command execution. |
| syslog | 514/udp | # outbound only |

| Service | Port/ Protocol | Description |
|---------|----------------|-------------|
| route | 520/udp | # for RIP routing protocol |
| ldap-ssl | 636/tcp | # LDAP over SSL |
| kerberos-sec | 750/udp | # outbound only, if at all |
| kerberos-sec | 750/tcp | # outbound only, if at all |
| nfsd | 2049/udp | # primary NFS service |
| nfsd | 2049/tcp | # primary NFS service |
| nrv | 2050/tcp | # Remote Volume protocol, used in FlexCache and Restore-On-Demand |
| iscsi-target | 3260/tcp | # iSCSI target port |
| nlockmgr | 4045/tcp | # NLM |
| nlockmgr | 4045/udp | # NLM |
| mountd | 4046/tcp | # NFS mountd protocol |
| mountd | 4046/udp | # NFS mountd protocol |
| status | 4047/tcp | # CTP state |
| status | 4047/udp | # CTP state |
| pcnfsd | 4048/tcp | # PCNFS protocol |
| pcnfsd | 4048/udp | # PCNFS protocol |
| rquotad | 4049/udp | # NFS rquotad protocol |
| ndmp | 10000/tcp | # for network backups |
| sm-ics | 10565/tcp | # Snapmirror Multipath |
| snapmirror | 10566/tcp | # SnapMirror and SnapVault |
| sm-sync-block | 10567/tcp | # Snapmirror Sync Block Data |
| sm-sync-trans | 10568/tcp | # Snapmirror Sync Transaction Data |

| Service | Port/ Protocol | Description |
|---------|----------------|-------------|
| `sm-sync-ctrl` | 10569/tcp | # Snapmirror Sync Control Data |

# /etc/services NNTP and TTCP ports

The NNTP and TTCP ports are not used by your storage system and should never be detected by a port scanner.

# NFS-enabled ports

Some ports (port numbers in the 600 range) on the storage system are NFS-enabled.

The following ports are found on the storage system with NFS-enabled:

| UDP | 602 | NFS mount daemon (`mountd`) |
|-----|-----|------------------------------|
| TCP | 603 | NFS mount daemon (`mountd`) |
| UDP | 604 | NFS status daemon (`statd`, `statmon`) |
| TCP | 605 | NFS status daemon (`statd`, `statmon`) |
| UDP | 606 | NFS lock manager (`lockd`, `nlockmgr`) |
| TCP | 607 | NFS lock manager (`lockd`, `nlockmgr`) |
| UDP | 608 | NFS quota daemon (`quotad`, `rquotad`) |

On other systems, the ports appear as follows:

| UDP | 611 | NFS mount daemon (`mountd`) |
|-----|-----|------------------------------|
| TCP | 612 | NFS mount daemon (`mountd`) |
| UDP | 613 | NFS status daemon (`statd`, `statmon`) |
| TCP | 614 | NFS status daemon (`statd`, `statmon`) |
| UDP | 615 | NFS lock manager (`lockd`, `nlockmgr`) |
| TCP | 616 | NFS lock manager (`lockd`, `nlockmgr`) |
| UDP | 617 | NFS quota daemon (`quotad`, `rquotad`) |

The following command on UNIX systems obtains the correct information by querying the port mapper on port 111:

```
toaster# rpcinfo -p sys1

   program vers proto    port  service
    100011    1   udp     608  rquotad
    100021    4   tcp     607  nlockmgr
    100021    3   tcp     607  nlockmgr
    100021    1   tcp     607  nlockmgr
    100021    4   udp     606  nlockmgr
    100021    3   udp     606  nlockmgr
    100021    1   udp     606  nlockmgr
    100024    1   tcp     605  status
    100024    1   udp     604  status
    100005    3   tcp     603  mountd
    100005    2   tcp     603  mountd
    100005    1   tcp     603  mountd
    100005    3   udp     602  mountd
    100005    2   udp     602  mountd
    100005    1   udp     602  mountd
    100003    3   udp    2049  nfs
    100003    2   udp    2049  nfs
    100000    2   tcp     111  rpcbind
    100000    2   udp     111  rpcbind
```

**Note:** The port numbers listed for mountd, statd, lockd, and quotad are not committed port numbers. These services can be running on other ports of the storage systems. Because the system selects these port numbers at random when it boots, they are not listed in the /etc/services file.

## Ports not listed in /etc/services

Some ports appear in a port scan but are not listed in the /etc/services file, for example, TCP ports 22 and 443.

The following ports appear in a port scan but are not listed in the /etc/services file.

| Protocol | Port | Service |
|----------|------|---------|
| TCP | 22 | SSH (SecureAdmin) |
| TCP | 443 | SSL (SecureAdmin) |
| TCP | 3260 | iSCSI-Target |

**Note:** Disable the open ports that you do not need.

## FTP

File Transfer Protocol (FTP) uses TCP ports 20 and 21.

If you use FTP to transfer files to and from your storage system, the FTP port is required; otherwise, use FilerView or the following CLI command to disable the FTP port:

**`options ftpd.enable off`**

FTP is not a secure protocol for two reasons:

* When users log in to the system, user names and passwords are transmitted over the network in clear text format that can easily be read by a packet sniffer program.
  These user names and passwords can then be used to access data and other network resources. You should establish and enforce policies that prevent the use of the same passwords to access storage systems and other network resources.
* FTP server software used on platforms other than storage systems contains serious security-related flaws that allow unauthorized users to gain administrative (root) access and control over the host.

For a detailed description of the FTP support for your storage system, see the *Data ONTAP 7-Mode File Access and Protocols Management Guide*.

# SSH

Secure Shell (SSH) protocol is a secure replacement for RSH and runs on TCP port 22. This port appears in a port scan only if the SecureAdmin software is installed on your storage system.

There are three commonly deployed versions of the SSH protocol:

* SSH version 1—is secure than RSH or Telnet, but is vulnerable to TCP session attacks.
  This vulnerability to attack lies in the SSH protocol version 1 itself and not in the associated storage system products.
* SSH version 2—has a number of feature improvements over SSH version 1 and is less vulnerable to attacks.
* SSH version 1.5—is used to identify clients or servers that support both SSH versions 1 and 2.

To disable SSH support or to close TCP port 22, you must use the following CLI command:

**`secureadmin disable ssh`**

# Telnet

Telnet is used for administrative control of your storage system and uses TCP connections on port 23. Telnet is more secure than RSH, as secure as FTP, and less secure than SSH or Secure Socket Layer (SSL).

Telnet is less secure than SSH and SSL because:

* When users log in to a system, such as your storage system, user names and passwords are transmitted over the network in clear text format.
  Clear text format can be read by an attacker by using a packet sniffer program. The attacker can use these user names and passwords to log in to your storage system and execute unauthorized

administrative functions, including destruction of data on the system. If administrators use the same passwords on your storage system as they do on other network devices, the attacker can use these passwords to access the resources of the storage system as well.

**Note:** To reduce the potential for attack, you must establish and enforce policies preventing administrators from using the same passwords on your storage system that they use to access other network resources.

• Telnet server software used on other platforms (typically in UNIX environments) have serious security-related flaws that allow unauthorized users to gain administrative (root) control over the host.

Telnet is also vulnerable to the same type of TCP session attacks as SSH protocol version 1. However, TCP session attacks are less common because a packet sniffing attack is easier.

To disable Telnet, you must set `options telnet.enable` to `off`.

# SMTP

Simple Mail Transport Protocol (SMTP) uses TCP port 25. Your storage system does not listen on this port but makes outgoing connections to mail servers using this protocol when sending AutoSupport e-mail.

# Time service

Your storage system supports the Network Time Protocol (NTP) for time synchronization.

To synchronize the system time with a network time server, you must ensure that the `timed.enable` option is set to `on`, the `timed.proto` option is set to `ntp`, and the `timed.servers` option is set with at least one valid time server.

**Note:** Both controllers in an HA pair should have time synchronization enabled to ensure time consistency in the event of controller failover.

For more information about synchronizing the system time, see the *Data ONTAP 7-Mode System Administration Guide*.

# DNS

The Domain Name System (DNS) uses UDP port 53 and TCP port 53. Your storage system does not typically listen on these ports because it does not run a domain name server. However, if DNS is enabled on your storage system, it makes outgoing connections using UDP port 53 for host name and IP address lookups.

The storage system never uses TCP port 53 because this port is used explicitly for communication between DNS servers. Outgoing DNS queries by your storage system are disabled by turning off

DNS support. Turning off DNS support protects against receiving bad information from another DNS server.

Because your storage system does not run a domain name server, the name service must be provided by one of the following:

• Network information service (NIS)
• An /etc/hosts file
• Replacement of host names in the configuration files (such as /etc/exports, /etc/usermap.cfg, and so on) with IP addresses

DNS must be enabled for participation in an Active Directory domain.

# DHCP

Clients broadcast messages to the entire network on UDP port 67 and receive responses from the Dynamic Host Configuration Protocol (DHCP) server on UDP port 68. The same ports are used for the BOOTP protocol.

DHCP is used only for the first-time setup of your storage system. Detection of DHCP activity on your storage system by a port scan other than the activity during the first-time setup indicates a serious configuration or software error.

# TFTP

Trivial File Transfer Protocol (TFTP) uses TCP port 69. It is used mostly for booting UNIX or UNIX-like systems that do not have a local disk (this process is also known as netbooting) and for storing and retrieving configuration files for devices such as Cisco routers and switches.

Transfers are not secure on TFTP because it does not require authentication for clients to connect and transfer files.

Your storage system's TFTP server is not enabled by default. When TFTP is enabled, the administrator must specify a directory to be used by TFTP clients, and these clients cannot access other directories. Even within the TFTP directory, access is read-only. TFTP should be enabled only if necessary. You can disable TFTP using the following option:

**options tftpd.enable off**

You can configure the maximum number of simultaneous connections handled by the TFTP server by using the tftpd.max_connections option. The default number of TFTP connections is 8. The maximum number of connections supported is 32.

# HTTP

Hypertext Transport Protocol (HTTP) runs on TCP port 80 and is the protocol used by Web browsers to access Web pages.

Your storage system uses HTTP to access the following:

- Files when HTTP is enabled
- FilerView for graphical user interface (GUI) administration
- Secure FilerView when SecureAdmin is installed

The SecureAdmin SSL interface accepts connections on TCP port 443. SecureAdmin manages the details of the SSL network protocol, encrypts the connection, and then passes this traffic through to the normal HTTP FilerView interface through a loopback connection. This loopback connection does not use a physical network interface. HTTP communication takes place inside your storage system, and no clear text packets are transmitted.

HTTP is not vulnerable to security attacks because it provides read-only access to documents by unauthenticated clients. Although authentication is not typically used for file access, it is frequently used for access to restricted documents or for administration purposes, such as FilerView administration. The authentication methods defined by HTTP send credentials, such as user names and passwords, over the network without encryption. The SecureAdmin product is provided with SSL support to overcome this shortcoming.

> **Note:** You can stop your storage system from listening for new connections by setting the options `httpd.enable` and `httpd.admin.enable` to `off`. If either of the options is set to `on`, your storage system will listen for new connections.

# Kerberos

There are four Kerberos ports in the `/etc/services` file: TCP port 88, UDP port 88, TCP port 750, and UDP port 750. These ports are used only for outbound connections from your storage system. Your storage system does not run Kerberos servers or services and does not listen on these ports.

Kerberos is used by your storage system to communicate with the Microsoft Active Directory servers for both CIFS authentication and, if configured, NFS authentication.

# NFS

Network File System (NFS) is used by UNIX clients for file access. NFS uses port 2049.

NFSv3 and NFSv2 use the portmapper service on TCP or UDP port 111. The portmapper service is consulted to get the port numbers for services used with NFSv3 or NFSv2 protocols such as mountd, statd, and nlm. NFSv4 does not require the portmapper service.

NFSv4 provides the delegation feature that enables your storage system to grant local file access to clients. To delegate, your storage system sets up a separate connection to the client and sends callbacks on it. To communicate with the client, your storage system uses one of the reserved ports (port numbers less than 1024). To initiate the connection, the client registers the callback program on a random port and informs the server about it.

With delegations enabled, NFSv4 is not firewall-friendly because several other ports need to be opened up as well.

You can disable the TCP and UDP ports by setting the `nfs.tcp.enable` and `nfs.udp.enable` options to `off`.

To disable NFS, you should use the `nfs off` command.

# CIFS

Common Internet File Service (CIFS) is the successor to the server message block (SMB) protocol. CIFS is the primary protocol used by Windows systems for file sharing.

CIFS uses UDP ports 137 and 138, and TCP ports 139 and 445. Your storage system sends and receives data on these ports while providing CIFS service. If it is a member of an Active Directory domain, your storage system must also make outbound connections destined for DNS and Kerberos.

CIFS is required for Windows file service. You can disable CIFS using FilerView or by issuing the `cifs terminate` command on your storage system console.

**Note:** If you disable CIFS, be aware that your storage system's `/etc/rc` file can be set up to automatically enable CIFS again after a reboot.

# SSL

The Secure Sockets Layer (SSL) protocol provides encryption and authentication of TCP connections. Data ONTAP supports SSLv2, SSLv3, and Transport Layer Security (TLS) version 1.0. You should use TLSv1.0 or SSLv3 because it offers better security than previous SSL versions.

When SecureAdmin is installed and configured on your storage system, it listens for SSL connections on TCP port 443. It receives secure Web browser connections on this port and uses unencrypted HTTP, running on TCP port 80, through a loopback connection to pass the traffic to FilerView. This loopback connection is contained within your storage system and no unencrypted data is transmitted over the network.

You can enable or disable SSL by using FilerView or with the following command:

```
secureadmin {enable|disable} ssl
```

For TLS to be used for communication, both the client requesting the connection and the storage system must support TLS.

TLS is disabled by default, and setting up SSL does not automatically enable TLS. Before enabling TLS, ensure that SSL has been set up and enabled. To enable or disable TLS, enter the following command:

```
options tls.enable {on|off}
```

# SNMP

Simple Network Management Protocol (SNMP) is an industry-standard protocol used for remote monitoring and management of network devices over UDP port 161.

SNMP is not secure because of the following reasons:

- Instead of using encryption keys or a user name and password pair, SNMP uses a community string for authentication. The community string is transmitted in clear text format over the network, making it easy to capture with a packet sniffer.
  Within the industry, devices are typically configured at the factory to use `public` as the default community string. The public password allows users to make queries and read values but does not allow users to invoke commands or change values. Some devices are configured at the factory to use `private` as the default community string, allowing users full read-write access.
- Even if you change the read and write community string on a device to something other than `private`, an attacker can easily learn the new string by using the read-only `public` community string and asking the router for the read-write string.

There are three versions of SNMP:

- SNMPv1 is the original protocol and is not commonly used.
- SNMPv2 is identical to SNMPv1 from a network protocol standpoint and is vulnerable to the same security problems. The only differences between the two versions are in the messages sent, messages received, and types of information. These differences are not important from a security perspective.
- SNMPv3 is the latest protocol version and includes security improvements but is difficult to implement and many vendors do not yet support it. SNMPv3 supports several different types of network encryption and authentication schemes. It allows for multiple users, each with different permissions, and solves SNMPv1 security problems while maintaining an important level of compatibility with SNMPv2.

SNMP is required if you want to monitor a storage system through an SNMP monitoring tool, such as DataFabric Manager. The SNMP implementation in the storage system allows read-only access. Regardless of the community string used, the user cannot issue commands or change variables using SNMP on your storage system.

You should use the `snmp.access` option to restrict SNMP access to a named set of trusted hosts.

You can disable SNMP entirely by setting the `snmp.enable` option to `off` to disable SNMP entirely.

The snmp community delete and snmp community add commands are used to change the community string to something other than the default value.

# RSH

Remote Shell (RSH) protocol is used for remote command execution. It is less secure than TFTP and uses TCP port 514.

RSH is not secure because passwords are not required for login and commands are easy to misconfigure. Therefore, you should disable RSH by setting the rsh.enable option to off.

You should use the SSH supplied with SecureAdmin for remote command execution and login. If this is not possible, Telnet is preferred to RSH.

If RSH is the only alternative, follow these guidelines when using RSH:

- Specify only secure, trusted hosts in the /etc/hosts.equiv file.
- Always use IP addresses rather than host names in the /etc/hosts.equiv file.
- Always specify a single IP address with a single user name on each line in /etc/hosts.equiv file.
- Use the rsh.access option instead of the trusted.hosts option for access control.
- Make sure the ip.match_any_ifaddr option is set to off.

# Syslog

Your storage system sends messages to hosts specified by the user in the /etc/syslog.conf file by using the syslog protocol on UDP port 514. It does not listen on this port, nor does it act as a syslog server.

# The routed daemon

The routed daemon, routed, listens on UDP port 520. It receives broadcast messages from routers or other hosts using Routing Information Protocol (RIP). These messages are used by your storage system to update its internal routing tables to determine which network interfaces are optimal for each destination.

Your storage system never broadcasts RIP messages containing routes because Data ONTAP is not capable of acting as a router.

RIP is not secure because an attacker can easily send artificial RIP messages and cause hosts running the routed daemon (such as your storage system) to redirect network traffic to the attacker. The attacker can then receive and shift this traffic for passwords and other information and send it on to the actual destination, where the intrusion is undetected. This method can also be used as a starting point for TCP session attacks.

Because of these security issues, use static routes (those set up using the `route` command on your storage system) instead of using the routed daemon.

# NDMP

Network Data Management Protocol (NDMP) runs on TCP port 10000 and is used primarily for backup of network-attached storage (NAS) devices, such as storage systems.

The protocol defines three authentication methods:

- NONE—allows authentication without restriction
- TEXT—sends a clear text password over the network, similar to Telnet or FTP
- MD5—uses the MD5 message digest algorithm along with a challenge-response message exchange to implement a secure login mechanism

Your storage system supports both the TEXT and MD5 authentication methods. Most NDMP-enabled backup software uses MD5 by default.

To entirely disable the TEXT authentication method, you should set the `ndmpd.authtype` option to `challenge`.

To restrict NDMP commands to certain authorized backup hosts, you should use the `ndmp.access` option.

Regardless of the authentication method used, NDMP sends backup data in decrypted format over the network, as does most other backup software. A separate network optimized for backup is a common means to increase performance while retaining data security.

To disable NDMP, you should set the `ndmp.enable` option to `off`.

# SnapMirror and SnapVault

SnapMirror and SnapVault use TCP port 10566 for data transfer. Network connections are always initiated by the destination system; that is, SnapMirror and SnapVault *pull* data rather than *push* data.

Authentication is minimal with both SnapMirror and SnapVault. To restrict inbound TCP connections on port 10566 to a list of authorized hosts or IP addresses, you should configure the `snapmirror.access` or `snapvault.access` option. When a connection is established, the destination storage system communicates its host name to the source storage system, which then uses this host name to determine if a transfer is allowed. You should confirm a match between the host name and its IP address. To confirm that the host name and the IP address match, you should set the `snapmirror.checkip.enable` option to `on`.

To disable SnapMirror, you should set the `snapmirror.enable` option to `off`. To disable SnapVault, you should set the `snapvault.enable` option to `off`.

# Error codes for the netdiag command

Network error codes are generated by the `netdiag` command. They describe the network problems and suggest the actions that you can take.

The following table lists some network error codes, describes the problems that the error codes point to, and suggests actions that you can take to fix the problems.

**Note:** Only a small fraction of the possible network error messages are presented in the following table. If you receive any problem code not listed in this table, contact your technical support.

| Error code | Description | Recommended actions |
|---|---|---|
| 201 | Link not detected. | Complete the following steps until you detect a link: 1. Ensure that the cable is connected between the switch port and your storage system interface, and that both ends are securely attached. 2. Ensure that the switch port and interface are both configured to the `up` status, and one of the following is true: • Autonegotiation is enabled on both sides • Autonegotiation is disabled on both sides, and the duplex and speed settings match 3. Because the switch port, cable, or NIC might be faulty, replace them, one by one, to locate the fault. 4. If the problem persists, contact your technical support. |
| 203 | No link is detected because of a speed mismatch. | Change the interface configuration or peer switch port configuration to match the speed. |
| 204 | The interface is not configured to the `up` status. | Configure the interface state to the `up` status. |
| 205 | Duplex mismatch. | Change the interface or peer switch port duplex setting so that they match. |
| 206 | Link capacity problem. | Upgrade to a faster interface. |

| Error code | Description | Recommended actions |
|---|---|---|
| 207 | The interface is not transmitting or receiving. | Complete the following steps:<br><br>1. Pull the network cable out from the network interface card.<br><br>2. Reinsert the cable.<br><br>3. Use `ifstat` to display statistics.<br><br>   • Link errors, such as CRC, are caused by a faulty switch port, cable, or NIC; replace them one by one to locate the fault.<br>   • Out-of-resource errors are caused by heavy loads.<br><br>4. If the problem persists, contact your technical support. |
| 208 | Excessive I/O errors. | Complete the following steps:<br><br>1. Reset the interface card.<br><br>2. Check the cables.<br><br>3. If the problem persists, contact your technical support. |
| 209 | Excessive unsupported protocol packets are being sent to your storage system. | The problem is not with your storage system.<br>Contact your network administrator to resolve the problem. |
| 301 | The IP address and the netmask are inconsistent with the assigned broadcast address. | Change the configuration by using the `ifconfig` command. |
| 302 | The broadcast address reaches a larger set of hosts than the standard broadcast computed from the IP address and netmask. | If this behavior is erroneous, change the configuration. |
| 303 | There are excessive IP reassembly errors. | Switch from NFS over UDP to NFS over TCP. |

| Error code | Description | Recommended actions |
|---|---|---|
| 401 | The TCP window advertised by the client is too small. | The problem is not with your storage system. Reconfigure the client. |
| 402 | There is excessive packet loss on the sending side. | The problem is not with your storage system. Examine the network and the client for congestion. |
| 403 | There is excessive packet loss on the receiving side. | The problem is not with your storage system. Examine the network and the client for congestion. |
| 404 | The average TCP packet size is poor on the receiving side because the network, client, or both are not enabled to support jumbo frames. | The problem is not with your storage system. Enable support for jumbo frames in network devices and the client. |
| 405 | The average TCP packet size is poor on the receiving side because of a problem with the network, client, or both. | The problem is not with your storage system. Examine the network and client for configured MTUs. |
| 406 | The average TCP packet size is poor on the receiving side because of a client application problem. | The problem is not with your storage system. Examine the client application data transmission strategy. |
| 407 | Excessive TCP listen socket drops because the system is overloaded or under security attack. | Contact your network administrator to resolve the problem. |
| 408 | There are excessive filtered TCP port drops because the system is under security attack. | Check your network. Contact your network administrator to resolve the problem. |

| Error code | Description | Recommended actions |
|---|---|---|
| 409 | There are excessive embryonic TCP connection drops because the system is under security attack or because a client has a bug. | A packet trace might assist in locating the problem. Contact your network administrator to resolve the problem. |
| 410 | Excessive TCP checksum errors. These errors can be caused by bad hardware on the client, in the network infrastructure (for example, blade in switch or router), or on the NIC. These errors can also be caused by a bug in the client. | • Check your client system for bugs.<br>• Replace hardware components until the problem is resolved.<br>• Contact your network administrator to resolve the problem. |
| 411 | There are packets because of a client. Your system might be under a security attack. | The problem is not with your storage system.<br>• Check your client system for bugs.<br>• Check for a security attack. |
| 451 | There are excessive UDP checksum errors. | Switch from NFS over UDP to NFS over TCP. |
| 601 | The DNS server is not reachable. | Examine the DNS server and the path to the DNS server. |
| 602 | The NIS server is not reachable. | Examine the NIS server and the path to the NIS server. |

# Copyright and trademark information

# 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