

Data ONTAP® 7.1

Network Management Guide

Network Appliance, Inc.
495 East Java Drive
Sunnyvale, CA 94089 USA
Telephone: +1 (408) 822-6000
Fax: +1 (408) 822-4501
Support telephone: +1 (888) 4-NETAPP
Documentation comments: doccomments@netapp.com
Information Web: <http://www.netapp.com>

Part number 210-01082_B0
December 2005

Copyright and trademark information

Copyright information

Copyright © 1994–2005 Network Appliance, Inc. All rights reserved. Printed in the U.S.A.

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

Portions of this product are derived from the Berkeley Net2 release and the 4.4-Lite-2 release, which are copyrighted and publicly distributed by The Regents of the University of California.

Copyright © 1980–1995 The Regents of the University of California. All rights reserved.

Portions of this product are derived from NetBSD, which is copyrighted by Carnegie Mellon University.

Copyright © 1994, 1995 Carnegie Mellon University. All rights reserved. Author Chris G. Demetriou.

Permission to use, copy, modify, and distribute this software and its documentation is hereby granted, provided that both the copyright notice and its permission notice appear in all copies of the software, derivative works or modified versions, and any portions thereof, and that both notices appear in supporting documentation.

CARNEGIE MELLON ALLOWS FREE USE OF THIS SOFTWARE IN ITS “AS IS” CONDITION. CARNEGIE MELLON DISCLAIMS ANY LIABILITY OF ANY KIND FOR ANY DAMAGES WHATSOEVER RESULTING FROM THE USE OF THIS SOFTWARE.

Software derived from copyrighted material of The Regents of the University of California and Carnegie Mellon University is subject to the following license and disclaimer:

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notices, this list of conditions, and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notices, this list of conditions, and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. All advertising materials mentioning features or use of this software must display the following acknowledgment:

This product includes software developed by the University of California, Berkeley and its contributors.
4. Neither the name of the University nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE REGENTS AND CONTRIBUTORS “AS IS” AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE REGENTS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER

IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

This software contains materials from third parties licensed to Network Appliance Inc. which is sublicensed, and not sold, and title to such material is not passed to the end user. All rights reserved by the licensors. You shall not sublicense or permit timesharing, rental, facility management or service bureau usage of the Software.

Portions developed by the Apache Software Foundation (<http://www.apache.org/>). Copyright © 1999 The Apache Software Foundation.

Portions Copyright © 1995–1998, Jean-loup Gailly and Mark Adler

Portions Copyright © 2001, Sitraka Inc.

Portions Copyright © 2001, iAnywhere Solutions

Portions Copyright © 2001, i-net software GmbH

Portions Copyright © 1995 University of Southern California. All rights reserved.

Redistribution and use in source and binary forms are permitted provided that the above copyright notice and this paragraph are duplicated in all such forms and that any documentation, advertising materials, and other materials related to such distribution and use acknowledge that the software was developed by the University of Southern California, Information Sciences Institute. The name of the University may not be used to endorse or promote products derived from this software without specific prior written permission.

Portions of this product are derived from version 2.4.11 of the libxml2 library, which is copyrighted by the World Wide Web Consortium.

Network Appliance modified the libxml2 software on December 6, 2001, to enable it to compile cleanly on Windows, Solaris, and Linux. The changes have been sent to the maintainers of libxml2. The unmodified libxml2 software can be downloaded from <http://www.xmlsoft.org/>.

Copyright © 1994–2002 World Wide Web Consortium, (Massachusetts Institute of Technology, Institut National de Recherche en Informatique et en Automatique, Keio University). All Rights Reserved. <http://www.w3.org/Consortium/Legal/>

Software derived from copyrighted material of the World Wide Web Consortium is subject to the following license and disclaimer:

Permission to use, copy, modify, and distribute this software and its documentation, with or without modification, for any purpose and without fee or royalty is hereby granted, provided that you include the following on ALL copies of the software and documentation or portions thereof, including modifications, that you make:

The full text of this NOTICE in a location viewable to users of the redistributed or derivative work.

Any pre-existing intellectual property disclaimers, notices, or terms and conditions. If none exist, a short notice of the following form (hypertext is preferred, text is permitted) should be used within the body of any redistributed or derivative code: "Copyright © [\$date-of-software] World Wide Web Consortium, (Massachusetts Institute of Technology, Institut National de Recherche en Informatique et en Automatique, Keio University). All Rights Reserved. <http://www.w3.org/Consortium/Legal/>.

Notice of any changes or modifications to the W3C files, including the date changes were made.

THIS SOFTWARE AND DOCUMENTATION IS PROVIDED "AS IS," AND COPYRIGHT HOLDERS MAKE NO REPRESENTATIONS OR WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO, WARRANTIES OF MERCHANTABILITY OR FITNESS

FOR ANY PARTICULAR PURPOSE OR THAT THE USE OF THE SOFTWARE OR DOCUMENTATION WILL NOT INFRINGE ANY THIRD PARTY PATENTS, COPYRIGHTS, TRADEMARKS OR OTHER RIGHTS.

COPYRIGHT HOLDERS WILL NOT BE LIABLE FOR ANY DIRECT, INDIRECT, SPECIAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF ANY USE OF THE SOFTWARE OR DOCUMENTATION.

The name and trademarks of copyright holders may NOT be used in advertising or publicity pertaining to the software without specific, written prior permission. Title to copyright in this software and any associated documentation will at all times remain with copyright holders.

Software derived from copyrighted material of Network Appliance, Inc. is subject to the following license and disclaimer:

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

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

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

Trademark information

NetApp, the Network Appliance logo, the bolt design, NetApp—the Network Appliance Company, DataFabric, Data ONTAP, FAServer, FilerView, MultiStore, NearStore, NetCache, SecureShare, SnapManager, SnapMirror, SnapMover, SnapRestore, SnapVault, SyncMirror, and WAFL are registered trademarks of Network Appliance, Inc. in the United States, and/or other countries. gFiler, Network Appliance, SnapCopy, Snapshot, and The Evolution of Storage are trademarks of Network Appliance, Inc. in the United States and/or other countries and registered trademarks in some other countries. ApplianceWatch, BareMetal, Camera-to-Viewer, ComplianceClock, ComplianceJournal, ContentDirector, ContentFabric, EdgeFiler, FlexClone, FlexVol, FPolicy, HyperSAN, InfoFabric, LockVault, Manage ONTAP, NOW, NOW NetApp on the Web, ONTAPI, RAID-DP, RoboCache, RoboFiler, SecureAdmin, Serving Data by Design, SharedStorage, Simulate ONTAP, Smart SAN, SnapCache, SnapDirector, SnapDrive, SnapFilter, SnapLock, SnapMigrator, SnapSuite, SnapValidator, SohoFiler, vFiler, VFM, Virtual File Manager, VPolicy, and Web Filer are trademarks of Network Appliance, Inc. in the United States and other countries. NetApp Availability Assurance and NetApp ProTech Expert are service marks of Network Appliance, Inc. in the United States. Spinnaker Networks, the Spinnaker Networks logo, SpinAccess, SpinCluster, SpinFS, SpinHA, SpinMove, and SpinServer are registered trademarks of Spinnaker Networks, LLC in the United States and/or other countries. SpinAV, SpinManager, SpinMirror, SpinRestore, SpinShot, and SpinStor are trademarks of Spinnaker Networks, LLC in the United States and/or other countries.

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

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

Network Appliance is a licensee of the CompactFlash and CF Logo trademarks.

Network Appliance NetCache is certified RealSystem compatible.

Table of Contents

	Preface	xi
Chapter 1	Network Interface Configuration	1
	Understanding the network interfaces on your storage system.	2
	Understanding frame size, MTU size, and jumbo frames	5
	Understanding Ethernet media types	8
	Understanding flow control.	10
	Configuring network interfaces.	12
	Configuring aliases for an interface	18
	Changing the status of an interface to Up or Down	20
	Displaying network interface information	21
	Diagnosing network problems	23
Chapter 2	ATM Configuration	25
	About ATM and ATM LANE	26
	Preparing the ATM adapter for LANE	31
	Verifying that the ATM adapter is installed and functioning	33
	Verifying a working connection to the ATM network.	34
	Verifying that the UNI is operational	35
	Configuring the LANE Configuration Server address.	37
	Configuring the ATM adapter for an Emulated LAN	40
	Adding an Emulated LAN to the ATM adapter	41
	Configuring the logical Ethernet interface	43
	Deleting an Emulated LAN from an ATM adapter	44
	Checking and completing the Emulated LAN configuration.	45
	Verifying the communications link	46
	Checking the configuration settings	47
	Checking the other elements of the Emulated LAN	48
	Modifying load balancing and failover	50
	Saving the ATM configuration commands in the /etc/rc file	52
	Saving the host and IP address data in the /etc/hosts file	53
	Understanding FORE/IP over SPANS	54
	Managing FORE/IP and PVCs	55

	Establishing FORE/IP PVCs on your storage system	56
	Displaying information about a FORE/IP PVC	58
	Displaying the FORE/IP configuration	59
	Changing the ATM adaptation layer for FORE/IP and SPANS	61
	Deleting a FORE/IP PVC	62
Chapter 3	Network Routing Configuration	63
	About routing in Data ONTAP	64
	About fast path	65
	About the routing table.	67
	Enabling and disabling routing mechanisms	70
	Displaying the routing table and default route information	72
	Modifying the routing table.	75
Chapter 4	Host-Name Resolution	77
	Maintenance of host information	78
	Using the /etc/hosts file to maintain host information	79
	Using DNS to maintain host information.	82
	Using dynamic DNS to update host information	89
	Using NIS to maintain host information	92
	Changing the host name search order	100
Chapter 5	Storage System Monitoring Using SNMP	103
	Understanding SNMP implementation in Data ONTAP	104
	Understanding traps in Data ONTAP	106
	Contents of the custom MIB	109
	Contents of the iSCSI MIB.	112
	Managing the SNMP agent	113
	Creating SNMP traps	119
	Understanding user-defined traps	120
	Defining or modifying a trap.	121
	SNMP trap parameters	126

Chapter 6	Virtual LAN (VLAN) Configuration	133
	Understanding VLANs	134
	VLANs in Data ONTAP	138
	Managing VLANs on your storage system	140
	Creating and configuring a VLAN on your storage system	141
	Adding an interface to a VLAN	144
	Deleting a VLAN	145
	Modifying VLAN interfaces	147
	Viewing VLAN statistics	148
 Chapter 7	 Configuring vifs	 151
	Understanding vifs	152
	Types of vifs	154
	Managing vifs	159
	Creating a single-mode vif	161
	Selecting an active interface in a single-mode vif	163
	Creating a multimode vif	165
	Adding interfaces to a vif	167
	Deleting an interface from a vif	168
	Displaying the status of a vif	169
	Displaying statistics of a vif	173
	Destroying a vif	174
	Second-level vifs	175
	Understanding second-level vifs on a single storage system	176
	Creating a second-level vif on a single storage system	177
	Understanding second-level vifs in a cluster	179
	Creating a second-level vif in a cluster	181
 Chapter 8	 Internet Protocol Security Configuration	 185
	Understanding IPsec	186
	Setting up IPsec	190
	Managing security policies	192
	Viewing security associations	198
 Appendix A	 Network Interface Statistics	 199
	Fast Ethernet interface statistics	200

	Gigabit Ethernet and Ethernet Controller IV interface statistics	.204
	Statistics for FAS250/ FAS270 network interfaces	.209
	FDDI interface statistics	.213
	ATM interface statistics	.217
Appendix B	Improving storage system performance	.219
Appendix C	IP port usage on NetApp storage systems	.221
Appendix D	Netdiag Error Codes	.235
	Index	.241

Preface

About this guide

This guide describes how to configure and manage network interfaces, virtual network interfaces (vifs), virtual LANs (VLANs), and routing on storage systems that run Data ONTAP® 7.1 software. The guide describes all Data ONTAP storage system models; however, some models do not support all of the networking interfaces. See the hardware guide for your storage system to identify which interfaces are supported on your system.

Audience

This guide is for system administrators who are familiar with operating systems that run on the Data ONTAP storage system clients, such as UNIX®, Windows 95™, Windows NT®, and Windows® 2000. It also assumes that you are familiar with how the Network File System (NFS), Common Interface File System (CIFS), and HyperText Transfer Protocol (HTTP) protocols are used for file sharing or transfers. This guide does not cover basic system or network topics, such as IP addressing, routing, and network topology; it emphasizes the characteristics of the Data ONTAP storage systems.

Data ONTAP user interfaces

You can perform Data ONTAP administrative procedures described in this guide using either of two kinds of user interfaces:

- ◆ The command-line interface
You enter commands at the storage system command line, from one of three places:
 - ❖ A system console
 - ❖ A client computer that can access the storage system through a Telnet session
 - ❖ A client computer that can access the storage system through a Remote Shell connection
- ◆ The FilerView® administration tool's interface
You use the FilerView Web-based graphical management interface to select, view, or enter information.

In this guide, administrative procedures are described for both the command-line and FilerView interfaces, except where a particular procedure can only be performed at the command line.

The FilerView descriptions in this guide assume that you have already started FilerView in a web browser as described in the *System Administration Guide*.

For more information about administering a storage system using these methods, see the *System Administration Guide*.

Accessing Data ONTAP man pages

Data ONTAP provides manual (man) pages for the types of information listed in the following table. The man pages are grouped into sections according to standard UNIX naming conventions.

Types of information	Man page section
Commands	1
Special files	4
File formats and conventions	5
System management and services	8

Man pages can be viewed in the following ways:

- ◆ At the storage system command line, by entering
`man command_or_file_name`
- ◆ From the FilerView main navigational page
- ◆ In the *Command Reference Guide*

Note

All Data ONTAP man pages are stored on the storage system in files whose names are prefixed with the string “na_” to distinguish them from client man pages. The prefixed names are used to refer to Data ONTAP man pages from other man pages and sometimes appear in the NAME field of the man page, but the prefixes are not part of the command, file, or services.

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

Terminology and conventions

Network Appliance’s storage products (filers, FAS appliances, and NearStore® systems) are all storage systems—also sometimes called *filers* or *storage appliances*.

In examples that illustrate commands executed on a UNIX workstation, the command syntax and output might differ, depending on your version of UNIX.

This guide uses the term “type” to mean pressing one or more keys on the keyboard. It uses the term “enter” to mean pressing one or more keys and then pressing the Enter key, or clicking in a field in a graphical interface and typing information into it.

Keyboard conventions

When describing key combinations, this guide uses the hyphen (-) to separate individual keys. For example, “Ctrl-D” means pressing the “Control” and “D” keys simultaneously. Also, this guide uses the term “Enter” to refer to the key that generates a carriage return, although the key is named “Return” on some keyboards.

Typographic conventions

The following table describes typographic conventions used in this guide.

Convention	Type of information
<i>Italic font</i>	Words or characters that require special attention. Placeholders for information you must supply. For example, if the guide says to enter the <code>arp -d hostname</code> command, you enter the characters “arp -d” followed by the actual name of the host. Book titles in cross-references.
Monospaced font	Command and daemon names. Information displayed on the system console or other computer monitors. The contents of files.
Bold monospaced font	Words or characters you type. What you type is always shown in lowercase letters, unless you must type it in uppercase letters.

Special messages

This guide contains special messages that are described as follows:

Note_____

A note contains important information that helps you install or operate the system efficiently.

Attention

An attention contains instructions that you must follow to avoid damage to the equipment, a system crash, or loss of data.

About this chapter

This chapter discusses the following:

- ◆ The types of interfaces supported on your storage system
- ◆ Concepts related to setting up and using network interfaces on your storage system
- ◆ How the interfaces are named
- ◆ How to configure the network interfaces on your storage system
- ◆ How you can obtain detailed statistics on various interfaces supported on your storage system

Topics in this chapter

This chapter covers the following topics:

- ◆ [“Understanding the network interfaces on your storage system”](#) on page 2
- ◆ [“Understanding frame size, MTU size, and jumbo frames”](#) on page 5
- ◆ [“Understanding Ethernet media types”](#) on page 8
- ◆ [“Understanding flow control”](#) on page 10
- ◆ [“Configuring network interfaces”](#) on page 12
- ◆ [“Configuring aliases for an interface”](#) on page 18
- ◆ [“Changing the status of an interface to Up or Down”](#) on page 20
- ◆ [“Displaying network interface information”](#) on page 21
- ◆ [“Diagnosing network problems”](#) on page 23

Understanding the network interfaces on your storage system

Types of interfaces your storage system supports

Your storage system supports the following interface types:

- ◆ Ethernet—including quad-port Ethernet adapters
- ◆ Gigabit Ethernet (GbE)
- ◆ Fiber Distributed Data Interface (FDDI)
- ◆ Asynchronous Transfer Mode (ATM)—Emulated LAN and FORE/IP

Your storage system also supports the following virtual network interface types:

- ◆ Virtual interface (vif)
- ◆ Virtual local area network (VLAN)
- ◆ Virtual hosting (vh)

Data ONTAP imposes a limit of 128 network interfaces (including physical, vif, VLAN, vh, and loopback interfaces) per storage system.

How interfaces are named

For physical interfaces, the interface names are assigned automatically based on the slot in which the network adapter is installed.

VLAN interfaces are displayed in the *interfaceID_and_slot_number-vlan_id* format, where *slot_number* is the slot in which the network adapter is installed and *vlan_id* is the identifier of the VLAN configured on the interface. For example, e8-2, e8-3, and e8-4 are three VLAN interfaces for VLANs 2, 3, and 4, configured on interface e8.

You can assign names for vifs and the emulated LAN interfaces.

You can use the `ifconfig` command-line interface (CLI) command or FilerView to display network interfaces on your storage system. For more information, see [“Configuring network interfaces”](#) on page 12.

How multiple ports are identified

Some Ethernet adapters support four ports. In the Data ONTAP context, these adapters are referred to as *quad-port Ethernet interfaces*, and sometimes shortened to just *quad-port interfaces*. Data ONTAP uses a letter to refer to each port on a quad-port interface. The following table shows the relationship of port numbers to letters.

Port number	Letter
1	a
2	b
3	c
4	d

Interface naming conventions

The following table lists interface types, their identifiers, and examples of names that use the identifiers.

Interface type	Interface type ID	Examples of names
Ethernet (single) and Gigabit Ethernet	e	e0 e1
Ethernet (quad-port)	e	e0a e0b e0c e0d e1a e1b
vif	Any user-specified string that meets the criteria specified in “Prerequisites” on page 161.	web_vif proxy_vif
VLAN	e	e8-2 e8-3
FDDI	f	f0 f1

Interface type	Interface type ID	Examples of names
ATM	a (used only to configure clusters)	a0 a1
ATM—Emulated LAN	e1 (default)	e10
ATM—Fore IP	fa	fa0

How Data ONTAP creates host names

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

Examples of host names

Example 1: A storage system named *toaster* with a single Ethernet interface in slot 0 and a quad-port Ethernet interface in slot 1 has the host names given in the following table.

Interface	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

Example 2: A storage system named *toaster* with a Fiber Distributed Data Interface (FDDI) interface in slot 1 has the host name *toaster-f1*.

Understanding frame size, MTU size, and jumbo frames

When to change the default frame size

The standard Ethernet (IEEE 802.3) frame size is 1,518 bytes. The default frame size can be changed on the following types of network interfaces:

Gigabit Ethernet interfaces: Increasing the default frame size for any Gigabit Ethernet interface supported on your storage system, as well as the Gigabit Ethernet infrastructure to which it connects, can significantly increase performance depending upon the activity.

ATM ELAN interfaces: If you need to change the frame size for an ATM Emulated LAN (ELAN) interface, you cannot do it on your storage system; you must change it on the switch to which the storage system connects.

FDDI interfaces: You can set the frame size on an FDDI interface when the network infrastructure does not support the default FDDI frame size.

The default frame size can be changed by reconfiguring the Maximum Transmission Unit (MTU) value on storage systems and clients.

Frame size and MTU size definitions

Two commonly used terms to describe frame characteristics are *frame size* and *MTU size*.

Frame size: 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).

MTU size: The MTU size specifies the maximum number of bytes of data (the 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 your storage systems. However, a *jumbo* frame, with an MTU size of 9,000 bytes or more, can also be configured.

About jumbo frames

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 sizes are 9,018 bytes and higher.

Because jumbo frames are larger than standard frames, fewer frames are needed and therefore CPU processing overhead is reduced.

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

The default frame size defined for a jumbo frame on your storage system is 9,018 bytes, and the maximum supported size is 9,210 bytes. These sizes correspond to MTU sizes of 9,000 and 9,192, respectively.

Ways to set up jumbo frames on your storage system

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. For information about how to set the MTU size, see [“Configuring network interfaces”](#) on page 12.

Network infrastructure requirements

Before you enable jumbo frames on your storage system, clients and intermediate routers on the network must have jumbo frames enabled. In particular, the following network infrastructure requirements (as appropriate) must be satisfied:

- ◆ The switch ports must have jumbo frames enabled.
- ◆ If your storage system and the client are on different subnets, the next-hop router must be configured for jumbo frames.
- ◆ Jumbo frames must be enabled on client interfaces.

Client configuration guidelines

Follow these guidelines in configuring clients to work with jumbo frames:

- ◆ Configure jumbo frames on the client as well as on your storage system. Find out 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.

- ◆ Ensure that the User Datagram Protocol (UDP) clients are configured with the same MTU size as your storage system.

UDP clients do not communicate their MTU size. Therefore, your storage system and the client should be configured with the same MTU size, or the storage system might send packets that the clients cannot receive.

- ◆ Check the MTU 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.

Understanding Ethernet media types

About media types You can configure the speed and the duplex setting, or specify autonegotiation, for an Ethernet interface. The media types available for your storage system interfaces are described in the following table.

Note For 10Base-T and 100Base-T interfaces, the `mediatype` option of the interface and its link partner (the interface on the other end of the connection) must be the same; that is, both interfaces must be configured either for speed and duplex or for autonegotiation. Otherwise, a duplex mismatch occurs, which can lead to poor performance.

Media-type value	Description
tp	10Base-T, half-duplex
tp-fd	10Base-T, full-duplex
100tx	100Base-T, half-duplex
100tx-fd	100Base-T, full-duplex
auto	Autonegotiate speed, duplex, and flow control. See “How media type auto works” on page 8.

For information about setting media-type values, see [“Configuring network interfaces”](#) on page 12.

How media type auto works The behavior of media type `auto` is determined by the type of network adapter installed on your storage system. The following table lists the parameters that are autonegotiated and the possible values for those parameters for each type of network adapter.

Note You can use the `ifstat` command to determine the speed, duplex, and flow control settings that are negotiated between the interface of your storage system and the link partner. For more information, see [“Displaying network interface information”](#) on page 21.

Network adapter	Parameters that are autonegotiated (with possible values)
10Base-T/100Base-T	Speed and duplex (half or full)
100Base-T/1000Base-T	Speed. If speed is 100 Mbps, duplex (half or full) and flow control are negotiated. If speed is 1000 Mbps, flow control is negotiated.
10Base-T/100Base-T/ 1000Base-T	Speed. If speed is 10 Mbps, duplex (half or full) is negotiated. If speed is 100 Mbps, duplex and flow control are negotiated. If speed is 1000 Mbps, flow control is negotiated.
Gigabit Ethernet	Flow control.

Understanding flow control

About flow control

Flow control is the management of the flow of frames between two directly connected link-partners. To achieve flow control, you 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. Thus, flow control can reduce or eliminate dropped packets due to overrun.

What the `flowcontrol` option is used for

Flow control can be configured for the following interfaces:

- ◆ Gigabit Ethernet
- ◆ 100Base-T/1000Base-T and 10Base-T/100Base-T/1000Base-T

This *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.

Flow control types

The following table describes the types you can specify for the `flowcontrol` option.

Flow control value	Description
none	No flow control
receive	Ability to receive flow control frames
send	Ability to send flow control frames
full	Ability to send and receive flow control frames

How to configure and view flow control

See “[Configuring network interfaces](#)” on page 12 for information on configuring and displaying flow control settings. You can also use the `ifstat` command to view the operational flow control setting. If you do not specify the `flowcontrol` option when configuring a network interface, the configured flow control setting defaults to full.

Configuring network interfaces

What network interface configuration includes

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 and broadcast address
- ◆ Set hardware-dependent values such as media type, MTU size, and flow control
- ◆ Specify whether the interface is attached to a network with firewall security protection
- ◆ Specify whether the network interface is to 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 on a cluster partner for takeover mode
- ◆ View the current configuration of a specific interface or all interfaces that exist on your storage system

Additional network interface configuration tasks include the following:

- ◆ [“Configuring aliases for an interface”](#) on page 18
- ◆ [“Changing the status of an interface to Up or Down”](#) on page 20
- ◆ [“Displaying network interface information”](#) on page 21

About configuration tools

The following tools are available for storage system network configuration and management.

Command-line interface	Graphical interface
<code>ifconfig</code> command For more information, see the <code>na_ifconfig(1)</code> man page.	FilerView Network windows For more information, see FilerView help.

How interface configuration works

You assign initial network interface configuration values when new interfaces are created. The method you use to configure the interface depends on your preference of command-line interface (the `ifconfig` command) versus graphical user interface (FilerView).

An `ifconfig` command is included in the `/etc/rc` file of the root volume for each storage system interface that you specify an IP address for 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 interfaces on subsequent storage system reboots.

Note

You can use the `ifconfig` command to change values of parameters for an 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.

When you use FilerView to make changes, the changes are automatically written to the `/etc/rc` file.

Viewing and modifying interface settings at the command line

To view or modify interfaces with the `ifconfig` command, complete the following step.

Step	Action
1	At your storage system command line, enter <code>ifconfig interface_name parameters</code> For more information on <code>ifconfig</code> parameters, see ◆ “Command syntax for viewing interface settings” on page 14 ◆ “Command syntax for modifying interface settings” on page 14

Viewing and modifying interface settings with FilerView

To view or modify interfaces with FilerView, complete the following steps.

Step	Action
1	In FilerView, click Network in the list on the left.
2	In the list under Network, click Manage Interfaces.

Step	Action	
3	If you want to....	Then...
	View interface configuration details	Click Show All Interface Details.
	Modify an interface configuration	Click Modify for the interface you want. Examples of configuration values are listed in “ Command syntax for modifying interface settings ” on page 14

Command syntax for viewing interface settings

The following table describes the command syntax for viewing network interface configuration settings.

To view ...	Use this command syntax ...
A single interface	<code>ifconfig interface_name</code>
All interfaces	<code>ifconfig -a</code>

Command syntax for modifying interface settings

The following table shows how to set specific network interface parameters using the `ifconfig` command. For more information about each task, see the `na_ifconfig(1)` man page.

To modify this parameter ...	What the parameter is for...	Use this command syntax ...
IP address	To configure an IP address for the specified interface	<code>ifconfig interface_name IP_address</code> Example 1: To configure a quad-port Ethernet interface <code>e3a</code> with an IP address <code>192.168.25.10</code> , enter <code>ifconfig e3a 192.168.25.10</code>

To modify this parameter ...	What the parameter is for...	Use this command syntax ...
Network mask	To specify a subnet mask for the specified interface	<pre>ifconfig interface_name netmask mask</pre> Example 1: To configure a 24-bit mask for the interface e3a configured in the previous example, enter <pre>ifconfig e3a netmask 255.255.255.0</pre> Note_____ By default, your storage system creates a network mask based on the class of the address (Class A, B, C, or D). However, if you have created subnets that do not match the class boundary of the IP address, you must specify a network mask.
Broadcast address	To specify an address that when used enables you to send a message to all machines on a network	<pre>ifconfig interface_name broadcast address</pre> Example: To set a broadcast address of 192.168.25.250 for the network 192.168.25.10 with subnet mask 255.255.255.0, enter <pre>ifconfig e3a broadcast 192.168.25.250</pre>
Media type	To configure speed and duplex for an interface	<pre>ifconfig interface_name mediatype value</pre> Example: To configure the interface e2 as a 100Base-TX full-duplex interface, enter <pre>ifconfig e2 mediatype 100tx-fd</pre>
MTU	To specify an MTU size for transmission between your storage system and its clients	<pre>ifconfig interface_name mtusize size</pre> Example: To specify an MTU size of 9000 for a Gigabit Ethernet interface e8, enter <pre>ifconfig e8 mtusize 9000</pre>

To modify this parameter ...	What the parameter is for...	Use this command syntax ...
Flow control	To specify the flow control type For more information, see the <code>na_ifconfig(1)</code> man page.	<pre>ifconfig interface_name flowcontrol value</pre> Example: To turn off flow control on interface e8, enter <pre>ifconfig e8 flowcontrol none</pre>
Trusted/ Untrusted	To declare an interface to be 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 initiate a ping, the ICMP response packets received on the interface will be dropped.	<pre>ifconfig interface_name trusted untrusted</pre> Example: To specify that the network attached to interface e8 is not trusted for firewall security, enter <pre>ifconfig e8 untrusted</pre>
WINS	To enable an interface to register with WINS when CIFS is running. By default, network interfaces are registered with a WINS server when CIFS is running.	<pre>ifconfig interface_name wins -wins</pre> Example: To disable interface e8 from registering with WINS servers, enter <pre>ifconfig e8 -wins</pre>
Partner IP address	To specify the IP address of an interface on the cluster partner and to specify the interface that will assume this interface during takeover	<pre>ifconfig interface_name partner address</pre> Example: To specify the IP address of an interface on the cluster partner that will be assumed by interface e8 if the cluster partner fails, enter <pre>ifconfig e8 partner 192.168.25.10</pre>

To modify this parameter ...	What the parameter is for..	Use this command syntax ...
nfo / -nfo	To specify whether negotiated failover is turned On or Off on the interface This parameter works in tandem with the cluster failover option <code>cf.takeover.on_network_interface_failure</code>. If the <code>nfo</code> parameter is turned On for an interface and the cluster failover option is enabled, negotiated takeover between the cluster nodes can occur if this network interface fails. This parameter cannot be used for an interface that is part of a vif. You must include this option in the <code>/etc/rc</code> file for it to persist across reboots. For more information about the <code>nfo</code> parameter, see the <code>na_ifconfig(1)</code> man page. For more information about the <code>cf.takeover.on_network_interface_failure</code> option, see the <code>na_options(1)</code> man page.	<pre>ifconfig interface_name nfo</pre> Example: To enable negotiated failover on an interface <code>e8</code> of a cluster, enter <pre>ifconfig e8 nfo</pre> Note _____ Remember to enable the <code>cf.takeover.on_network_interface_failure</code> option after enabling negotiated failover with the above command. _____

Configuring aliases for an interface

About aliases

An alias is an alternative IP address for an interface. An alias can be useful when you are changing the IP address of an interface to a new address, but also want to keep accepting packets addressed to the old IP address.

There are two alias options available for the `ifconfig` command:

- ◆ `alias`—Establishes an alternative IP address for an interface.
- ◆ `-alias`—Removes an alternative IP address (alias) for an interface.

Note

Aliases for interfaces cannot be managed with FilerView.

Using the alias options

You can use the `alias` option at your storage system command line. However, the IP address configured using the `alias` option at the command line is lost if the storage system reboots. If you want to make your changes persistent across reboots, include these changes in the `/etc/rc` file of the root volume.

You cannot set up an IP address and an alias for an interface with one `ifconfig` command; you must configure the IP address for the interface before setting up the alias.

The `-alias` option is useful when you want to stop using the IP address originally configured on an interface but do not want to reboot your storage system.

Configuring and removing aliases

To set or remove an alias for an interface, complete the following step.

Step	Action
1	Enter the following command: <pre>ifconfig <i>interface_name</i> [<i>alias</i> -<i>alias</i>] <i>address</i> [<i>netmask mask</i>]</pre> Example: In the following example, the interface e0 (already configured with an IP address of 172.28.50.21) is set up with an alias IP address of 172.28.50.30: <pre>ifconfig e0 alias 172.28.50.30 netmask 255.255.255.0</pre> Example: The following example removes the 172.28.50.30 alias for the interface e0 set in the previous example: <pre>ifconfig e0 -alias 172.28.50.30</pre>

Changing the status of an interface to Up or Down

When you might change the status of an interface

You might have to change the status of an interface to Up or to Down in the course of doing one of the following:

- ◆ Installing a new interface
- ◆ Upgrading an interface
- ◆ Troubleshooting network connectivity issues
- ◆ Disabling a failed interface

Changing the status with ifconfig

To change the status of an interface to Up or to Down at the command line, complete the following step.

Step	Action
1	Enter the following command: <code>ifconfig interface {up down}</code>

Changing the status with FilerView

To change the status of an interface to Up or to Down using FilerView, complete the following steps.

Step	Action
1	In FilerView, click Network in the list on the left.
2	In the list under Network, click Manage Interfaces.
3	Click Up or Down in the Status field for the interface you want.

Displaying network interface information

Commands for displaying network interface statistics

Data ONTAP provides several commands that you can use to display statistics about network interface status and performance. The following table lists the commands and key information they display.

Command	Information displayed
<code>ifconfig -a</code>	◆ Interface status (up or down) ◆ Configuration parameters
<code>ifstat</code>	◆ Packets sent and received ◆ Collisions and other errors ◆ Negotiated media type settings between storage system interfaces and link partners
<code>netstat</code>	◆ Active sockets for each protocol ◆ Memory buffer (mbuf) pool usage ◆ Protocol-specific statistics for all protocols or a single protocol ◆ Cumulative or continuous packet traffic for all interfaces or a single interface ◆ Routing tables; for more information, see “Displaying the routing table and default route information” on page 72.

For more information, see the man pages on your storage system for these commands or see the Data ONTAP 7.1 *Command Reference Guide*.

You can also use FilerView to display selected interface and routing information. See [“Displaying interface information with FilerView”](#) on page 22 for more information.

About the ifstat command

The `ifstat` command displays statistics maintained by the networking code, network adapter, and network driver. The statistics displayed are gathered from the time of the last reboot or from the last time you cleared them.

Note

If you use the `ifstat` command on a storage system that is part of a cluster, the resulting information pertains only to the storage system on which the command was run. The information does not include statistics for the cluster partner.

The output of the `ifstat` command might contain many kinds of information, because different types of interfaces—for example, Ethernet, Gigabit Ethernet, FDDI, and ATM—generate different types of statistics. For the detailed statistics displayed for each network interface, see Appendix A, “[Network Interface Statistics](#),” on page 199.

Displaying interface information with FilerView

The Network Report in FilerView presents selected network interface statistics and routing information. It provides the information you would get by running all the following commands:

- ◆ `netstat -i`
- ◆ `routed status`
- ◆ `netstat -rn`

To display the Network Report, complete the following steps.

Step	Action
1	In FilerView, click Network in the list on the left.
2	In the list under Network, click Report.

Diagnosing network problems

About diagnosing network problems

The `netdiag` command specifies that any network problems be continuously diagnosed.

After you enter this command, Data ONTAP continuously gathers and analyzes statistics and performs diagnostic tests to identify and report problems related to the physical, network, or transport layers. If any problems are found, the command output also suggests remedial actions.

For information about all the options available with the `netdiag` command, see the `na_netdiag(1)` man page.

For a list of the `netdiag` error codes, see Appendix D, “[Netdiag Error Codes](#),” on page 235.

Diagnosing transport layer problems

To diagnose transport layer problems in your storage system, complete the following step.

Step	Action
1	Enter the following command: <pre>netdiag -t</pre> Sample result: A storage system whose TCP window size is smaller than the recommended value displays the following output: <pre>Performing transport layer diagnostics..... The TCP receive window advertised by CIFS client 10.10.10.10 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.</pre>

Other ways to troubleshoot network problems

To test whether your storage system can reach other hosts on your network, you can use the `ping` command.

About this chapter

This chapter outlines the features and concepts of Asynchronous Transfer Mode (ATM) and describes the key components of an Emulated LAN, including the LAN Emulation (LANE) Client, LANE Server, LANE Configuration Server, and Broadcast and Unknown Server (BUS). It also describes how to use FORE/IP over Simple Protocol for ATM Network Signaling (SPANS).

Topics in this chapter

This chapter discusses the following topics:

- ◆ [“About ATM and ATM LANE”](#) on page 26
- ◆ [“Preparing the ATM adapter for LANE”](#) on page 31
- ◆ [“Configuring the ATM adapter for an Emulated LAN”](#) on page 40
- ◆ [“Checking and completing the Emulated LAN configuration”](#) on page 45
- ◆ [“Understanding FORE/IP over SPANS”](#) on page 54
- ◆ [“Managing FORE/IP and PVCs”](#) on page 55

About ATM and ATM LANE

What ATM is

ATM is a network technology that combines the features of cell-switching and multiplexing to offer reliable and efficient network services. ATM provides an interface between the network and devices such as workstations and routers. The asynchronous nature of ATM means that bandwidth is made available on demand instead of slots of transmission time allocated to network devices, as in a synchronous system employing Time-Division Multiplexing (TDM).

ATM employs fixed-sized cells of 53 bytes each as the basic unit of transmission. Each cell consists of a 5-octet header, identifying the source of the transmission and other information, and a 48-octet payload containing the user data and headers for higher-level protocols. This architecture permits text, voice, graphics, and video to share the same network without any one source dominating network bandwidth.

ATM employs a star topology with an ATM switch acting as the hub of the network. All devices are connected directly to this hub, making network configuration and troubleshooting more straightforward, as well as offering dedicated bandwidth to the central switch.

Ways to use ATM on your storage system

You can use ATM in two ways on your storage system:

- ◆ ATM LANE, which provides the services of an Ethernet LAN to higher-level network application software
- ◆ FORE/IP over Permanent Virtual Connection (PVC) or Switched Virtual Connection (SVC), using SPANS to establish the SVCs

Your storage system can simultaneously support FORE/IP and LANE over User-Network Interface (UNI) 3.0 or 3.1 on the same physical interface.

Note

Data ONTAP uses conventional IP routing table lookups for routing all traffic on a FORE/IP ATM interface. For more information, see [“About Data ONTAP routing”](#) on page 64.

Differences between LANs and ATM

In addition to the issue of connection-based versus connectionless service, LANs differ from ATM in the following ways:

- ◆ The shared medium approach of LANs makes them ideal for broadcast and multicast messages.
- ◆ The Media Access Control (MAC) addresses used to identify a network interface of a LAN are typically based on the manufacturer's serial numbers. This means that the address is constant, independent of the network to which the interface is connected.

About LANE

Many organizations use a LAN for their internal data communications. Examples of these LANs include Ethernet/IEEE 802.3 and IEEE 802.5 (Token Ring). However, LANs typically offer a connectionless service, while ATM is always connection-oriented. This means that to use LAN-based applications using ATM, some form of LANE is required.

Benefits provided by LANE

LANE is an ATM service that offers the following benefits:

- ◆ You can run LAN-based application software on an ATM network.
- ◆ You can interconnect ATM networks to conventional LANs with existing bridging methods.

This permits applications running on ATM-connected end systems to interoperate with those running on traditional LAN-based devices. These LAN-based end systems can also communicate with each other across the ATM network.
- ◆ You can run more than one Emulated LAN on the same ATM network, with each Emulated LAN independent of the others.

About ATM cause codes

Data ONTAP displays cause code strings when ATM connections for LANE Configuration Server, LANE Server, Broadcast and Unknown Server, or LAN Emulation Client normally or abnormally terminate. They describe the reason for the connection termination or rejection. For more information about these cause codes, see the ATM Forum's UNI 3.0 and 3.1 specifications.

What an Emulated LAN is

An Emulated LAN comprises a group of ATM-attached devices, logically analogous to a group of LAN stations attached to an Ethernet/802.3 LAN segment. You can configure several Emulated LANs within an ATM network, and membership in an Emulated LAN is independent of where the end station is physically connected. The end station can belong to multiple Emulated LANs.

Components of an Emulated LAN

An Emulated LAN consists of the following components:

- ◆ A single LANE Service, which itself consists of a LANE Server and a BUS. Each of these components is discussed in more detail later in this chapter.
- ◆ A set of LANE Clients.
A LANE Client communicates with other LANE Clients and with the LANE Service using Virtual Channel Connections (VCCs) in an ATM SVC environment.

What a LANE Client is

The LANE Client is part of an ATM end station or a MAC bridge. It performs data forwarding as well as address resolution, among other control functions. The LANE Client supplies higher-level software with an Ethernet/IEEE 802.3 MAC layer interface that enables LAN-based application software to communicate over ATM networks just as it would over a traditional LAN.

How LANE Clients communicate

LANE Clients communicate with other clients using the LANE Service and represent users by their MAC addresses. A LANE Client employs separate VCCs for data and control communication, including LAN Emulation Address Routing Protocol (LE_ARP) requests for address resolution. User data intended for another end station is encapsulated in IEEE 802.3 frames.

What LANE Service is

The LANE Service, consisting of a LANE Server, BUS, and LANE Configuration Server, can be implemented as part of one or more end systems or as part of the ATM switch. When you implement the service in a distributed fashion over multiple devices, benefits include parallel operation as well as better error recovery through redundancy.

Within the LANE Service, the LANE Server is responsible for coordinating the control functions, while the LANE Configuration Server serves network clients by supplying Emulated LAN configuration information. The BUS forwards broadcast and multicast frames and handles unresolved unicast frames.

What a LANE Server does

The LANE Server performs control services within an Emulated LAN, including registering MAC addresses and resolving these addresses to corresponding ATM addresses when requested. LANE Clients can also request the LANE Server to resolve a route descriptor to an ATM address. The LANE Server responds to the client or forwards the request to other clients, which might be in a better position to service the query.

The LANE Server also coordinates the process of a LANE Client joining an Emulated LAN. There is a single LANE Server per Emulated LAN, and each LANE Server has a unique ATM address.

How a LANE Configuration Server functions in an administrative domain

The LANE Configuration Server maintains information concerning all the Emulated LANs in an administrative domain, and supplies the LANE Client with the ATM address for the LANE Server in the domain. Before joining an Emulated LAN, the LANE Client first exchanges configuration information with the LANE Configuration Server.

Upon successfully communicating with the LANE Configuration Server, the LANE Client receives a list of Emulated LANs that are available to join. There is a single LANE Configuration Server per administrative domain for Emulated LANs within a domain.

What a BUS does

The BUS accepts and processes data sent by a LANE Client to the broadcast MAC address 'FFFFFFFFFFFF'. The BUS also handles all multicast messages, as well as initial unicast frames sent by a LANE Client before the ATM address has been resolved.

The BUS thereby offers services that emulate the shared medium capabilities typical of a LAN. The BUS does this by serializing the frames and retransmitting them to the appropriate LANE Clients within the Emulated LAN.

Although there might be multiple BUSES defined within an Emulated LAN, each LANE Client is associated with only a single BUS per Emulated LAN.

UNI load balancing

The User-Network Interface (UNI), which serves as an interface point between ATM end systems and the ATM switch, supports both automatic adapter failover and load balancing across multiple adapters connected to the same physical ATM switch. This means that the UNI signaling module automatically detects which adapters are connected to a single physical network and places all adapters connected to that network in a failover group.

Both incoming and outgoing connections are directed to the least used adapter in a load-balancing group. If an adapter in a load-balancing group fails, the connections on that adapter are automatically transferred to another adapter.

UNI load balancing and adapter failover do not require any configuration. However, you can statically configure or disable UNI load balancing.

How LANE handles addressing and address resolution

LANs use a MAC address to designate the source and destination addresses for end stations. For LANE to function transparently, it must offer similar functionality. In practical terms, this means that each LANE Client has a MAC address, and when more than one LANE Client uses the same network interface, each LANE Client is assigned a different MAC address.

When the LANE Client needs to send data to another MAC address, it must first resolve that address to an ATM address, thus enabling it to establish a data-direct VCC to that LANE Client. To do so, it sends an LE_ARP_REQUEST to the LANE Server. The LANE Server can either respond to this request or forward it to other LANE Clients. If the specified MAC address is known anywhere on the Emulated LAN, the originating LANE Client gets an LE_ARP_RESPONSE frame containing the corresponding ATM address.

LANE standards supported in this release

This release of Data ONTAP supports the following features and standards:

- ◆ ATM Forum LANE Version 1.0 LANE Client Support
- ◆ UNI 3.0 and 3.1
- ◆ Integrated Local Management Interface (ILMI) Address Registration
- ◆ ILMI Management Information Base (MIB) extensions for LANE

The software works with the FORE OC3 ATM network interface. The software provides Ethernet LANE services, with the capability to configure multiple Emulated LANs on each available network interface.

The current release does not support ATM LANE 2.0, Multiprotocol Over ATM (MPOA), Classical IP (CLIP), or Token Ring LANE services.

Preparing the ATM adapter for LANE

Preparing for ATM LANE

Before the ATM adapter can communicate using ATM LANE, you need to ensure that the ATM adapter is installed correctly and that it can communicate with the network. This section describes the steps you take to enable the ATM adapter on your storage system to communicate using ATM LANE.

Prerequisites for configuring ATM adapters

Before you start configuring the ATM adapters in your storage system, ensure that you meet the prerequisites in the following table.

Prerequisite	Explanation
Complete the normal setup procedure for your storage system, run it automatically when you first install your storage system, or run the <code>setup</code> command for an existing installation.	You need an ATM switch with one or more Emulated LANs already configured on the switch (with the corresponding configurations for the LANE Configuration Server, LANE Server, and BUS).
Know the LANE Configuration Server address for the Emulated LAN that you want your storage system to join.	In most cases, the LANE Configuration Server has been configured to use the “well-known” address; however, this might be different at your site.
If your site has multiple Emulated LANs, know the ATM address of the LANE Configuration Server for each Emulated LAN you want a client to join.	You can configure each ATM adapter in your storage system to communicate over multiple Emulated LANs on the network.

Note

If you need more information about creating an Emulated LAN or configuring the LANE Configuration Server, LANE Server, and BUS, see the documentation that came with your switch.

**For detailed
information**

The following sections discuss how to configure the ATM adapter for LANE.

- ◆ “[Verifying that the ATM adapter is installed and functioning](#)” on page 33
- ◆ “[Verifying a working connection to the ATM network](#)” on page 34
- ◆ “[Verifying that the UNI is operational](#)” on page 35
- ◆ “[Configuring the LANE Configuration Server address](#)” on page 37

Verifying that the ATM adapter is installed and functioning

Verifying that the adapter works

To verify that the ATM adapter is functioning, complete the following step.

Step	Action
1	Enter the following command: atm adinfo You can perform the <code>atm adinfo</code> command again at any time to determine the correct unit number.

Example of the atm adinfo command

Sample output from the `adinfo` command is as follows:

```
atm adinfo
a1:unit 1:      PCA-200E Media=OC3-MM-SC HW=2.0.1 FW=4.2.0
Serial=4201600 Slot=1
                MAC=00:20:48:40:1C:80
a2:unit 2:      PCA-200E Media=OC3-MM-SC HW=2.0.1 FW=4.2.0
Serial=4201621 Slot=2
                MAC=00:20:48:40:1C:95
a3:unit 3:      PCA-200E Media=OC3-MM-SC HW=2.0.1 FW=4.2.0
Serial=4201955 Slot=3
                MAC=00:20:48:40:1D:E3
```

Interpreting the output

You should see lines for each ATM adapter in your storage system that is functioning properly. The presence of the lines indicates that the adapter has passed its self-test procedure and that your storage system initialized the adapter. The `adinfo` command also displays the device name for each of the installed adapters, as well as the unit number, at the beginning of each line.

The unit number uniquely identifies the ATM adapter in your storage system, and there is a one-to-one mapping between the device names and unit numbers. The device name consists of the prefix “a” followed by the physical slot number. The unit number is the slot number.

Preparing the ATM adapter for LANE

Verifying a working connection to the ATM network

Verifying that the connection works

To check that the ATM adapter in your storage system is properly connected, complete the following step.

Step	Action
1	Enter the following command: atm adstat -d device <i>device</i> is the name of the ATM adapter whose connection you want to check.

Example of the atm adstat -d command

The following command displays statistics about the ATM adapter in slot 1 of your storage system:

```
atm adstat -d a1
```

Sample output from this command is as follows:

```
Device statistics:
      Buffer Allocation Failures
      Type 1           Type 2
      Small   Large   Small   Large   Receive Queue Full   Carrier
      0         0       0       0       0                   0       ON
```

Interpreting the output

The Carrier column should indicate ON. If it does not, your cabling is incorrectly connected or faulty, or your ATM network is malfunctioning or misconfigured.

Note _____
If you need information about connecting the cabling to your storage system's ATM adapter, see the appropriate section in the hardware guide that came with your storage system.

Preparing the ATM adapter for LANE

Verifying that the UNI is operational

Checking whether the UNI is operational

Your storage system ATM address is automatically registered with the switch; therefore, you use the `uniconfig` command only to display configuration parameters, check the UNI version number, and ensure that the UNI is operational.

To check whether the UNI is operational, complete the following step.

Step	Action
1	Enter the following command: <pre>atm uniconfig show [-unit unit_name]</pre> <i>unit_name</i> is the UNI you want to check. Note_____ If you do not specify the unit number, the UNI information for all ATM adapters in your storage system is displayed. _____

Example of the atm uniconfig show command

Abbreviated sample output from the `atm uniconfig show` command is as follows:

```
atm uniconfig show -unit unit3
UNI parameters for unit3
=====
VPI/VCI                : 0/5
AAL type                : 5
QoS                    : UBR
UNI configured version  : 3.1
UNI operating version   : 3.1
SSCOP operational state : operational
Primary ATM address     :
47.0005.80.ffe100.0000.f21a.4d19.002048401de3.00
UNI failover configuration
```

=====

Status: dynamic

Groups: (0) (1*) (2)

Interpreting the output

The following items should enable you to verify that the ATM interface is operational:

- ◆ The UNI configured version and UNI operating version values should be 3.1.
- ◆ The SSCOP operational state should indicate that the UNI is operational. If you see *inoperational* instead, the ATM card is improperly connected to the network or the switch is improperly configured.
- ◆ The Primary ATM address should be a valid ATM address for your network. If you see an address consisting entirely of zeros, the ATM card is improperly connected to the network or there might be a configuration problem.

Configuring the LANE Configuration Server address

About configuring the LANE Configuration Server address

For your storage system to participate in an Emulated LAN, you must configure the ATM adapter with the address of the LANE Configuration Server. The LANE Client joins an Emulated LAN by first exchanging configuration information with the LANE Configuration Server. The LANE Configuration Server then supplies the client with the ATM address for the LANE Server.

Knowing the LANE Configuration Server address, the system can now determine all existing Emulated LANs, as well as the ATM address of the LANE Server. However, the LAN Type remains unknown until you configure the adapter to join the Emulated LAN, which is discussed in “[Configuring the ATM adapter for an Emulated LAN](#)” on page 40.

Configuring the LANE Configuration Server address

To configure the LANE Configuration Server address for your ATM adapter, complete the following steps.

Step	Action
1	Enter the following command to set the LANE Configuration Server address: <pre>atm elconfig set -lecs <i>ATM_address</i> -wellknown -manual -unit <i>unit_number</i></pre> <i>ATM_address</i> is the ATM address of the LANE Configuration Server on the network. -wellknown indicates that the well-known ATM address will be used. -manual places the host in manual configuration mode; configuration information is not retrieved from the LANE Configuration Server. <i>unit_number</i> is the unit designator for the ATM adapter in your storage system. Note You do not have to specify the unit number if only one ATM adapter is installed in your storage system.

Step	Action
2	Enter the following command to check that the adapter is communicating with the LANE Configuration Server: atm elconfig show -all
3	Study the output. If a separate line fails to appear for each configured Emulated LAN, grouped by adapter, verify that each Emulated LAN is configured properly for the ATM switch. For more information, see your switch vendor's documentation.

Example of the atm elconfig set command with a well-known address

The following command sets the LANE Configuration Server address to the well-known ATM address for the adapter in slot 2:

```
atm elconfig set -lecs -wellknown -unit 2
```

Example of the atm elconfig set command without a well-known address

If the LANE Configuration Server on your network does not use the well-known address, specify the LANE Client Server ATM address in place of `wellknown`, as shown in the following command:

```
atm elconfig set -lecs
47.0079.00.000000.0000.0000.0000.00a03e000001.00 -unit 2
```

Example of the atm elconfig show command

Abbreviated sample output (showing Emulated LANs available through the LANE Configuration Server for Adapter 2 only) from the `atm elconfig show -all` command is as follows:

```
ELANs on Adapter 2
=====
LECS (current): 47.0079.00.000000.0000.0000.0000.00a03e000001.00
      ELAN           LAN Type      LES ATM Address
      ====           =====      =====
      eighteenKMTU   Unknown
47.0005.80.ffe100.0000.f21a.4d19.0020481a4d19.14
```

```
=> default          Unknown
47.0005.80.ffe100.0000.f21a.4d19.0020481a4d19.f0

    nineKMTU        Unknown
47.0005.80.ffe100.0000.f21a.4d19.0020481a4d19.11
```

Interpreting the output

Each Emulated LAN on the network should appear in the output, and the LAN Type and LANE Server ATM addresses should appear as expected. There should be a separate line for each configured Emulated LAN on the network, grouped and arranged for each ATM adapter that is installed in your storage system.

An arrow to the left of an Emulated LAN signifies that the ATM adapter has been configured to operate on that Emulated LAN. For information about configuring adapters to operate on an Emulated LAN, see [“Adding an Emulated LAN to the ATM adapter”](#) on page 41.

The LANE Server ATM address should appear valid. If it does, you know that the ATM adapter in your storage system is communicating properly with the switch. If the LANE Server ATM address is all zeros, it might mean that the cable connection is not working or something else is improperly configured at the switch.

The LANE Configuration Server ATM address should match the address that you specified earlier in [“Configuring the LANE Configuration Server address”](#) on page 37.

Configuring the ATM adapter for an Emulated LAN

About the configuration

You must configure the ATM adapter to enable it to operate on one or more Emulated LANs.

For detailed information

The following sections discuss the actions you take to configure an ATM adapter for an Emulated LAN:

- ◆ [“Adding an Emulated LAN to the ATM adapter”](#) on page 41
- ◆ [“Configuring the logical Ethernet interface”](#) on page 43
- ◆ [“Deleting an Emulated LAN from an ATM adapter”](#) on page 44

Configuring the ATM adapter for an Emulated LAN

Adding an Emulated LAN to the ATM adapter

Prerequisite

Before you configure the ATM adapter to operate on an Emulated LAN, you must configure the LANE Configuration Server ATM address for each adapter, as described in “[Configuring the LANE Configuration Server address](#)” on page 37. Also, the Emulated LAN you specify must already have been configured at the switch.

Adding an Emulated LAN to the adapter

To add the Emulated LAN to the adapter, complete the following steps.

Step	Action
1	Enter the following command: <pre>atm elconfig add ELAN -if interface -les ATM_address -type ethernet -unit unit_number</pre> <i>ELAN</i> is the Emulated LAN that you want the adapter to join. <i>interface</i> is the logical network interface. The <i>-les</i> flag is for joining Emulated LANs whose configuration information is not returned by a LANE Configuration Server. Note You only use the <i>-les</i> flag when the <i>-manual</i> flag is set in the <code>atm elconfig set</code> command. Do not use the <i>-les</i> flag if the LANE Configuration Server address is set to <code>wellknown</code>. <i>ATM_address</i> is the LANE Server ATM address. <i>unit_number</i> is the unit designator for the ATM adapter in your storage system. Note If there is only a single ATM adapter in your storage system, you do not need to specify the unit number in the command. The <code>atm elconfig</code> command sets it automatically.

Step	Action
2	To add more than one Emulated LAN to an adapter, repeat Step 1.

Example of the atm elconfig add command

The following command adds the adapter unit 2 to the nineKMTU Emulated LAN of type Ethernet, using interface el1:

```
atm elconfig add nineKMTU -if el1 -type ethernet -unit 2
```

Interpreting the output

This example assumes that the nineKMTU Emulated LAN already exists on the switch.

The el1 interface refers to a logical interface, thereby enabling you to configure more than one logical interface for the same physical ATM adapter. This means that you can use the `atm elconfig add` command repeatedly to configure your storage system to communicate over multiple Emulated LANs using a single physical ATM adapter. Only Ethernet emulated networks are supported.

Configuring the ATM adapter for an Emulated LAN

Configuring the logical Ethernet interface

Configuring the interface

After the ATM adapter joins an Emulated LAN, you need to assign an IP address to the (logical) network interface and configure additional parameters.

To configure a logical Ethernet interface, complete the following steps.

Step	Action
1	Enter the following command: <pre>ifconfig interface address netmask mask up</pre> <i>interface</i> is the name of the logical network interface. <i>address</i> is the IP address associated with the interface. <i>mask</i> is the network mask that is selected according to the class of the IP address. For more information about the <code>netmask</code> parameter and the <code>ifconfig</code> command, see the <i>Data ONTAP 7.1 Command Reference Guide</i>. Example: The following command configures the <code>e10</code> logical Ethernet network interface, assigning a corresponding IP address and netmask: <pre>ifconfig e10 172.20.12.19 netmask 255.255.252.0 up</pre>
2	To configure more than one logical Ethernet interface, repeat Step 1.

Deleting an Emulated LAN from an ATM adapter

Deleting an Emulated LAN from an adapter

To delete an Emulated LAN from an ATM adapter, complete the following steps.

Step	Action
1	Enter the following command to mark the interface down: <pre>ifconfig interface down</pre> <i>interface</i> is the name of the logical network interface that you want to delete. Example: The following command marks the el1 logical network interface as down: <pre>ifconfig el1 down</pre>
2	Enter the following command: <pre>atm elconfig delete ELAN -unit unit_number</pre> <i>ELAN</i> is the ELAN that you want to delete from the adapter. <i>unit_number</i> is the unit designator for the ATM adapter in your storage system. Example: The following command deletes the nineKMTU Emulated LAN for the adapter designated by unit 2: <pre>atm elconfig delete nineKMTU -unit 2</pre>

Checking and completing the Emulated LAN configuration

About completing the configuration

After you configure the ATM adapter for an Emulated LAN, you should verify your configuration to ensure that it is correct.

For detailed information

The following sections discuss the actions you take to check and complete the Emulated LAN configuration:

- ◆ [“Verifying the communications link”](#) on page 46
- ◆ [“Checking the configuration settings”](#) on page 47
- ◆ [“Checking the other elements of the Emulated LAN”](#) on page 48
- ◆ [“Modifying load balancing and failover”](#) on page 50
- ◆ [“Saving the ATM configuration commands in the /etc/rc file”](#) on page 52
- ◆ [“Saving the host and IP address data in the /etc/hosts file”](#) on page 53

Verifying the communications link

After you add an Emulated LAN to an adapter and configure the interface, you need to check to ensure that your storage system can communicate with other clients through the Emulated LAN. The easiest way to do this is to ping another LANE Client on the Emulated LAN to ensure that information is traveling out through the ATM adapter and back again.

To ping a LANE Client, or any other client, complete the following step.

Step	Action
1	Enter the following command: ping host <i>host</i> is the IP address of the computer to which you want to send an ICMP ECHO_REQUEST datagram.

Example

The following command sends the datagram to host 204.125.14.45, and waits for a response:

```
ping 204.125.14.45
```

If the host responds, ping prints “*host* is alive.” Otherwise, ping resends the ECHO_REQUEST once a second. If the host does not respond after 20 seconds, ping prints the following output:

```
no answer from host.
```

Checking and completing the Emulated LAN configuration

Checking the configuration settings

Verifying adapter configurations

After you verify the communication link, you should check the state of the adapters in your storage system to ensure that the configuration is correct.

To check the configuration settings, complete the following steps.

Step	Action
1	Enter the following command: atm elconfig show -all
2	Verify there is an arrow preceding the Emulated LAN name, which indicates that the adapter has joined the Emulated LAN.
3	Check that the LAN Type is Ethernet.
4	Check that the LANE Server ATM address is a valid ATM address. If the address is all zeros, it indicates a configuration problem at the switch.

Example of the atm elconfig show -all command

Abbreviated sample output (showing Emulated LANs on Adapter 2 only) from this command is as follows:

```
atm elconfig show -all
ELANs on Adapter 2
=====
LECS (current): 47.0079.00.000000.0000.0000.0000.00a03e000001.00
      ELAN           LAN Type      LES ATM Address
      ====           =====      =====
      eighteenKMTU   Ethernet
47.0005.80.ffe100.0000.f21a.4d19.0020481a4d19.14
=> default          Unknown
47.0005.80.ffe100.0000.f21a.4d19.0020481a4d19.f0
=> nineKMTU         Ethernet
47.0005.80.ffe100.0000.f21a.4d19.0020481a4d19.11
```

Checking and completing the Emulated LAN configuration

Checking the other elements of the Emulated LAN

Checking the other elements You should also check the other elements of the Emulated LAN to ensure they are configured and operating correctly.

To check the other elements of the Emulated LAN, complete the following steps.

Step	Action
1	Enter the following command: atm elconfig show -configured
2	Check the output to verify that the LANE Server, LANE Configuration Server, and BUS ATM addresses are all valid, and that the state of the Emulated LAN is operational.
3	Use the <code>netstat -i</code> command to check the MTU of each Emulated LAN to verify that the setting matches the configuration created on the switch.

Example of the atm elconfig show -configured command Abbreviated sample output (showing information related to the eighteenKMTU Emulated LAN on adapter 2 only) from the `elconfig show -configured` command is as follows:

```
atm elconfig show -configured
ELAN Name                : eighteenKMTU
Interface                : el1
Configured Unit          : 2
MAC Address              : 00:20:48:08:12:c3
LEC Address              :
47.0005.80.ffe100.0000.f20f.6d4c.0020480812c3.00
LECS Address             :
c5.0079.00.000000.00000000000000a03e000001.00
Configuration Direct VCC : unit=2 vpi/vci=0/279
LES Address              :
47.0005.80.ffe100.0000.f21a.4d19.0020481a4d19.f0
```

```
Control Direct VCC      : unit=2 vpi/vci=0/280
Control Distribute VCC  : unit=2 vpi/vci=0/281
BUS Address             :
47.0005.80.ffe100.0000.f21a.4d19.0020481a4d19.f0
Multicast Send VCC      : unit=2 vpi/vci=0/282
Multicast Forward VCC   : unit=2 vpi/vci=0/283
LAN Type                : Ethernet/IEEE 802.3
Maximum Frame Size      : 1516
State                   : operational
MPOA                    : disabled
LECID                   : 13
```

Modifying load balancing and failover

About load balancing

Load balancing enables incoming and outgoing traffic to be spread across ATM adapters in a group.

By default, all ATM adapters are dynamically assigned to load-balancing groups. No configuration is necessary to activate the load balancing and failover features. However, you can override the default by changing the adapters to static mode so groups can be manually configured—for instance, when a failover group contains adapters on different switches, resulting in reduced network performance.

Requirements for load balancing

Load balancing does not depend on any nonstandard extensions to the UNI. However, the switch must support the following:

- ◆ Registering the same ATM address on multiple ports
- ◆ Registering multiple ATM addresses on a single port

If the switch does not support these features, load balancing and failover are automatically disabled.

Modifying load-balancing groups

To modify load-balancing groups, complete the following step.

Step	Action
1	Enter the following command: <pre>atm uniconfig set failover [-state off static dynamic] [-group unit ...]</pre> -state dynamic is the default. Use -state off to disable load balancing and failover -state static to put UNI load balancing in a static mode unit specifies the load-balancing group membership. Note _____ The parameters to the -group option of the atm uniconfig set failover command specify the ATM adapters (units) that should be logically assigned to a load-balancing and failover group. If you specify a unit that has already been assigned to another group, the unit is automatically removed from the original group before being assigned to the new group.

Example of the atm uniconfig set failover command

The following example demonstrates how you disable load balancing and failover:

```
atm uniconfig set failover -state off
```

Saving the ATM configuration commands in the /etc/rc file

Saving the configuration commands

By saving the ATM configuration information in the /etc/rc file, you avoid having to reconfigure the adapters manually each time your storage system is restarted.

To save the configuration commands in the /etc/rc file for automatic execution at boot time, complete the following steps.

Step	Action
1	Mount the root file system and add the configuration commands to the /etc/rc file using a text editor, such as vi.
2	Save your changes.

Sample /etc/rc file with ATM configuration commands

Following is a sample portion of an /etc/rc file containing configuration commands for three ATM adapters:

```
# unit 1
elconfig set -lecs -wellknown -unit 1
elconfig add default -if el0 -type ethernet -unit 1
ifconfig el0 172.20.12.19 netmask 255.255.252.0 up
# unit 2
elconfig set -lecs -wellknown -unit 2
elconfig add nineKMTU -if el1 -type ethernet -unit 2
ifconfig el1 201.201.201.219 netmask 255.255.255.0 up
# unit 3
elconfig set -lecs -wellknown -unit 3
elconfig add eighteenKMTU -if el2 -type ethernet -unit 3
ifconfig el2 201.201.210.219 netmask 255.255.255.0 up
elconfig add nineKMTU -if el3 -type ethernet -unit 3
ifconfig el3 201.201.210.220 netmask 255.255.255.0 up
elconfig wait
```

Saving the host and IP address data in the /etc/hosts file

Saving the host and IP address

To save the host and IP address information of the Emulated LAN configuration in the /etc/hosts file, complete the following steps.

Step	Action
1	Mount the root file system and add the host and IP address to the /etc/hosts file using a text editor, such as vi.
2	Save your changes.

Sample /etc/hosts file entry

Following is a sample portion of the /etc/hosts file containing the host and IP address information for the ATM adapters:

```
172.20.12.19      myfiler-el0
201.201.201.219   myfiler-el1
201.201.210.219   myfiler-el2
201.201.210.220   myfiler-el3
```

Additional information for storage systems with only ATM adapters

If you have a storage system that has only ATM adapters, the /etc/hosts file must contain an entry for your storage system’s host name.

The host name is not displayed as part of the command prompt until you add the host name in one IP entry and reboot your storage system. On storage systems that include other types of network interfaces, the installation setup procedure automatically adds the host name entry to the /etc/hosts file.

Example

Following is the first line in a sample /etc/hosts file:

```
172.20.12.19      myfiler  myfiler-el0
```

Understanding FORE/IP over SPANS

Version supported	Data ONTAP currently supports only FORE/IP 5.3. For more information about the differences between FORE/IP 5.3 and older versions, see the FORE documentation.
Types of SPANS connections ATM supports	ATM supports two types of SPANS connections: ◆ SVC ◆ PVC
When SVCs get assigned	Data ONTAP dynamically assigns SVCs when interoperating with ATM hosts and with switches that support the FORE/IP SPANS protocols.
When to use PVCs	You use ATM PVCs to interoperate with ATM hosts and with switches that do not support FORE/IP SPANS. For example, if you are not using a FORE systems switch, PVCs can connect FORE equipment at each end through non-FORE switches.
How FORE/IP interfaces allow communication	For each physical ATM interface, Data ONTAP creates a FORE/IP interface, called fa, at boot time. The fa interface supports FORE/IP on top of SPANS signaling. FORE/IP allows communication as follows: ◆ Using AAL4 or AAL5 ATM adaptation layer types with no encapsulation ◆ Using a broadcast Address Resolution Protocol (ARP) for SPANS address resolution ◆ Using direct communication of all hosts on a physical ATM network without the use of IP routers

Note

Data ONTAP does not support FORE/IP load balancing or failover options.

Managing FORE/IP and PVCs

About establishing and deleting PVCs

PVCs are static; for each destination, you must establish (attach the IP layer to) a PVC explicitly and delete (detach the IP layer from) the PVC explicitly.

For each destination that needs to establish a PVC with your storage system, you must establish an outgoing PVC and an incoming PVC in three places:

- ◆ On your storage system
- ◆ On the destination ATM host
- ◆ On all interconnecting ATM switches

For detailed information

The following sections describe the actions involved in managing FORE/IP PVCs:

- ◆ [“Establishing FORE/IP PVCs on your storage system”](#) on page 56
- ◆ [“Displaying information about a FORE/IP PVC”](#) on page 58
- ◆ [“Displaying the FORE/IP configuration”](#) on page 59
- ◆ [“Changing the ATM adaptation layer for FORE/IP and SPANS”](#) on page 61
- ◆ [“Deleting a FORE/IP PVC”](#) on page 62

What this section does not discuss

The following topics are not discussed in this section:

- ◆ Establishing a FORE/IP PVC on the remote ATM host
Set up a FORE/IP PVC on the remote ATM host according to the documentation for that host.
- ◆ Establishing a FORE/IP PVC on interconnecting ATM switches
On the interconnecting ATM switches, assign virtual channels corresponding to the virtual path identifier (VPI) and virtual channel identifier (VCI) entries made on your storage system and the remote ATM host according to the documentation for those switches.

Establishing FORE/IP PVCs on your storage system

Process for establishing FORE/IP PVCs

The process for establishing FORE/IP PVCs on your storage system includes the following tasks:

- ◆ Establishing an outgoing FORE/IP PVC
- ◆ Establishing an incoming FORE/IP PVC

FORE/IP PVC configuration variables

When establishing an outgoing or incoming FORE/IP PVC, replace the following variables with their respective values in the command line.

Variable	Description
<i>hostname</i>	Name or IP address of the remote host.
<i>iface</i>	Name of the ATM interface. This is usually <i>fan</i> , where <i>n</i> is a number.
<i>vpi</i>	VPI (virtual path identifier); this must be 0.
<i>vci</i>	VCI (virtual channel identifier); this number must have the following properties: ◆ It must not be in use on your storage system. ◆ It must be less than 1,024. ◆ It must obey the limits of the destination host and interconnecting devices.
<i>aal</i>	ATM adaptation layer (AAL) type. It must be 4 or 5, and should be the AAL type supported by the destination host, which is typically 5. The default is 5. Note —— AAL4 is not supported on ForeRunner HE622 (OC-12) adapters.

Variable	Description
<i>encap</i>	Encapsulation type. Specify one of the following encapsulation types: ◆ null (no encapsulation; this is the default) ◆ llc_routed (IEEE LLC encapsulation for routed protocol data units [PDUs]) ◆ llc_bridged_8023 (IEEE LLC encapsulation for Ethernet/802.3 bridged PDUs) It should be the same as the encapsulation type used by the destination host. If the encapsulation type is llc_bridged_8023, you must include <i>addr</i>.
<i>addr</i>	Six-byte colon-separated destination MAC address.

Establishing an outgoing FORE/IP PVC on your storage system

To establish an outgoing FORE/IP PVC on your storage system, complete the following step.

Step	Action
1	Enter the following command: <pre>atm atmarp -s hostname pvc iface vpi vci [aal [encap [addr]]]</pre>

Establishing an incoming FORE/IP PVC on your storage system

To establish an incoming FORE/IP PVC on your storage system, complete the following step.

Step	Action
1	Enter the following command: <pre>atm atmarp -l pvc iface vpi vci [aal [encap]]</pre>

Displaying information about a FORE/IP PVC

Displaying information about FORE/IP PVCs

Data ONTAP enables you to display address resolution information for incoming and outgoing PVCs so that you can verify the current settings.

To display information about all FORE/IP PVCs and other interfaces on a host, complete the following step.

Step	Action
1	Enter the following command: <pre>atm atmarp [hostname -a]</pre> <i>hostname</i> is the name or IP address of a specific remote host; use the <i>-a</i> flag to display information about all the current FORE/IP PVCs.

Example

If you use the *-a* flag, a display similar to the following appears:

```
atm atmarp -a

iface=a5 switch.port=f21a2420.56 vpi.vci=0.114 aal=5
encapsulation=NULL
iface=a5 switch.port=f21a2420.25 vpi.vci=0.113 aal=5
encapsulation=NULL
```

Managing FORE/IP and PVCs

Displaying the FORE/IP configuration

**FORE/IP
information
displayed**

You can display the FORE/IP configuration information to verify the current ATM adapter settings. The following types of information are displayed:

- ◆ Fore/IP parameters
- ◆ Connectionless VC parameters
- ◆ SPANS signaling VC parameters

**Displaying FORE/IP
configuration
information**

To display the current FORE/IP configuration information on an ATM adapter, complete the following step.

Step	Action
1	Enter the following command: atm atmconfig device <i>device</i> is the ATM adapter name.

**Sample atm
atmconfig
command output**

The following is sample output from the `atm atmconfig` command:

```
atm atmconfig fa0
FORE IP parameters for fa0
=====
MTU:                9188
SVC peak rate:      (unlimited)

Connectionless VC parameters
=====
VPI/VCI:            0/14
AAL:                5
peak rate:          (unlimited)
```

```
SPANS signaling VC parameters
=====
VPI/VCI:      0/15
AAL:          5
peak rate:    (unlimited)
```

Changing the ATM adaptation layer for FORE/IP and SPANS

When to change the AAL

You can change the FORE/IP ATM adaptation layer, for instance, when you install an OC-12 adapter and you need to change the AAL from 4 to 5.

See “[FORE/IP PVC configuration variables](#)” on page 56 for a description of command variables.

Changing the FORE/IP AAL

To change the FORE/IP AAL, complete the following step.

Step	Action
1	Enter the following command: <pre>atm atmconfig -c vpi vci aal device</pre> <i>device</i> is the ATM adapter name; use the -c switch to display information for FORE/IP. Example: <code>atm atmconfig -c 0 14 5 fa3</code>

Changing the SPANS AAL

To change the SPANS AAL, complete the following step.

Step	Action
1	Enter the following command: <pre>atm atmconfig -s vpi vci aal device</pre> <i>device</i> is the ATM adapter name; use the -s switch to display information for SPANS. Example: <code>atm atmconfig -s 0 15 5 fa3</code>

Deleting a FORE/IP PVC

Deleting an outgoing FORE/IP PVC

To delete an outgoing FORE/IP PVC entry, complete the following step.

Step	Action
1	Enter the following command: atm atmarp -d <i>hostname</i> <i>hostname</i> is the name or IP address of the remote host.

Deleting an incoming FORE/IP PVC

To delete an incoming FORE/IP PVC for a remote host, complete the following step.

Step	Action
1	Enter the following command: atm atmarp -x <i>iface vpi vci</i> <i>iface</i> is the name of the interface. <i>vpi</i> and <i>vci</i> are the values of the VPI and VCI of the FORE/IP PVC to be deleted for the specific interface.

About this chapter

This chapter discusses how Data ONTAP manages routing, how it handles different types of packet requests, and how you can modify the routing table.

Topics in this chapter

This chapter discusses the following topics:

- ◆ [“About routing in Data ONTAP”](#) on page 64
- ◆ [“Enabling and disabling routing mechanisms”](#) on page 70
- ◆ [“Displaying the routing table and default route information”](#) on page 72
- ◆ [“Modifying the routing table”](#) on page 75

About routing in Data ONTAP

About Data ONTAP routing

Although your storage system can have multiple network interfaces, it does not function as a router. The Data ONTAP software does not route packets between the interfaces of your storage system on behalf of other network hosts; however, Data ONTAP can route its own outbound packets.

Data ONTAP uses two routing mechanisms:

- ◆ fast path
To route Network File System (NFS) packets over User Datagram Protocol (UDP) and to route all TCP traffic, Data ONTAP uses a mechanism called fast path. See “[About fast path](#)” on page 65.
- ◆ routing table
To route all other IP traffic, Data ONTAP uses the information available in the local routing table. See “[About the routing table](#)” on page 67.

About fast path

What fast path is

Fast path is an alternate routing mechanism available in Data ONTAP. Instead of using the routing table of your storage system to route, this mechanism uses

- ◆ The source Media Access Control (MAC) address of the incoming packet as the destination MAC address of the outgoing packet for NFS-over-UDP and all TCP traffic transmitted from your storage system
- ◆ The same interface for incoming and outgoing traffic

Using this mechanism provides the following advantages:

- ◆ Load balancing between multiple storage system interfaces on the same subnet

The load balancing is achieved by sending responses on the same interface of your storage system as incoming requests.

- ◆ Increasing storage system performance

The increase in storage system performance is achieved by skipping routing table lookups.

Fast path is enabled automatically on your storage system; however, you can disable it.

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

TCP: Because TCP is connection-oriented and because data is acknowledged as part of the TCP protocol, Data ONTAP can use fast path on every TCP packet transmitted except the very first SYN packet (if Data ONTAP initiates a connection). For fast path, the interface used to transmit a packet is the same interface the last packet was received on.

For TCP connections, Data ONTAP automatically turns off fast path if it detects that using fast path in a network setup is not optimal.

**Effect of fast path:
Telnet works, but
ping fails**

If fast path is enabled and the default router stops working, you might notice that Telnet sessions to your storage system can still be established from a non-local subnet, even though you cannot use the `ping` utility to communicate with your storage system's interfaces. This happens because the fast path mechanism skips routing table lookups. As a result, the routing table is not used to respond to any NFS-over-UDP or TCP connection requests. Therefore, Telnet requests (which use the TCP protocol) succeed, while ping requests—which use the Internet Control Message Protocol (ICMP)—fail.

**Effect of fast path
on 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 forwarding packets to your storage system is not the one forwarding the packets that the storage system sends back. In this case, you must disable fast path.

About the routing table

What the routing table contains

The routing table contains the current routes that have been established and are currently in use, as well as the default route specification.

Default route setup in Data ONTAP

Data ONTAP uses a default route entry to route to destinations that it does not explicitly know about in its routing table. You can set the default route in Data ONTAP either during the initial setup or later by modifying the `/etc/rc` file.

If you are upgrading your storage system to this Data ONTAP release and currently use the `/etc/dgateways` file to set a default route, you should now use the `/etc/rc` file, router discovery, or Routing Information Protocol (RIP) instead. The `/etc/dgateways` file was deprecated in Data ONTAP 6.0 (that is, it is still supported for backward compatibility but its use is not recommended).

Example: The following sample `/etc/rc` file shows the `route add` command used to add a default route:

```
hostname tpubs-f720
ifconfig e0 172.28.50.21 netmask 255.255.255.0 mediatype 100tx-fd
route add default 172.28.50.1 1
routed on
```

Managing the routing table

You can manage the routing table in two ways:

- ◆ Automatically, using the `routed` daemon
The `routed` daemon is enabled by default.
- ◆ Manually, using the `route` command

The routing table might also be modified when one of the following occurs:

- ◆ A new interface is configured with the `ifconfig` command and there are no existing entries for the new network number in the routing table.
- ◆ Your storage system receives an ICMP redirect packet with notification of a better first-hop router for a particular destination.
- ◆ Your storage system is rebooted after the default route in the `/etc/rc` file is modified.

What the `routed` daemon provides

The `routed` daemon enables these functions:

- ◆ Deletion of redirected routes after a specified period
- ◆ Router discovery with Internet Router Discovery Protocol (IRDP), which is useful only if there is no static default route
- ◆ Listening for RIP packets
- ◆ Migration of routes to alternate interfaces when multiple interfaces are available on the same subnet

In addition, `routed` can be configured to

- ◆ 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

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

When the `routed` daemon can be turned off

In some circumstances, it might be desirable to turn the `routed` daemon off. For example, if you have multiple interfaces on the same subnet and you want to direct network traffic to specific interfaces, you must turn `routed` off because `routed` sees all interfaces on a subnet as equivalent.

You can safely turn off `routed` if you

- ◆ Do not use RIP or router discovery (they can be disabled by setting values in the `/etc/gateways` file)
- ◆ Have a single router per subnet or a network in which redirects are not sent
- ◆ Are able to manage your routing table directly

Note

Unless you have specific routing needs and you understand network routing configuration, you are advised to always keep `routed` on, even if you do not want Data ONTAP to make routing decisions based on routing updates. Turning `routed` off could cause unexpected routing behavior in Data ONTAP.

Routing tables in a vFiler unit environment

If you enable the MultiStore® license, Data ONTAP disables the `routed` daemon. Therefore, routing tables in a vFiler™ unit environment must be managed manually with the `route` command.

All vFiler units in an IPspace (the IP address space in which vFiler units can function) share a routing table. Therefore, any commands that display or manipulate the routing table apply to all vFiler units in that IPspace.

For more information, see the section on network considerations in the Data ONTAP 7.1 *MultiStore Management Guide*.

Enabling and disabling routing mechanisms

Controlling routing Both the fast path mechanism and the `routed` daemon are enabled by default in Data ONTAP. To enable or disable these routing mechanisms, use the command line or FilerView methods described below.

Note If you disable both fast path and `routed`, you must be prepared to configure routing manually; see “[About routing in Data ONTAP](#)” on page 64.

Turning fast path on or off To turn fast path on or off, complete the following step. (You cannot turn fast path on or off in FilerView.)

Step	Action
1	Enter the following command at your storage system command line: <code>options ip.fastpath.enable {on off}</code>

Note You can use the `-x` option with the `netstat` command to see if the fast path mechanism is enabled for a specific connection.

Turning routed on or off at the command line To turn the `routed` daemon on or off, complete the following step.

Step	Action
1	Enter the following command at your storage system command line: <code>routed {on off}</code>

Note If you use the command-line method, you must also edit the `/etc/rc` file in the root volume to specify the same `routed` daemon behavior across storage system reboots.

Turning routed on or off with FilerView

To turn the `routed` daemon on or off with FilerView, complete the following steps.

Step	Action
1	In FilerView, click Network in the list on the left.
2	In the list under Network, click Configure.
3	Select Yes (for on) or No (for off) in the Routed Enabled drop-down list, then click Apply.

Note

If you make changes to `routed` configuration in FilerView, the changes are saved automatically in the `/etc/rc` file and therefore become persistent across reboots.

Displaying the routing table and default route information

Displaying the routing table at the command line

To display the Data ONTAP routing table at the command line, complete the following step.

Step	Action
1	Enter the following command at your storage system command line: netstat -rn

Displaying default route information at the command line

To display information about whether `routed` is on or off, default route information, and routing protocols at the command line, complete the following step.

Step	Action
1	Enter the following command at your storage system command line: routed status

Displaying routing information with FilerView

To display the routing table, the default route information, and routing protocols using FilerView, complete the following steps.

Step	Action
1	In FilerView, click Network in the list on the left.
2	In the list under Network, click Report. The Routing section of the Network Report shows the default route and protocols in effect, and then shows routing tables.

Example of a routing table

The following example shows a routing table output as a response to the command:

```
netstat -rn

Internet:
Destination      Gateway           Flags      Refs      Use  Interface
default          172.28.50.1      UGS        5         5860  e0
127.0.0.1        127.0.0.1        UH         1          262  lo
172.28.100/24     link#1           UC         0           0  e0
172.28.50.1      0:e0:52:1:dd:66  UHL        1           0  e0
172.28.50.3      8:0:20:9b:37:e6  UHL        0           4  e0
172.28.50.18     8:0:20:94:1c:ce  UHL        0           0  e0
172.28.50.255    ff:ff:ff:ff:ff:ff UHL        0        3903  e0
172.28.255.255   ff:ff:ff:ff:ff:ff UHL        1        1733  e0
```

In the previous example, the destination can be a host, 172.28.50.1, a network, 172.28.100/24, 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
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
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

Flag	Description
S	Static—Route added with the <code>route</code> command
2	Proxy ARP—Host is configured to respond to ARP requests for a host other than itself

For more information about the routing table display, see the `na_netstat(1)` man page.

Modifying the routing table

About the route command

The routing table can be managed directly using the `route` command. The command enables you to

- ◆ Add and delete routes or modify existing ones
- ◆ Remove all gateways in the routing table

You can also list routes with the `route -s` command, which yields the same output as `netstat -rn`.

Note_____

You cannot modify the routing table using FilerView.

Modifying the routing table

To modify the routing table, complete the following step.

Step	Action
1	Enter the following command: <pre>route [add delete] [inet inet6 prefixlen length] [host net] destination [gateway metric]</pre> For more information about the <code>route</code> command and options, see the <code>na_route(1)</code> man page.

Modifying the routing table in a cluster environment

As in other aspects of cluster management, the routing tables of clustered storage system partners must be synchronized.

In takeover mode, each storage system in a cluster retains its own routing table. You can make changes to the routing table on the active storage system in the standard way, or you can make changes to the routing table on the failed storage system using the `route` command in partner mode. However, the changes you make in partner mode are lost after a giveback.

About this chapter

This chapter discusses how you can use the Data ONTAP configuration files, Domain Name System (DNS), and Network Information Service (NIS) to resolve host names.

Topics in this chapter

This chapter discusses the following topics:

- ◆ [“Maintenance of host information”](#) on page 78
- ◆ [“Using the /etc/hosts file to maintain host information”](#) on page 79
- ◆ [“Using DNS to maintain host information”](#) on page 82
- ◆ [“Using dynamic DNS to update host information”](#) on page 89
- ◆ [“Using NIS to maintain host information”](#) on page 92
- ◆ [“Changing the host name search order”](#) on page 100

Maintenance of host information

Ways to maintain host information

Host 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
For detailed information, see [“Using the `/etc/hosts` file to maintain host information](#)” on page 79.
- ◆ On a Domain Name System (DNS) server
For detailed information, [“Using DNS to maintain host information](#)” on page 82.
- ◆ On a Network Information Service (NIS) server
For detailed information, see [“Using NIS to maintain host information](#)” on page 92.

Search order for host information

If you use more than one of the above ways to maintain host information, the ways are used in the order determined by the `/etc/nsswitch.conf` file. For detailed information about this file, see [“Changing the host name search order](#)” on page 100.

The role of host-name resolution in Data ONTAP

Data ONTAP relies on correct host-name resolution to provide basic connectivity for storage systems on the network, including

- ◆ Processing NFS mount requests
- ◆ Establishing CIFS sessions
- ◆ Authenticating Remote Shell (RSH) protocol sessions to storage systems

If you are unable to access storage system data or establish sessions, there might be problems with host-name resolution on your storage system or on a name server.

Using the `/etc/hosts` file to maintain host information

About the `/etc/hosts` file

Data ONTAP uses the `/etc/hosts` file to resolve host names to IP addresses, including host names used in any of the following files:

- ◆ `/etc/rc`
- ◆ `/etc/syslog.conf`
- ◆ `/etc/exports`
- ◆ `/etc/netgroup`
- ◆ `/etc/hosts.equiv`

You must ensure that the `/etc/hosts` file is kept up-to-date. If you update the file, you do not need to reboot your storage system—the changes to the 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

Note

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

For more information on file format, see the `na_hosts(5)` man page.

Ways to add entries to the `/etc/hosts` file

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

- ◆ Locally

You might want to add entries to the local `/etc/hosts` file if the number of entries is small. You can do so in the following ways:

- ❖ At the command line

See “[Editing the `/etc/hosts` file manually](#)” on page 80.

- ❖ Using FilerView

See “[Editing the `/etc/hosts` file with FilerView](#)” on page 81.

- ◆ Remotely using the NIS makefile master

If the number of entries is large and you have access to an NIS makefile master, you might want to use the makefile master to create the `/etc/hosts` file. This method prevents errors that could be introduced in the manual creation process. For details, see [“Creating /etc/hosts from the NIS master”](#) on page 81.

Note

Using NIS to distribute the `/etc/hosts` file is different from looking up host names on an NIS server. For more information about network lookups, see [“Using NIS to maintain host information”](#) on page 92.

Editing the `/etc/hosts` file manually

To edit the `/etc/hosts` file manually, complete the following steps.

Step	Action
1	From a workstation that has access to your storage system’s root volume, open the <code>/etc/hosts</code> file using a text editor.
2	Edit the file to your needs. The format of the file is as follows: <i>IP address Host-name aliases</i>
3	Save the file.

Example: The following shows how the entries might look in the `/etc/hosts` file on a storage system:

```
192.16.3.145    toaster    toaster-e0
192.16.4.155    toaster-e2
192.16.5.165    toaster-e4
192.16.6.175    toaster-e8
```

In the first line, your storage system’s host name itself is used as an alias for the first network interface. That is, network traffic addressed to `toaster` will be received on the `toaster-e0` interface.

Editing the /etc/hosts file with FilerView

To edit the /etc/hosts file with FilerView, complete the following steps.

Step	Action
1	In FilerView, click Network in the list on the left.
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 wish to add and click OK.
5	Click Apply in the Manage Hosts File window.

Creating /etc/hosts from the NIS master

To modify the makefile for the NIS master to create a hosts file and copy it to the /etc directory on your storage system's default volume, complete the following steps.

Step	Action
1	On the NIS server, open the NIS makefile with an editor.
2	Locate the section for hosts.time.
3	Add the following lines at the end of the hosts.time section, replacing <i>dirname</i> with a directory name of your choice, and <i>toaster 1</i>, <i>toaster2</i>, and so on with names of your storage systems: <pre>@mntdir=/tmp/<i>dirname_etc_mnt</i>_\$\$\$\$;\nif [! -d \$\$mntdir]; then rm -f \$\$mntdir; \nmkdir \$\$mntdir; fi;\nfor s_system in <i>toaster1 toaster2 toaster3</i> ; do \nmount \$\$s_system:/etc \$\$mntdir;\nmv \$\$mntdir/hosts \$\$mntdir/hosts.bak;\ncp /etc/hosts \$\$mntdir/hosts;\numount \$\$mntdir;\ndone;\nrmdir \$\$mntdir</pre>
4	Save the NIS makefile. The /etc/hosts file on your storage system is updated whenever the NIS makefile is run.

Using DNS to maintain host information

Advantage of using DNS

DNS enables you to maintain host information centrally. As a result, you do not have to update the `/etc/hosts` file every time you add a new host to the network. 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.

A conventional storage system policy for efficient host-name resolution is to do both of the following:

- ◆ Maintain a short `/etc/hosts` file containing local interfaces, as described in [“Ways to add entries to the `/etc/hosts` file”](#) on page 79.
- ◆ Enable DNS with DNS caching, as described in [“About configuring DNS”](#) on page 82 and [“What DNS name caching does”](#) on page 86.

About configuring DNS

You can configure your storage system to use one or more DNS servers either during the `setup` procedure or later using the command line or FilerView.

If you configure DNS during the `setup` procedure, your storage system’s DNS domain name and name server addresses are configured

- ◆ Automatically if you use Dynamic Host Configuration Protocol (DHCP) to configure onboard interfaces
- ◆ Manually if you do not use DHCP—you must enter the values when prompted

If you configure DNS later, you need to take these 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 these ways:

- ◆ Using FilerView
See [“Configuring DNS with FilerView”](#) on page 83.

- ◆ At the command line
See the appropriate instructions:
 - ❖ “[Creating or editing /etc/resolv.conf](#)” on page 85
 - ❖ “[Specifying the DNS domain name](#)” on page 85
 - ❖ “[Disabling or enabling DNS](#)” on page 86

If you want to use primarily DNS for host-name resolution, specify it ahead of other methods in the hosts map in the `/etc/nsswitch.conf` file. For information about how to edit the `nsswitch.conf` file, see “[Changing the host name search order](#)” on page 100.

Correct host-name resolution depends on the correct configuration of 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(8)` man page.

Configuring DNS with FilerView

To set or modify DNS configuration values with FilerView, complete the following steps.

Step	Action
1	In FilerView, click Network in the list on the left.
2	In the list under Network, click Manage DNS and NIS Name Service.

Step	Action	
3	If you want to....	Then...
	Enable DNS	Select Yes in the DNS Enabled field.
	Set or modify the DNS domain name	Enter a name in the DNS Domain Name field. Examples of configuration values are listed in “ Specifying the DNS domain name ” on page 85.
	Specify or modify DNS servers	Enter up to three IP addresses in the DNS Servers fields. Examples of configuration values are listed in “ Creating or editing /etc/resolv.conf ” on page 85.
	Specify or modify the search list for host name lookup	Enter a name in the DNS Domain Search field.

Creating or editing /etc/resolv.conf

To create or edit the /etc/resolv.conf file, complete the following step.

Step	Action	
1	If...	Then...
	You are creating the /etc/resolv.conf file	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: nameserver <i>ip_address</i> Example: nameserver 192.9.200.10 nameserver 192.9.200.20 nameserver 192.9.200.30
	You are editing the /etc/resolv.conf file	From a workstation that has access to your storage system's root volume, edit the /etc/resolv.conf file using a text editor.

You can optionally set or modify the domain search list for DNS host name lookup. For more information, see the `na_resolv.conf(5)` man page.

Specifying the DNS domain name

To specify or change the DNS domain name, complete the following step at your storage system command line.

Step	Action
1	Enter the following command: options dns.domainname <i>domain</i> <i>domain</i> is the new domain name, which follows your storage system's host name in the fully qualified domain name. For example, the domain name of the storage system <code>system1.company.com</code> is <code>company.com</code> .

Disabling or enabling DNS

To disable or enable DNS, complete the following step at your storage system command line.

Step	Action
1	Enter the following command: <code>options dns.enable {off on}</code> Use <code>off</code> to disable DNS or <code>on</code> to enable DNS.

If you did not configure DNS during the Data ONTAP `setup` procedure, DNS is disabled by default.

Once enabled, DNS should be disabled only when you change host-name resolution procedures or when you troubleshoot problems with the DNS name server or Windows Active Directory server.

Note _____
Your storage system’s CIFS implementation depends on DNS to provide the Windows Active Directory service. Therefore, disabling DNS might interrupt CIFS services.

What DNS name caching does

DNS name caching enables the DNS name resolver to speed up the process by which it converts host names into IP addresses. DNS name caching stores DNS requests by caching them so that they are easy to find the next time. Name caching improves DNS performance in the case of name server failure as well as reducing the time it takes for cluster takeover and giveback.

DNS name caching is enabled by default.

Disabling or enabling DNS name caching

To disable or enable DNS name caching, complete the following step at your storage system command line.

Attention _____
Disabling DNS name caching clears the DNS name cache.

Step	Action
1	Enter the following command: <code>options dns.cache.enable {on off}</code> Use <code>on</code> to enable DNS name caching or <code>off</code> to disable DNS name caching.

Flushing the DNS cache

Entries in the DNS cache have a set expiration. If an entry that has expired is needed again, your storage system contacts the DNS server to get an updated entry. However, if a DNS entry changes before it has expired, you must flush the DNS cache to force the storage system to get the new DNS record.

If some of your DNS records change often, you should make sure that your DNS server transmits them with a low Time To Live (TTL). (You set the TTL in the DNS server.) You can also disable DNS caching on your storage system with the `dns.cache.enable` option, but doing so might reduce performance.

To flush the DNS cache, complete the following step.

Step	Action
1	Enter the following command: <code>dns flush</code>

Displaying DNS information

You can display the following types of DNS information:

- ◆ Status of the DNS resolver
- ◆ List of DNS servers configured in the `/etc/resolv.conf` file
- ◆ State of each DNS server
- ◆ Timestamp when the DNS server was last polled
- ◆ Average round-trip time of a DNS query
- ◆ Total number of DNS queries made

- ◆ Number of failed DNS queries
- ◆ Default domain configured on your storage system
- ◆ List of other domains that will be used with unqualified names for name lookup

To display DNS information, complete the following step.

Step	Action
1	Enter the following command: dns info

For more information about the `dns info` display, see the `na_dns(1)` man page.

Using dynamic DNS to update host information

About dynamic DNS updates

Dynamic DNS updates enable your storage system to send new or changed DNS information to the primary master DNS server for your storage system's zone.

Need for dynamic DNS updates

Without dynamic DNS updates, system administrators have to 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 not only slow, but also error-prone.

Additionally, in a disaster-recovery situation when a storage system with a large number of vFiler units is brought online, manual configuration of DNS information for those vFiler units can result in a longer-than-needed downtime.

By enabling dynamic DNS updates on your storage system, you allow your storage system to automatically send information to the DNS servers as soon as the information changes on the system.

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. StorageSystem1 automatically sends updated information to primary master DNS server for StorageSystem1.

How dynamic DNS updates work in Data ONTAP

If dynamic DNS updates are enabled on your storage system, it periodically sends updates to the primary master DNS server for its zone. Your storage system finds out the primary master DNS server for its zone by querying the DNS servers configured in 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.

In addition to periodic updates, DNS updates are also sent if any DNS information changes on your storage system.

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 using dynamic DNS updates in Data ONTAP, the following conditions apply:

- ◆ 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:
 - ❖ Vif and VLAN interfaces
 - ❖ vFiler units
- ◆ You cannot set TTL values for individual vFiler units. All vFiler units inherit the TTL value 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 dynamic DNS updates

To enable your storage system to send dynamic DNS updates automatically, complete the following step on your storage system.

Step	Action
1	Enter the following command: options dns.update.enable [off on secure] off—Disable dynamic DNS updates on—Enable dynamic DNS updates secure—Enable secure dynamic DNS updates Note— Secure dynamic DNS updates are supported for Windows DNS servers only.

Changing the time-to-live setting for DNS entries

To change the TTL for the DNS entries, complete the following step.

Step	Action
1	Enter the following command: options dns.update.ttl time where <i>time</i> 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. For example, to set the TTL to two hours, enter the following command: options dns.update.ttl 2h

Using NIS to maintain host information

Advantage of using NIS

Like DNS, NIS enables you to centrally maintain host information. NIS provides two methods for storage system host-name resolution:

- ◆ Using a makefile master on the NIS server, which creates a `/etc/hosts` file and copies it to your storage system's default volume for local host name lookup
This method is described in "[Creating /etc/hosts from the NIS master](#)" on page 81.
- ◆ Using a hosts map, maintained as a database on the NIS server, which your storage system queries in a host lookup request across the network
This method is described in this section.

NIS also enables you to maintain user information. For more information, see the *Data ONTAP System Administration Guide*.

Using NIS slave for name resolution

Host-name resolution 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. To improve performance, you can enable an NIS slave on your storage system.

The NIS slave establishes a contact with an NIS master server and does the following two tasks:

- ◆ Downloads the maps from the NIS master server
Once the maps have been downloaded, they are stored in the `/etc/yp/nis_domain_name/` directory. All NIS requests from your storage system are then serviced by the NIS slave 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.
- ◆ Listens for updates from the NIS master
When the maps on the NIS master are changed, the NIS master administrator can choose to notify all slaves. Therefore, in addition to periodically checking for updates from the NIS master, the NIS slave also listens for updates from master.

Note

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

Selection of an NIS master

When the NIS slave is enabled 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 listed with the `nis.servers` option. If that is the case, the servers listed with the `nis.servers` option inform the slave about the master server.

Note

Either the NIS server must have an entry in the hosts map for the master or the `/etc/hosts` file on your storage system must be able to resolve the IP address of the master. Otherwise, the NIS slave on the storage system cannot contact the master.

Guidelines for using the NIS slave

Keep the following guidelines in mind when using the NIS slave on your storage system:

- ◆ 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/nswitch.conf` file must contain DNS as an option to use for host name lookup.

If you have your NIS server configured to do 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, 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 as described previously, the lookups succeed.

- ◆ You can use the NIS slave for the following:
 - ❖ Vif and VLAN interfaces
 - ❖ vFiler units
 - ❖ Storage system clusters

Note

Ensure that the `nis.servers` options value is the same on both cluster nodes and that the `/etc/hosts` file on both cluster nodes can resolve the name of the NIS master server.

About configuring NIS for host lookups

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 or FilerView.

If you configure NIS later, 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

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 the option `nis.slave.enable` to On. For more information about enabling NIS slave, see [“Enabling an NIS slave on your storage system”](#) on page 98.

Data ONTAP interfaces to configure NIS

You can enable NIS and set NIS configuration values in either of these ways:

- ◆ Using FilerView
 - See [“Configuring NIS with FilerView”](#) on page 95.
 - You cannot use FilerView to configure the NIS slave.
- ◆ At the command line
 - See the appropriate instructions:
 - ❖ [“Specifying NIS servers to bind to”](#) on page 96
 - ❖ [“Specifying the NIS domain name”](#) on page 96
 - ❖ [“Enabling or disabling NIS using the command-line interface”](#) on page 96

If you want to use primarily NIS for host-name resolution, specify it ahead of other methods in the hosts map in the /etc/nsswitch.conf file. For information about editing the /etc/nsswitch.conf file, see [“Changing the host name search order”](#) on page 100.

Correct host-name resolution depends on the correct configuration of 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 storage system’s NIS client, see the na_nis(8) man page.

Configuring NIS with FilerView

To set or modify NIS configuration values with FilerView, complete the following steps.

Step	Action	
1	In FilerView, click Network in the list on the left.	
2	In the list under Network, click Manage DNS and NIS Name Service.	
3	If you want to....	Then...
	Enable or disable NIS	Select Yes or No in the NIS Enabled field.
	Set or modify the NIS domain name	Enter a name in the NIS Domain Name field. Examples of configuration values are listed in “Specifying the NIS domain name” on page 96.
	Specify or modify NIS servers	Enter one or more IP addresses in the NIS Servers fields. Examples of configuration values are listed in “Specifying NIS servers to bind to” on page 96.

Enabling or disabling NIS using the command-line interface

To enable or disable NIS on your storage system, complete the following step.

Step	Action
1	Enter the following command: options nis.enable {on off} Use <i>on</i> to enable and <i>off</i> to disable NIS.

Specifying the NIS domain name

To specify the NIS domain name, complete the following step.

Step	Action
1	Enter the following command: options nis.domainname <i>domain</i> <i>domain</i> is the NIS domain name to which your storage system belongs; for example, typical NIS domain names might be <i>sales</i> or <i>marketing</i> . The NIS domain name is usually not the same as the DNS domain name.

Specifying NIS servers to bind to

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 furthest one.

To specify an ordered list of NIS servers you want your storage system to bind to, complete the following step.

Note

You can specify NIS servers by IP address or host name. If host names are used, make sure each host name, along with its IP address, is listed in the */etc/hosts* file of your storage system. Otherwise, the binding with host name will fail.

Step	Action
1	Enter the following command to specify the NIS servers and their order: <pre>options nis.servers ip_address, server_name, *</pre> The asterisk (*) specifies that broadcast is used to bind to NIS servers if the servers in the list are not responding. This is the default. If you do not specify broadcasting (that is, if you do not add the asterisk), and none of the listed servers is responding, NIS services are disrupted until one of the preferred servers responds. You can specify only IPv4 addresses or server names that resolve to IPv4 addresses using the /etc/hosts file on your storage system. Attention _____ Using the NIS broadcast feature can incur security risks. _____

Example of specifying NIS servers to bind to: The following lists two servers and uses the broadcast default:

```
options nis.servers 172.15.16.1,nisserver-1, *
```

Your storage system first tries to bind to 172.15.16.1. 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. While bound to the NIS server that responded to the broadcast, the storage system continues to poll the preferred servers. As soon as one of the preferred servers is found, the storage system binds to the preferred server.

Enabling an NIS slave on your storage system

To enable an NIS slave on your storage system, complete the following step.

Step	Action
1	Enter the following command: <pre>options nis.slave.enable {on off}</pre> Use <code>on</code> to enable the NIS slave and <code>off</code> to disable it. Note_____ If the NIS slave is disabled, your storage system reverts back to the original configuration, in which it contacts an NIS server to resolve host names. _____

Displaying NIS information

You can display the following types of NIS information:

- ◆ NIS domain name
- ◆ Last time the local group cache was updated
- ◆ The following 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
- ◆ Whether an NIS slave is enabled or not
- ◆ NIS master server
- ◆ Last time the NIS map was checked by the NIS slave

To display NIS information, complete the following step.

Step	Action
1	Enter the following command: <pre>nis info</pre>

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

NIS administrative commands

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

Command	Function
<code>ypcat</code>	Prints an entire NIS map
<code>ypgroup</code>	Displays the NIS group cache entries
<code>ypmatch</code>	Looks up specific entries in an NIS map
<code>ypwhich</code>	Returns the name of the current NIS server

Changing the host name search order

How the host name search order is determined

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.

The default `/etc/nsswitch.conf` file

Data ONTAP creates a default `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 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

If you want to change this order, you can do so in either of these ways:

- ◆ By using FilerView
See [“Changing the host name search order with FilerView”](#) on page 101.
- ◆ By editing the `/etc/nsswitch.conf` file
See [“Editing the `/etc/nsswitch.conf` file”](#) on page 101.

Changing the host name search order with FilerView

To change the host name search order with FilerView, complete the following steps.

Step	Action
1	In FilerView, click Network in the list on the left.
2	In the list under Network, click Manage DNS and NIS Name Service.
3	In the Name Service section, select the desired values in the Hosts drop-down lists.

Editing the /etc/nsswitch.conf file

To change the order in which Data ONTAP searches for host information, complete the following steps.

Step	Action
1	If the /etc/nsswitch.conf file does not exist in your storage system's root volume, create it.
2	Edit the file, entering each line in the following format: <pre>map: service ...</pre> <i>map</i> for the host-name resolution service is <code>hosts</code>. <i>service</i> is one or more of the following: <code>files</code>, <code>dns</code>, <code>nis</code>. For example, to change the resolution order to use NIS exclusively, change the hosts line to read as follows: <pre>hosts: nis</pre>
3	Save the file.

About this chapter This chapter describes how Data ONTAP supports SNMP on your storage system and how you can use SNMP to monitor your storage system.

Topics in this chapter This chapter discusses the following topics:

- ◆ [“Understanding SNMP implementation in Data ONTAP”](#) on page 104
- ◆ [“Managing the SNMP agent”](#) on page 113
- ◆ [“Creating SNMP traps”](#) on page 119

Understanding SNMP implementation in Data ONTAP

SNMP process

If Simple Network Management Protocol (SNMP) is enabled in Data ONTAP, SNMP managers can query your storage system's SNMP agent for information (specified in your storage system's MIBs or the MIB-II specification). In response, the SNMP agent gathers information and forwards it to the SNMP managers using the SNMP protocol. The SNMP agent also generates trap notifications whenever specific events occur and sends these traps to the SNMP managers. The SNMP managers can then carry out actions based on information received in the trap notifications.

SNMP agent and MIB groups supported

For diagnostic and other network management services, Data ONTAP provides an SNMP agent compatible with SNMP version 1. This agent 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.

For more information about protocol support, see the `na_snmpd(8)` man page.

Types of traps in Data ONTAP

There are two types of traps in Data ONTAP:

- ◆ Built-in—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

For more information, see “[Understanding traps in Data ONTAP](#)” on page 106.

- ◆ User-defined—User-defined traps exist only after they are defined by a series of `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.

For more information, see “[Creating SNMP traps](#)” on page 119.

About the Data ONTAP MIBs

A Management Information Base (MIB) file is a textual description of SNMP objects and traps. Therefore, the Data ONTAP MIB files document the SNMP capabilities of the Data ONTAP version running on your storage system. MIBs are not configuration files—that is, values in the MIBs are not read by Data ONTAP, and changes to the MIB files do not affect SNMP functionality.

Data ONTAP provides two MIB files:

- ◆ A custom MIB (/etc/mib/netapp.mib)
See “[Contents of the custom MIB](#)” on page 109.
- ◆ An internet SCSI (iSCSI) MIB (/etc/mib/iscsi.mib)
See “[Contents of the iSCSI MIB](#)” on page 112.

Data ONTAP also provides a short cross-reference between object identifiers (OIDs) and object short names in the /etc/mib/traps.dat file. This is useful for creating user-defined traps, as discussed in “[Defining or modifying a trap](#)” on page 121.

Note

The latest versions of the Data ONTAP MIBs and traps.dat files are available online on the NetApp on the Web™ (NOW) site at <http://now.netapp.com/NOW/download/tools/mib/filer.shtml>. However, the versions of these files on the web site do not necessarily correspond to the SNMP capabilities of your Data ONTAP version. They are provided to help you evaluate SNMP features in the latest Data ONTAP release.

About traps

Traps are mechanisms that alert you to significant events on your storage system. If SNMP is configured, traps are fired when a defined event, such as a network traffic interruption or line power failure, occurs. Trap information, in the form of MIB Object Identifiers (OIDs), is sent from your storage system’s agent to an SNMP management station.

About built-in traps in Data ONTAP MIBs

Built-in traps in Data ONTAP MIBs are identified by the string TRAP-TYPE. For example, the following is a complete trap definition from the Data ONTAP custom MIB:

```
upsLinePowerOff          TRAP-TYPE
ENTERPRISE               netapp
DESCRIPTION
    "UPS: Input line power has failed and UPS is now on battery."

 ::= 142
```

Traps in the custom MIB are provided in a number of categories, including the following:

Category	Examples of trap messages
Disk Health Monitor	Degraded I/O, disk predictive-failure event
Disks	Disk - failure alert, shutdown, repaired
Fan	Fan - failed, shutdown, warning, repaired
Power supply	Power supply - failed, shutdown, warning, repaired
CPU	CPU busy, OK
NVRAM	Battery discharged, low
Cluster	Node failed, repaired
Volumes	Nearly full, full, repaired

Category	Examples of trap messages
Temperature	Over temperature, shutdown, repaired
Shelf	Fault, repaired
Global	Not recoverable, critical, not critical, OK
Soft quotas	Exceeded, normal
Autosupport	Send error, configuration error, successful send

Note

These categories are examples of the MIB trap contents; it is not an exhaustive list. The most complete listings are provided in the MIBs themselves.

About MIB trap priority

By convention, the right-most digit of a trap ID number indicates its priority (degree of severity), using the same enumeration as syslog entries. For example, trap ID 142 `upsLinePowerOff` is priority 2, alert.

Trap priorities are listed in the following table.

Trap ID last digit	Priority
1	emergency
2	alert
3	critical
4	error
5	warning
6	notification
7	information
8	debug

For more information, see the `na_syslog.conf(5)` man page.

Where to get further information

See the following RFCs for more information:

- ◆ RFC 1157—Defines and describes SNMP.
- ◆ RFC 1213—Defines and describes the SNMP MIB-II specification.
- ◆ RFC 1155—Defines and describes the structure and identification of management information for TCP/IP-based internets.
- ◆ RFC 1215—Defines the convention for defining traps for use with SNMP.
- ◆ RFC 1212—Gives concise MIB definitions.

Contents of the custom MIB

About the custom MIB

The custom MIB provides detailed information about many aspects of storage system operation. The custom MIB file, `netapp.mib`, is located in the `/etc/mib` directory on your storage system.

The custom MIB was verified using `smilint` from the `libsmi` tool version 0.4.0.

The custom MIB groups

The top-level groups in the custom MIB that are relevant to your storage system are described in the following table.

Note Information about the objects described in this table is available for your storage system only if the corresponding feature is enabled on that storage system.

Group name	Contents
product	Product information, such as the software version and system ID.
sysStat	System-level statistics, such as CPU up-time and idle time, total number of kilobytes transmitted and received on all network interfaces, cluster takeover status, hardware temperature, and power supply status. Note If your storage system is not licensed for cluster setup, a value indicating no cluster license is returned.
nfs	Statistics such as those displayed by the <code>nfsstat</code> command, including statistics for each client if per-client statistics are enabled. The per-client statistics are indexed by client IP addresses.
quota	Information related to disk quotas, including the output of the <code>quota report</code> command.

Group name	Contents
filesystem	Information related to the file system, including the equivalent of the <code>maxfiles</code> and <code>df</code> commands, overall mirror status, number of plexes in a file system, and some of the information from the Snapshot™ <code>snap list</code> command.
raid	Redundant Array of Independent Disks (RAID) configuration and plex-specific information, such as plex ID, plex status, and parent volume.
cifs	Statistics such as those displayed by the <code>cifs stat</code> command.
snapmirror	SnapMirror® statistics, such as status, number of active backups and restores, and number of bytes read and written by SnapMirror.
ndmp	Information related to NDMP sessions, including the equivalent of the <code>ndmpd status</code> command, such as currently active sessions and number of backups and restores that succeeded.
fabric	Information about the Storage Area Network (SAN) fabric, including status and configuration.
dafs	Direct Access File System (DAFS) statistics, including status, requests, sessions, calls, and interface information.
vi	Information about the virtual interface.
backup	Information about dump and restore activities.
vfiler	Information about vFiler units, including status, numbers of vFiler units per physical storage system, and licensing information.
blocks	Information about block transfer activities, including read, write and ops statistics, protocols licensed, and LUNs configured.
nfscache	Information about <code>nfscache</code> .

Group name	Contents
snapvault	SnapVault® statistics, such as status, number of active backups and restores, and number of bytes read and written by SnapVault.
ftpd	ftpd statistics, including status, connections, and daemon information.

Contents of the iSCSI MIB

About the iSCSI MIB The iSCSI MIB provided with Data ONTAP is an SMIV1 (Structure of Management Information version 1) version of the SMIV2 iSCSI MIB draft 09. Because the Data ONTAP SNMP implementation does not support SMIV2 syntax, the iSCSI MIB is a port of the draft standard to SMIV1 in accordance with RFC 2576.

You can get the iSCSI MIB from the following sources:

- ◆ The `/etc/mib/iscsi.mib` file on your storage system, after you have installed the Data ONTAP software
- ◆ The NOW™ Web site at <http://now.netapp.com/>

A short cross-reference between iSCSI OIDs and short names is included in the `/etc/mib/traps.dat` file.

iSCSI management objects

The following list presents an overview of the iSCSI management objects in the iSCSI MIB. See the MIB file for more information.

- ◆ Header and data descriptors
- ◆ Instances
- ◆ Portals
 - ❖ Targets
 - ❖ Initiators
- ◆ Nodes
 - ❖ Targets
 - ❖ Target authorization
 - ❖ Initiators
 - ❖ Initiator authorization
- ◆ Sessions
- ◆ Connections

Managing the SNMP agent

About your storage system's SNMP agent

Your storage system's SNMP agent responds to queries and sends traps to network management stations. Your storage system's SNMP agent does not have write privileges—that is, it cannot be used to take corrective action in response to a trap.

What SNMP agent management includes

To configure the SNMP agent on your storage system, you must do all of the following:

- ◆ Verify that SNMP is enabled.
SNMP is enabled by default in Data ONTAP.
- ◆ Enable traps.
Although SNMP is enabled by default, traps are disabled by default.
- ◆ Specify one or more network management station host names.
No traps are sent unless at least one SNMP management station is specified as a trap host. Trap notifications can be sent to a maximum of eight network management stations.

You can optionally do any or all of the following:

- ◆ 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. See “[Setting SNMP access privileges](#)” on page 115.
- ◆ 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 SNMP agent authentication failure traps, which are generated when the agent receives queries with the wrong community string. The traps are sent to all hosts specified as trap hosts.

- ◆ Create and load user-defined traps.

For more information, see “[Creating SNMP traps](#)” on page 119.

You can also view current SNMP and trap configuration. The following sections explain how to perform these tasks.

Note

Storage systems in a cluster can have different SNMP configurations.

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

About configuration tools

The following tools are available for storage system SNMP configuration and management:

Command-line interface	Graphical interface
<code>snmp</code> command For more information, see the <code>na_snmp(1)</code> man page.	FilerView SNMP windows For more information, see FilerView Help.

Note

SNMP commands entered at the command line or in FilerView are persistent across reboots.

Enabling SNMP at the command line

To enable SNMP, complete the following step.

Step	Action
1	Enter the following command at your storage system command line: <code>options snmp.enable {on off}</code> Enables (with value <code>on</code>) or disables (with value <code>off</code>) SNMP in Data ONTAP. For more information about this option, see the <code>na_options(1)</code> man page.

Setting SNMP access privileges

To set SNMP access privileges on a host or interface basis, complete the following step. (You cannot set SNMP access privileges in FilerView.)

Step	Action
1	Enter the following command at your storage system command line: <code>options snmp.access options</code> For details about using this option, see the <code>na_protocolaccess(8)</code> man page.

Viewing and modifying SNMP configuration values at the command line

To view or modify SNMP configuration values, complete the following step.

Step	Action
1	Enter the following command at your storage system command line: <code>snmp {options values}</code> Examples of configuration values are listed in “Example of typical SNMP commands” on page 118. For more information about <code>snmp</code> parameters, see “Command syntax for SNMP configuration parameters” on page 116.

Viewing and modifying SNMP configuration values with FilerView

To view or modify SNMP configuration values with FilerView, complete the following steps.

Step	Action
1	In FilerView, click SNMP in the list on the left.
2	In the list under SNMP, click Configure.

Step	Action	
3	If you want to....	Then...
	View SNMP configuration values	The current configuration is displayed.
	Set or modify SNMP configuration values	Enter configuration values in drop-down lists or text fields. Click Apply when finished. Examples of configuration values are listed in “ Example of typical SNMP commands ” on page 118.

Command syntax for SNMP configuration parameters

The following table lists the SNMP configuration commands and parameters available in Data ONTAP. 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.

Command	Description
<code>snmp</code>	Displays the current values of all SNMP options, such as <code>init</code> , <code>community</code> , <code>contact</code> , and <code>traphost</code> .
<code>snmp authtrap [0 1]</code>	With a value: Enables (with value 1) or disables (with value 0) SNMP agent authentication failure traps. Without a value: Displays the current value of <code>authtrap</code> set in Data ONTAP.
<code>snmp community</code>	Displays the current list of communities.

Command	Description
<code>snmp community add ro community</code>	Adds a community. 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 <code>ro</code> (read-only).
<code>snmp community delete {all ro community}</code>	Deletes one or all communities.
<code>snmp contact [contact]</code>	With the option: Sets the contact name for your storage system. You must enclose the <i>contact</i> string in single quotes (' ') if the string contains spaces. You can enter a maximum of 255 characters for the contact information. Without the option: Displays the current contact name set in Data ONTAP.
<code>snmp init [0 1]</code>	With a value: Enables (with value 1) or disables (with value 0) built-in traps and the traps defined using the <code>snmp traps</code> command. Without a value: Displays the current value of <code>snmp init</code> in Data ONTAP. Default value: By default, SNMP traps are disabled in Data ONTAP; the system uses the equivalent of <code>snmp init 0</code>.

Command	Description
<code>snmp location [location]</code>	With the option: Sets the location of your storage system. You must enclose the <i>location</i> string in single quotes (' ') if the string contains spaces. Without the option: Displays the current location set in Data ONTAP.
<code>snmp traphost [{add delete} {hostname ipaddress}]</code>	With the option: Adds or deletes SNMP hosts that receive traps from Data ONTAP. Without the option: Displays the current trap hosts set in Data ONTAP.
<code>snmp traps [options]</code>	See “Command syntax for SNMP trap parameters” on page 123.

Example of typical SNMP commands

The following example shows a typical set of commands to configure SNMP monitoring. It assumes that SNMP remains enabled by default.

```
snmp contact 'jdoe@abc.com 415-555-1212'
snmp location 'ABC corporation, engineering lab'
snmp community add ro private
snmp traphost add snmp-mgr1
snmp init 1
```

Creating SNMP traps

Working with SNMP traps

You can create user-defined traps in Data ONTAP if the predefined built-in traps are not sufficient to create alerts for conditions you wish to monitor.

Note

Before you invest the effort to define a new trap, you are advised to consult the Data ONTAP MIBs to see if any existing traps serve your purpose. For more information, see [“Understanding traps in Data ONTAP”](#) on page 106.

The following sections explain

- ◆ [“Understanding user-defined traps”](#) on page 120
- ◆ [“Defining or modifying a trap”](#) on page 121
- ◆ [“SNMP trap parameters”](#) on page 126

Understanding user-defined traps

About user-defined traps

You can set 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 and have the SNMP application on your network management station show a flashing message on your console when a fan has stopped working.

Traps are persistent. After you set a trap, it exists across reboots until you remove it or modify it.

Guidelines for creating traps

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.
- ◆ Make sure the condition you intend to trap can be generated in your storage system's 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.

Defining or modifying a trap

Ways to define or modify a trap

You can define traps or modify traps you have already defined by entering values in one of the following ways:

- ◆ At the command line
See [“Viewing and modifying trap values at the command line”](#) on page 122.
- ◆ Using FilerView
See [“Viewing or modifying trap values with FilerView”](#) on page 122.
- ◆ In a configuration file
See [“Command syntax for SNMP trap parameters”](#) on page 123.

You must supply the following elements when creating or modifying traps.

- ◆ Trap name
This is the name of the user-defined trap you want to create or change.

Note

A trap name must have no embedded periods.

- ◆ Trap parameters
These are parameters defined in [“SNMP trap parameters”](#) on page 126.
- ◆ Parameter value
This is the value you assign to a trap parameter.

Note

When you create a user-defined trap, it is initially disabled by default. You must enable a trap before it can be triggered using the `snmp traps` command or FilerView.

Viewing and modifying trap values at the command line

To view or modify traps using the command-line interface, complete the following step.

Step	Action
1	At your storage system command line, enter the following command: snmp traps {options variables} For more information about snmp traps parameters, see “ Command syntax for SNMP trap parameters ” on page 123.

Viewing or modifying trap values with FilerView

To define or modify traps using FilerView, complete the following steps.

Step	Action	
1	In FilerView, click SNMP in the list on the left.	
2	In the list under SNMP, click Traps.	
3	If you want to....	Then...
	Create a new trap	1. Click Add. 2. In the Add an SNMP Trap window, enter the requested information and click Add.
	View or modify an existing trap	1. Click Manage for the trap you want. 2. In the Manage SNMP Traps window, click Modify.

Example of trap definitions

The following command-line example sets a group of traps. The trap descriptions are numbered in brackets.

The same parameters can be entered using FilerView.

Example:

```
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
```

Explanation: 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 fires and continues to fire every hour [5] until the total number of CIFS operations drops below one million.

Command syntax for SNMP trap parameters

The following table lists the SNMP trap commands available in Data ONTAP. 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.

Command	Description
<code>snmp traps</code>	Displays the list of user-defined traps set in Data ONTAP.
<code>snmp traps [enable disable reset delete] [trapname]</code>	Enables, disables, resets, or deletes the <i>trapname</i> trap. If you do not specify the <i>trapname</i> trap, all traps defined so far are acted on.
<code>snmp traps walk prefix</code>	Walks (traverses in order) the trap list by prefix; that is, lists all traps that have names beginning with <i>prefix</i> .
<code>snmp traps load trap_list_filename</code>	Loads a set of traps from a text file. The <i>trap_list_filename</i> file contains a list of traps without the <code>snmp traps</code> command preceding each trap. If the specified file name is defaults, traps are read from the <code>/etc/defaults/traps</code> file.

Command	Description
<code>snmp traps <i>trapname.parm value</i></code>	Defines or changes a user-defined trap parameter. See “ SNMP trap parameters ” on page 126.

Defining and modifying traps in a configuration file

You are advised to define traps in a configuration file, which is then loaded with the `snmp traps load` command. If you define and load traps this way, Data ONTAP automatically backs up your SNMP configuration in Snapshot copies, making it easy to transfer user-defined traps to other storage systems, and simplifying recovery of SNMP configurations if there is some kind of disaster.

To create a trap configuration file, follow these steps.

Step	Action
1	Create a traps configuration file on your storage system — for example, <code>/etc/mib/mytraps</code> . The name and location of the file is at your discretion.
2	Enter the traps in the configuration file in the following form: <i>trapname.parm value</i> That is, use parameters of the <code>snmp traps</code> command without the command name. For example, to set the <code>cifstotalops</code> trap listed in the previous command example, enter the following lines in your configuration file: <pre>cifstotalops.var snmp.1.3.6.1.4.1.789.1.7.3.1.1.1.0 cifstotalops.trigger level-trigger cifstotalops.edge-1 1000000 cifstotalops.interval 10 cifstotalops.backoff-calculator step-backoff cifstotalops.backoff-step 3590 cifstotalops.rate-interval 3600 cifstotalops.priority alert cifstotalops.message snmp.1.3.6.1.4.1.789.1.7.3.1.1.1.0</pre>

Step	Action
3	Test each line of the file by entering the <code>snmp traps</code> command at the command line or by specifying the trap with FilerView. Make corrections as needed.
4	Load the configuration file with the <code>snmp traps load</code> command. For example: <code>snmp traps load /etc/mib/mytraps</code>

SNMP trap parameters

SNMP trap parameters

The following table lists parameters that you use to create traps.

- ◆ The left-hand column lists parameters that you enter at the command line with the `snmp traps` command, as described in “[Command syntax for SNMP trap parameters](#)” on page 123.
- ◆ The right-hand column lists the equivalent parameters that you select in FilerView, as described in “[Viewing or modifying trap values with FilerView](#)” on page 122.

The sections following the table describe individual parameters. See also “[Example of trap definitions](#)” on page 122.

Parameter in commands	Equivalent in FilerView
var	OID
trigger	Trigger
edge-1 edge-2	Edge 1 Edge 2
edge-1-direction edge-2-direction	Edge 1 Direction 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 of the form `snmp.oid`, where *oid* is an Object Identifier (OID).

Note

The `traps.dat` file, 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 only necessary to enter the numerical OID, 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`—Fires 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`—Fires 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`—Fires 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 discussed later in this section). 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 let you 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 values for the edge-1-direction and the edge-2-direction parameters are

- ◆ edge-1-direction—up
- ◆ edge-2-direction—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 interval value for the Data ONTAP predefined traps is 60, or one minute.

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 (at system startup, for example).

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 (Backoff Style)

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 for 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 fire 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 can not set the message parameter in FilerView.

About this chapter

This chapter discusses the concepts underlying virtual local area networks (VLANs) and VLAN tagging, how VLANs are implemented in Data ONTAP, and how you manage VLANs on your storage system.

Topics in this chapter

This chapter discusses the following topics:

- ◆ [“Understanding VLANs”](#) on page 134
- ◆ [“VLANs in Data ONTAP”](#) on page 138
- ◆ [“Managing VLANs on your storage system”](#) on page 140

Understanding VLANs

What a VLAN is

A VLAN is a logical network segment that can span multiple physical network segments. The end-stations belonging to a VLAN are related by function or application. For example, end-stations 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.

About 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 *port-based*, or 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 will be flooded only among the members of that VLAN. Communication *between* VLANs, however, 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.

What GVRP is

GARP VLAN Registration Protocol (GVRP) uses the 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.

Because GVRP provides dynamic registration of VLAN membership, members can be added or removed from a VLAN on the fly, 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 802.1D:1998 edition).

What a VLAN tag is

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.

How VLAN tagging works

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 flooding of the frame to ports that belong to the identified VLAN.

For example, in the previous 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. Because the destination MAC address is not known to Switch 1, 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 thereby receives the frame.

Advantages of VLANs

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 can be isolated 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

The following requirements must be satisfied before you set up VLANs in a network:

- ◆ 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.

If an end-station cannot register or cannot be configured to belong to a VLAN, the end-station can belong only to one VLAN. This VLAN is configured on the switch port to which the end-station connects. The frames sent on this switch port are untagged.

VLANs in Data ONTAP

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 will support. This information (dynamic registration) is updated periodically thereafter. This information is also sent every time an interface comes up after being down or whenever there is a change in the VLAN configuration of the interface.

Guidelines for setting up VLANs in Data ONTAP

VLANs in Data ONTAP are implemented in compliance with the IEEE 802.1Q standard. Additionally, you must follow the following guidelines while setting up VLANs in Data ONTAP:

- ◆ You cannot set up VLANs using the `setup` procedure. You must use the command line or the FilerView interface to create, change, or destroy VLANs.
- ◆ You must add the commands to create VLANs on your 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 your storage system; however, Data ONTAP imposes a limit of 128 interfaces (including physical, vif, vlan, vh, and loopback interfaces) per storage system.
- ◆ You can create VLANs on physical interfaces as well as vifs. For more information about vifs, see Chapter 7, “[Configuring vifs](#),” on page 151.
- ◆ 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, different MTU sizes can be specified for individual VLANs.
- ◆ You can assign an identification number from 1 to 4,094 to a VLAN.
- ◆ You must ensure that the interface on your storage system is also a member of its partner’s VLANs in a cluster failover pair.
- ◆ You cannot configure any parameters except `mediatype` for the physical network interface configured to handle VLANs.

Interfaces that do not support VLANs

The FDDI and ATM interfaces do not support VLANs.

Reverting to earlier versions of Data ONTAP

Reverting to Data ONTAP 6.1 or 6.1.x: If your storage system is a member of a VLAN and you need to revert to Data ONTAP 6.1 or 6.1.x, you must ensure that the `ifconfig` commands in the `/etc/rc` file do not contain the `-g GVRP` flag or the `vlan modify` command.

Reverting to a version earlier than Data ONTAP 6.1: If your storage system is a member of a VLAN and you need to revert to a version earlier than Data ONTAP 6.1, you must make sure that the `ifconfig` commands in the `/etc/rc` file do not contain any VLAN configuration information.

Managing VLANs on your storage system

Command for managing VLANs on your storage system

You manage VLANs on your storage system using the `vlan` command. This command allows you to create, add interfaces to, delete, and display statistics of a VLAN.

The `vlan` command syntax

The `vlan` command syntax is as follows:

```
vlan create [-g {on|off}] ifname vlanid_list
vlan add ifname vlanid_list
vlan delete -q ifname [vlanid_list]
vlan modify -g {on|off} ifname
vlan stat ifname [vlanid_list]
```

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

Persistence of the `vlan` commands

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.

For detailed information

For detailed information on how to perform specific tasks using the `vlan` command, see the following topics:

- ◆ [“Creating and configuring a VLAN on your storage system”](#) on page 141
- ◆ [“Adding an interface to a VLAN”](#) on page 144
- ◆ [“Deleting a VLAN”](#) on page 145
- ◆ [“Modifying VLAN interfaces”](#) on page 147
- ◆ [“Viewing VLAN statistics”](#) on page 148

Creating and configuring a VLAN on your storage system

Commands for creating and configuring a VLAN

Creating and configuring a VLAN involves two commands: the `vlan create` command and the `ifconfig` command.

The `vlan create` command creates a VLAN interface, includes that interface in one or more VLAN groups as specified by the VLAN identifier, enables VLAN tagging, and enables (optionally) GVRP on that interface.

The `ifconfig` command enables you to configure the VLAN interface created by the `vlan` command.

About enabling and disabling GVRP on VLAN interfaces

By default, GVRP is disabled on VLAN interfaces created using the `vlan create` command; however, you can enable it with the `-g` flag available with the command.

If you enable GVRP on an interface that is configured down, the state of the interface and all associated VLAN interfaces is automatically configured up. This state change occurs so that the interface can start sending VLAN registration frames to register its VLAN membership with the switch.

Creating a VLAN on your storage system

To create a VLAN on your storage system, complete the following step.

Note— You must be familiar with “[Guidelines for setting up VLANs in Data ONTAP](#)” on page 138 before proceeding with the following procedure.

Step	Action
1	Enter the following command: <pre>vlan create [-g {on off}] ifname vlanid</pre> -g enables (on) or disables (off) GVRP on an interface. By default, GVRP is disabled on the interface. <i>ifname</i> is the name of the network interface. <i>vlanid</i> is the VLAN identifier to which the <i>ifname</i> interface belongs. You can include a list of VLAN identifiers. Result: A VLAN interface with the name <i>ifname-vlanid</i> is created. Note— VLANs created using the <code>vlan create</code> command are not persistent across reboots unless the <code>vlan</code> commands are added to the <code>/etc/rc</code> file.

Example of creating a VLAN interface

You can create VLANs with identifiers 10, 20, and 30 on interface e4 of a storage system 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:

```
e4: flags=80008042<BROADCAST,RUNNING,MULTICAST, VLAN> mtu 1500
```

Configuring an interface in a VLAN

Using the `ifconfig` command, you can configure all the parameters for a VLAN interface that you can for a physical interface. The parameters you can configure are

- ◆ IP address
- ◆ Network mask
- ◆ Interface status

- ◆ Media type
- ◆ MTU size
- ◆ Flow control
- ◆ Partner

For detailed information about the `ifconfig` command, see Chapter 1, “[Network Interface Configuration](#),” on page 1.

To configure the IP address and network mask for a VLAN interface, complete the following step.

Step	Action
1	Enter the following command: <pre>ifconfig ifname-vlanid IP_address netmask mask</pre> <i>ifname-vlanid</i> is the VLAN interface name. <i>IP_address</i> is the IP address for this interface. <i>mask</i> is the network mask for this interface. Example: You can configure a VLAN interface e4-10, created in the previous example, using the following command: <pre>ifconfig e4-10 172.25.66.11 netmask 255.255.255.0</pre>

Managing VLANs on your storage system

Adding an interface to a VLAN

Command for adding an interface to a VLAN

If a physical interface does not belong to any VLAN, you 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 must use the `vlan add` command to add the interface to subsequent VLANs.

Like the `vlan create` command, the `vlan add` command creates a VLAN interface that must be configured using the `ifconfig` command.

Adding an interface to a VLAN

To add an interface to a VLAN, complete the following step.

Step	Action
1	Enter the following command: <pre>vlan add ifname vlanid</pre> <i>ifname</i> is the name of the network interface. <i>vlanid</i> is the VLAN identifier to which the <i>ifname</i> interface belongs. You can include a list of VLAN identifiers. Result: A VLAN interface with the name <i>ifname-vlanid</i> is created. Note_____ VLANs created using the <code>vlan add</code> commands are not persistent across reboots unless the <code>vlan</code> commands are added to the <code>/etc/rc</code> file. _____

Example of adding an interface to a VLAN

You can add VLANs with identifiers 40 and 50 on interface `e4` of a storage system using the following command:

```
vlan add e4 40 50
```

As a result, VLAN interfaces `e4-40` and `e4-50` are created.

Managing VLANs on your storage system

Deleting a VLAN

Command for deleting a VLAN

The `vlan delete` command is used to delete the VLANs on an interface. You can delete either a specific VLAN or all VLANs associated with that interface. If all VLANs for an interface are deleted, the interface is available to be configured as a regular physical interface.

Deleting a VLAN

To delete a VLAN on your storage system, complete the following step.

Note—
By default, the `vlan delete` command prompts you to confirm the deletion. If you do not want to receive this prompt, use the `-q` flag. This action invokes *quiet mode*, which causes the operation to complete without prompting.

Step	Action	
1	If you want to delete...	Then...
	All VLANs	Enter the following command: vlan delete [-q] ifname <i>ifname</i> is the name of the network interface.
	Example: You delete all VLANs configured on interface e4 with the following command: vlan delete e4	

Step	Action	
	If you want to delete...	Then...
	A specific VLAN	Enter the following command: <pre>vlan delete [-q] ifname vlanid</pre> <i>ifname</i> is the name of the network interface. <i>vlanid</i> is the VLAN identifier to which the <i>ifname</i> interface belongs. You can include a list of VLAN identifiers.
	Example: You delete VLAN e4-30 with the following command: <pre>vlan delete e4 30</pre>	

Managing VLANs on your storage system

Modifying VLAN interfaces

**Command for
modifying VLAN
interfaces**

The `vlan modify` command enables or disables GVRP on all the interfaces of a network adapter. That is, you can enable GVRP on network adapter e8 of a storage system, but not on the VLAN interface e8-2. Once you enable GVRP on a network adapter, it is enabled on all associated VLAN interfaces.

**Modifying VLAN
interfaces**

To enable or disable GVRP on VLAN interfaces, complete the following step.

Step	Action
1	Enter the following command: <pre>vlan modify -g {on off} adap_name</pre> -g enables (on) or disables (off) GVRP. <i>adap_name</i> is the name of the network adapter. Note_____ VLANs modified using the <code>vlan modify</code> command are not persistent across reboots unless the <code>vlan</code> commands are added to the <code>/etc/rc</code> file. _____

Command for displaying VLAN statistics

The `vlan stat` command is used to display the statistics of network interfaces configured in VLANs on your storage system. In addition to displaying the frames received and transmitted on an interface, this command displays the number of frames that were rejected because the frames did not belong to any of the VLAN groups to which the interface belongs.

Viewing VLAN statistics

To view VLAN statistics on your storage system, complete the following step.

Step	Action	
1	If you want to view...	Then...
	Statistics of all VLANs configured on a network interface	Enter the following command: vlan stat ifname <i>ifname</i> is the name of the network interface.
	Statistics of a specific VLAN configured on a network interface	Enter the following command: vlan stat ifname vlanid <i>ifname</i> is the name of the network interface. <i>vlanid</i> is the VLAN identifier to which the <i>ifname</i> interface belongs. You can include a list of VLAN identifiers.

Example of the vlan stat command

The following example displays the statistics of all VLANs on a storage system named toaster:

```
toaster> 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

RECEIVE STATISTICS
  Total frames:      0 | Total bytes: 0 | Multi/broadcast:    0
TRANSMIT STATISTICS
  Total frames:      0 | Total bytes: 0 | Multi/broadcast:    0
  Queue overflows:    0

Vlan Interface: e4-5 (0 hours, 0 minutes, 7 seconds) --
ID: 5
MAC Address: 00:90:27:5c:58:14

RECEIVE STATISTICS
  Total frames:      0 | Total bytes: 0 | Multi/broadcast:    0
TRANSMIT STATISTICS
  Total frames:      0 | Total bytes: 0 | Multi/broadcast:    0
  Queue overflows:    0
```


About this chapter

This chapter discusses vifs (a feature that implements link aggregation on your storage system) and how you can manage various types of vifs on your storage system.

Topics in this chapter

This chapter discusses the following topics:

- ◆ “[Understanding vifs](#)” on page 152
- ◆ “[Types of vifs](#)” on page 154
- ◆ “[Managing vifs](#)” on page 159
- ◆ “[Second-level vifs](#)” on page 175

Understanding vifs

About vifs

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

Different vendors also refer to vifs using these terms:

- ◆ Virtual aggregations
- ◆ Link aggregations
- ◆ Trunks
- ◆ EtherChannel


Advantages of vifs

Using vifs provides several advantages over using individual network interfaces, such as the following:

- ◆ Higher throughput—Multiple interfaces work as one interface.
- ◆ Fault tolerance—If one interface in a vif goes down, your storage system can stay connected to the network using the other interfaces.
- ◆ No single point of failure—If the physical interfaces in a vif are connected to different switches and a switch goes down, your storage system stays connected to the network through the other switches.


Storage system interfaces before grouping into a vif

The following figure shows four separate storage system interfaces, e3a, e3b, e3c, and e3d, before grouping into a vif.



Storage system interfaces after grouping into a vif

The following figure shows the four storage system interfaces grouped into a single vif called Trunk1.



Types of vifs

Kinds of vifs on your storage system

There are two kinds of vifs:

- ◆ Single-mode
- ◆ Multimode


Single-mode vif operation

In a single-mode vif, only one of the interfaces in the vif is active. The other interfaces are on standby, ready to take over if the active interface fails. Failure means that the link status of the interface is down, which signals that the interface has lost connection with the switch.

There can be more than one interface on standby in a single-mode vif. 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 vif does not require any switch configuration or a switch that supports link aggregation.

All interfaces in a single-mode vif share a common Media Access Control (MAC) address.

Example: In the following figure, e0 and e1 are part of the SingleTrunk1 single-mode vif. The active interface, e0, fails. The standby e1 interface takes over and maintains the connection to the switch.



Multimode vif operation

Notes about multimode vif implementation: The multimode vif implementation in Data ONTAP is in compliance with IEEE 802.3ad (static). Dynamic 802.3ad, also known as Link Aggregation Control Protocol (LACP), is not supported for the multimode vifs. Additionally, Port Aggregation Protocol (PAgP), Cisco's proprietary link aggregation protocol, is not supported. Any switch that supports aggregates, but does not have control packet exchange for configuring an aggregate can be used for the multimode vifs.

How multimode vifs work: In a multimode vif, all interfaces in the vif are active and share a single MAC address. This logical aggregation of interfaces provides higher throughput than a single-mode vif. Multimode vifs can recover from a failure of up to $(n-1)$ interfaces, where n is the total number of interfaces that form the vif.


A multimode vif requires a switch that supports link aggregation over multiple switch ports. The switch is configured so that all ports to which links of a vif are connected are part of a single logical port. For information about configuring the switch, see your switch vendor's documentation. For an example of how you can configure a Cisco® Catalyst® 6500 series switch, see [“How to configure a Cisco Catalyst 6500 series switch for a multimode vif”](#) on page 156.

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 multimode vif. The load-balancing schemes are discussed in detail in [“Load balancing in multimode vifs”](#) on page 158.

Data ONTAP is only responsible for distributing outbound traffic and does not have control over how inbound packet arrive because each end of an aggregate is responsible for controlling the distribution of its outbound traffic.

Example of a multimode vif: In the following figure, e0, e1, e2, and e3 are part of the MultiTrunk1 multimode vif. All four interfaces in the MultiTrunk1 multimode vif are active.



If any three of the interfaces fail, either one by one or simultaneously, your storage system still stays connected to the network.

How to configure a Cisco Catalyst 6500 series switch for a multi-mode vif: You can configure a Cisco Catalyst 6500 series switch that has one module (installed in slot 2) and one Supervisor Engine installed (in slot 1) and is running CatOS for a multimode vif by using the `set port channel` command at the switch console.

Syntax of the Cisco `set port channel` command: The syntax of the `set port channel` command is as follows.

For usage guidelines and examples of the `set port channel` command, see the Catalyst 6500 series Command Reference.

Note

The syntax of the `set port channel` command is as it was available at the time of writing the *Data ONTAP 7.1 MultiStore Guide*. You are advised to consult your latest Catalyst documentation to ensure the validity of this information.

```

set port channel mod_num/port_num... [admin_group]
set port channel mod_num/port_num mode {on | off | desirable |
auto} [silent | non-silent]
set port channel all distribution {ip | mac} [source | destination
| both]

```

mod_num/port_num is the number of module and port(s).

admin_group is the number of administrative group. Valid values are 1 to 1024.

mode—Specifies the EtherChannel mode.

on—Enables and forces specified ports to channel without PAgP.

`off`—Prevents ports from channeling.

`desirable`—Sets a PAgP mode that places a port into an active negotiating state, in which the port initiates negotiations with other ports by sending PAgP packets.

`auto`—(Optional) Sets a PAgP mode that places a port into a passive negotiating state, in which the port responds to PAgP packets it receives, but does not initiate PAgP packet negotiation.

`silent`—Keyword to use with `auto` or `desirable` when no traffic is expected from the other device, to prevent the link from being reported to Spanning Tree Protocol (STP) as down.

`non-silent`—Keyword to use with `auto` or `desirable` when traffic is expected from the other device.

`all distribution`—Keyword to apply frame distribution to all ports in the switch.

`ip`—Keyword to specify the frame distribution method using IP address values.

`mac`—Keyword to specify the frame distribution method using MAC address values.

`source`—(Optional) Keyword to specify the frame distribution method using source address values.

`destination`—(Optional) Keyword to specify the frame distribution method using destination address values.

`both`—(Optional) Keyword to specify the frame distribution method using source and destination address values.

The default is EtherChannel, set to `auto` and `silent` on all module ports. The defaults for frame distribution are `ip` and `both`.

Example of how to configure a Cisco Catalyst 6500 series switch for a multimode vif: To configure EtherChannel on ports 2 through 5 of module installed in slot 2, enter the following command:

```
set port channel 2/2-5
```

To set the EtherChannel load-balancing feature so that IP addresses of both source and destination addresses are utilized to determine the outgoing frames, enter the following command:

```
set port channel all distribution ip both
```

Load balancing in multimode vifs

To ensure that all interfaces of a multimode vif are equally utilized for outgoing traffic, the following load-balancing methods are available:

- ◆ IP-address based
- ◆ MAC-address based
- ◆ Round robin

The load-balancing method to use for a multimode vif can be specified only when the vif is created. If no method is specified, the IP-address based load balancing method is used.

IP-address and MAC-address based: In both of these methods, the last byte of the source and destination address (IP address and MAC address) is used to determine the interface to use for the outgoing frame. The following formula is used:

$((source_address \text{ XOR } destination_address) \% number_of_links)$

If the result of this formula maps to an interface that is not in the UP link-state, the next active interface is used.

For example, a vif consisting of eight physical interfaces is created with the IP address-based load-balancing method. It is configured with IP address 10.0.0.10. Based on the above formula, an IP frame going through this vif to the destination IP address 172.26.15.224 will use interface #2, provided that this interface is in the UP link-state.

Note

Do not select the MAC-address based load-balancing method when creating vifs on a storage system that connects directly to a router. In such a setup, for every outgoing IP frame, the destination MAC address will be the MAC address of the router. As a result, only one interface of the vif will be used.

Round robin: Unlike the IP-address and MAC-address load-balancing methods, this method provides true load balancing. This method may cause out-of-order packet delivery and retransmissions due to overruns.

This method of load balancing is recommended for clients connected in a back-to-back configuration with your storage system.

Managing vifs

About managing vifs

You manage vifs on your storage system with the `vif` command. This command enables you to create, add interfaces to, delete interfaces from, display status and statistics of, and destroy a vif.

Guidelines for creating and configuring vifs on your storage system

The following guidelines apply to creating and configuring vifs on your storage system:

- ◆ You can group up to 16 physical Ethernet interfaces on your storage system to obtain a vif.
The network interfaces that are part of a vif do not have to be on the same network adapter, but it is best that all network interfaces be full-duplex.
- ◆ You cannot include a virtual LAN (VLAN) interface in a vif.
- ◆ The interfaces that form a vif must have the same Maximum Transmission Unit (MTU) size.

You can use the `ifconfig` command to configure the MTU size on the interfaces of a vif. You need to configure the MTU size only if you are enabling jumbo frames on the interfaces. For more information about jumbo frames, see [“Understanding frame size, MTU size, and jumbo frames”](#) on page 5.

- ◆ You can include any Gigabit Ethernet interface supported on your storage system, or any 10Base-T/100Base-TX Ethernet controller.

Note

Do not mix interfaces of different speeds or media in the same multimode vif.

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

The vif command syntax

The `vif` command syntax is as follows:

```
vif create [single|multi] vif_name -b [rr|mac|ip] [interface_list]
vif {favor|nofavor} interface
vif add vif_name interface_list
vif delete vif_name interface
```

```
vif destroy vif_name
vif status [vif_name]
vif stat vif_name [interval]
```

For detailed information about the `vif` command and all the options available with this command, see the `na_vif(1)` man page.

Persistence of the vif command

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

- ◆ `vif create`
- ◆ `vif add`
- ◆ `vif delete`
- ◆ `vif destroy`
- ◆ `vif favor`
- ◆ `vif nofavor`

For detailed information

For detailed information about how to perform specific tasks using the `vif` command, see the following topics:

- ◆ [“Creating a single-mode vif”](#) on page 161
- ◆ [“Selecting an active interface in a single-mode vif”](#) on page 163
- ◆ [“Creating a multimode vif”](#) on page 165
- ◆ [“Adding interfaces to a vif”](#) on page 167
- ◆ [“Deleting an interface from a vif”](#) on page 168
- ◆ [“Displaying the status of a vif”](#) on page 169
- ◆ [“Displaying statistics of a vif”](#) on page 173
- ◆ [“Destroying a vif”](#) on page 174

Creating a single-mode vif

About creating a single-mode vif

This procedure enables you to create a single-mode vif—in which only one interface is active at a time and the others are ready to take over if the active interface fails. If you want a specific interface in a vif to be active, you need to specify that interface as preferred, otherwise an interface in the vif is randomly selected to be the active interface. For more information, see [“Selecting an active interface in a single-mode vif”](#) on page 163.

Prerequisites

You need to meet the following prerequisites to create a single-mode vif:

- ◆ Decide on a case-sensitive name for the vif 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 a vif.
- ◆ Decide on a list of the interfaces you want to combine into the vif.
- ◆ Configure all interfaces that will be included in the vif to be down using the `ifconfig` command.

Creating a single-mode vif

To create a vif in which only one interface is active at a time, complete the following steps.

Note

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

Step	Action
1	Enter the following command: <pre>vif create single vif_name [interface_list]</pre> <i>vif_name</i> is the name of the vif. <i>interface_list</i> is a list of the interfaces you want the vif to consist of. Note_____ You must ensure that all interfaces to be included in the vif are configured down. You can use the <code>ifconfig</code> command to configure an interface down. _____ Example: You can create a single-mode vif with the following command: <pre>vif create single SingleTrunk1 e0 e1</pre>
2	Enter the following command: <pre>ifconfig vifname IP_address netmask mask</pre> <i>vifname</i> is the name of the vif. <i>IP_address</i> is the IP address for this interface. <i>mask</i> is the network mask for this interface. Example: You can configure an IP address of 10.120.5.74 and a netmask of 255.255.255.0 on the single-mode vif SingleTrunk1, created in the previous step, with the following command: <pre>ifconfig SingleTrunk1 10.120.5.74 netmask 255.255.255.0</pre>

Selecting an active interface in a single-mode vif

About selecting an active interface

When you create a single-mode vif, by default, an interface is selected randomly to be the active interface. However, if you want to specify another interface as active, you can use the `vif favor` command to override the random selection. Additionally, if you want to specify an interface *not* to be considered when random selection is made, you can use the `vif nofavor` command.

The active interface is also known as a preferred interface. There can be only one active interface in a single-mode vif.

For example, you might want to select an interface over another when you add a new, higher speed or higher bandwidth interface to the vif and want this new interface to be the preferred interface.

The interface that you designate as the one not to be considered during random selection is known as the “not favored” interface.

Selecting an active interface

To change the active interface in a single-mode vif, complete the following step.

Note_____

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

Step	Action
1	Enter the following command: <pre>vif favor interface</pre> <i>interface</i> is the name of the interface you want to be active. Example: You can specify the interface <code>e1</code> to be preferred with the following command: <pre>vif favor e1</pre>

About designating an interface “not favored”

The interface marked as “not favored” (using the `vif nofavor` command) can still become the active interface when all other interfaces in a single-mode vif have failed. Even after other interfaces come back up, a “not favored” interface continues to stay active until it fails or until you, the system administrator, change the active interface using the `vif favor` command.

Designating an interface as “not favored”

To designate an interface as “not favored” so it will not be considered during random selection for an active interface in a single-mode vif, complete the following step.

Note—
The operation performed using the `vif nofavor` command is not persistent across reboots unless the command is added to the `/etc/rc` file.

Step	Action
1	Enter the following command: <pre>vif nofavor interface</pre> <i>interface</i> is the name of the interface you don’t want to be considered during random selection for an active interface. Example: You can specify the interface e2 to be “not favored” with the following command: <pre>vif nofavor e2</pre>

Creating a multimode vif

About creating a multimode vif

This procedure enables you to create a multimode vif on your storage system. By default, the IP-address based load-balancing method is used for a multimode vif. However, you can select another method while creating the vif. Once a load-balancing method has been assigned to a vif, it cannot be changed.

Note

Do not select the MAC-address based load-balancing method when creating vifs on a storage system that connects directly to a router. In such a setup, for every outgoing IP frame, the destination MAC address will be the MAC address of the router. As a result, only one interface of the vif will be used.

For more information about load-balancing methods available for multimode vifs, see [“Load balancing in multimode vifs”](#) on page 158.

Prerequisites

You need to meet the following prerequisites to create a multimode vif:

- ◆ 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 vif 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 a vif.
- ◆ Decide on the interfaces you want the vif to consist of.
- ◆ Configure all interfaces that will be included in the vif to be down using the `ifconfig` command.

Creating a multimode vif

To create a vif in which all interfaces are active at once, complete the following steps.

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

Step	Action
1	Enter the following command: <pre>vif create multi vif_name -b {rr mac ip} [interface_list]</pre> -b specifies the type of load-balancing method. ◆ rr—Round robin ◆ mac—MAC-address based ◆ ip—IP-address based (default) <i>vif_name</i> is the name of the vif. <i>interface_list</i> is a list of the interfaces that make up the vif. Note—You must ensure that all interfaces to be included in the vif are configured down. You can use the <code>ifconfig</code> command to configure an interface down. Example: You can create a multimode vif, comprising interfaces e0, e1, e2, and e3 and using MAC-based load balancing, with the following command: <pre>vif create multi MultiTrunk1 -b mac e0 e1 e2 e3</pre>
2	Enter the following command: <pre>ifconfig vifname IP_address netmask mask</pre> <i>vifname</i> is the name of the vif. <i>IP_address</i> is the IP address for this interface. <i>mask</i> is the network mask for this interface.

Managing vifs

Adding interfaces to a vif

About adding interfaces

This procedure enables you to add one or more interfaces to a vif. You can add physical interfaces to a vif any time after you create it.

Requirement before adding interfaces

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 vif must be configured down using the `ifconfig` command.

Adding interfaces to a vif

To add one or more interfaces to a vif, complete the following step.

Note

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

Step	Action
1	Enter the following command: vif add vif_name interface_list <i>vif_name</i> is the name of a previously configured vif. <i>interface_list</i> is a list of the interfaces you want to add to the vif. Example: You can add the interface e4 to the multimode vif MultiTrunk1 with the following command: vif add MultiTrunk1 e4

Managing vifs

Deleting an interface from a vif

About deleting an interface from a vif

This procedure enables you to delete an interface from a vif. The vif must be configured down before you delete its interface.

Deleting interface of a vif

To delete an interface of a vif, complete the following steps.

Note_____

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

Step	Action
1	Bring the vif down by entering the following command: <code>ifconfig vif_name down</code> <i>vif_name</i> is the name of the vif you want to bring down.
2	Enter the following command: <code>vif delete vif_name interface</code> <i>vif_name</i> is the name of a vif. <i>interface</i> is the interface of the vif you want to delete. Example: You can delete the interface e4 from a multimode vif MultiTrunk1 with the following commands: <code>ifconfig MultiTrunk1 down</code> <code>vif delete MultiTrunk1 e4</code>

Managing vifs

Displaying the status of a vif

Displaying vif status

You can display the current status of a specified vif or all single-mode and multimode vifs on your storage system.

To display the status of a vif, complete the following step.

Step	Action
1	Enter the following command: vif status [vif_name] <i>vif_name</i> is the name of the vif whose status you want to display. If you don't specify the vif name, the status of all vifs is displayed.

Example of displaying vif status

The following example displays the status of vif1 on your storage system called toaster:

```
toaster> vif status vif1
default: transmit 'IP Load balancing', VIF Type 'multi_mode', fail
'log'
vif1: 1 link, transmit 'none', VIF Type 'single_mode' fail
'default'

VIF Status      Up      Addr_set
up:
e10: state up, since 05Oct2001 17:17:15 (05:23:05)
      mediatype: auto-1000t-fd-up
      flags: enabled
      input packets 20, input bytes 1280
      output packets 2, output bytes 84
      up indications 1, broken indications 0
      drops (if) 0, drops (link) 0
      indication: up at boot
              consecutive 3, transitions 1
down:
e5: state down, since 05Oct2001 17:17:03 (05:22:00)
      mediatype: auto-unknown-cfg_down
      flags: disabled
```

input packets 0, input bytes 0
output packets 0, output bytes 0
up indications 0, broken indications 0
drops (if) 0, drops (link) 0
indication: down at boot
consecutive 4, transitions 0

The following table describes the console output.

Field name	Sub-field name	Description	Value
default		Indicates the default values for fields such as transmit, VIF Type, and fail. These values apply if no values are specified for these fields when a vif is created.	
	transmit	Indicates the default load-balancing method.	IP Load balancing
	VIF Type	Indicates the default vif type.	multi_mode
	fail	Indicates the default location where the errors will be logged.	log (that is, system log)

Field name	Sub-field name	Description	Value
vif1		Indicates that the data following this field pertains to the vif vif1.	
	transmit	Indicates the load-balancing method used.	none (load-balancing methods are used for multimode vifs only)
	VIF type	Indicates the type of vif1.	single_mode
	fail	Indicates the location where errors will be logged for vif1.	default (system log)
	VIF Status	Indicates the current status of vif1.	Up
	Addr_set	Indicates that a MAC address has been configured for vif1 and all of its interfaces.	
	up	Indicates that the interface following this sub-field is up. In this example, because vif1 is a single-mode vif, e10 is up while e5 is down.	

Field name	Sub-field name	Description	Value
	state	Indicates the current link-state of the interface.	up
	since	Indicates the date, time, and number of hours since the interface has been up.	05Oct2001 17:17:15 (05:23:05)
	flags	Indicates that the interface is enabled to send and receive data (enabled).	enabled
	consecutive	Indicates the number of consecutively received Up or Broken indications from the switch and link interaction.	3
	transitions	Indicates the number of indications received that caused a state transition from Up to Broken or Down to Up.	1
	down	Indicates that the interface following this sub-field is down. In this example, e5 is the standby interface for the single-mode vif, vif1.	

For more information about the `vif status` command, see the `na_vif(1)` man page.

Managing vifs

Displaying statistics of a vif

Displaying vif statistics

You display statistics dynamically for a specific vif or for all vifs.

To display statistics, complete the following step.

Step	Action
1	Enter the following command: vif stat [vif_name] [interval] <i>vif_name</i> is the name of the vif. If you don't specify a vif, the status of all vifs is displayed. <i>interval</i> is the interval, in seconds. The default is one second.

Example of displaying vif statistics

The following example displays output of the `vif stat` command:

```
vif stat vif0
```

```
vif (trunk) vif0
      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 `vif stat` command was run, and the following rows show the total number of packets received and sent per second thereafter.

Managing vifs

Destroying a vif

About destroying a vif You destroy a vif when you no longer need it or when you want to use the interfaces that form the vif for other purposes. After you complete this procedure, the links in the vif act individually rather than as an aggregate.

Destroying a vif To destroy a vif, complete the following steps.

Note_____

The operation performed using the `vif destroy` command is not persistent across reboots. If you want to destroy a vif permanently, make sure that the `vif create` commands corresponding to this vif do not exist in the `/etc/rc` file.

Step	Action
1	Configure the vif down by entering the following command: <code>ifconfig vif_name down</code> <i>vif_name</i> is the name of the vif you want to bring down.
2	Enter the following command: <code>vif destroy vif_name</code> <i>vif_name</i> is the name of the vif to destroy.

Second-level vifs

About second-level vifs

You group multiple multimode vifs to obtain a second layer of vif called the second-level vif.

Second-level vifs enable you to provide a standby multimode vif in case the primary multimode vif fails. You can use second-level vifs on a single storage system or in a cluster.

For detailed information

For detailed information about second-level vifs and how to create them on a single storage system and in a cluster, see the following topics:

- ◆ [“Understanding second-level vifs on a single storage system”](#) on page 176
- ◆ [“Creating a second-level vif on a single storage system”](#) on page 177
- ◆ [“Understanding second-level vifs in a cluster”](#) on page 179
- ◆ [“Creating a second-level vif in a cluster”](#) on page 181

Understanding second-level vifs on a single storage system

About second-level vifs on a single storage system


You use a second-level vif on a single storage system to provide a standby multimode vif in case the primary vif fails. You can provide additional redundancy by using two switches configured for multiple-port connections and four or more interfaces on your storage system.

Example of a second-level vif on a single storage system

You can set up your storage system with two two-link multimode vifs. Each vif is connected to a different switch capable of link aggregation over multiple ports. Next, you can set up a second-level single-mode vif that contains both of the multimode vifs.

When you configure the second-level vif using the `vif create` command, only one of the two multimode vifs is brought up as the active link. If all the underlying interfaces in the active vif fail, the second-level vif activates the link corresponding to the other vif.

In the following illustration, Secondlev is the single-mode second-level vif comprising the Firstlev1 and Firstlev2 vifs. Firstlev1 is initially the active interface; if Switch 1 drops both links, Switch 2 and Firstlev2 take over and maintain the connection to the network. For information about the commands to use to create the vif shown in this example, see [“Example of creating a second-level vif on a single storage system”](#) on page 178.



Creating a second-level vif on a single storage system

Assumptions made in this procedure

The following procedure assumes that you want to create a second-level vif, called *vif_name*, on a single storage system with two multimode vifs, called *vif_name1* and *vif_name2*. The *vif_name1* vif is composed of two physical interfaces, *if1* and *if2*, and *vif_name2* is composed of two physical interfaces, *if3* and *if4*.

By default, IP-based load balancing will be used for the multimode vifs created in this procedure.

Prerequisites

You need to meet the following prerequisites to create a second-level vif:

- ◆ 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 each vif 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 a vif.
- ◆ Decide on a list of the interfaces you want the vif to consist of.
- ◆ Configure all interfaces that will be included in the vif to be down using the `ifconfig` command.

Creating a second-level vif on a single storage system

To create a second-level vif on a single storage system, complete the following steps.

Note_____

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

Step	Action
1	Enter the following commands to create two multimode interfaces: <pre>vif create multi -b {rr mac ip} vif_name1 if1 if2 vif create multi -b {rr mac ip} vif_name2 if3 if4</pre> -b specifies the type of load-balancing method. ◆ rr—Round robin ◆ mac—MAC-address based ◆ ip—IP-address based (default) Note_____ You must ensure that all interfaces to be included in the vif are configured down. You can use the <code>ifconfig</code> command to configure an interface down.
2	Enter the following command to create a single-mode interface from the multimode interfaces: <pre>vif create single vif_name vif_name1 vif_name2</pre>

Example of creating a second-level vif on a single storage system

The following commands create the second-level vif shown in “[Example of a second-level vif on a single storage system](#)” on page 176. In this example, IP-based load- balancing is used for the multimode vifs.

```
vif create multi Firstlev1 e0 e1  
vif create multi Firstlev2 e2 e3  
vif create single Secondlev Firstlev1 Firstlev2
```

Understanding second-level vifs in a cluster

Advantage of second-level vifs in a cluster


In a cluster configuration, you can access data from both storage systems even if one of the storage systems in the cluster fails. In a second-level vif connected in a single-mode configuration, you can maintain connectivity to your storage system even if one of the switches fails. Thus, by using the two configurations together, you can achieve a fully redundant storage system connectivity architecture.

Normal cluster operation with second-level vifs

The following figure shows second-level vifs in a cluster. 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.




Switch failure in a cluster with second-level vifs

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 a cluster. Firstlev2 takes over the MAC address of Firstlev1 and maintains the connectivity through Switch 2.



Creating a second-level vif in a cluster

**Assumptions made
in this procedure**

The following procedure assumes that you want to create two second-level vifs, *secondlev1* and *secondlev2*, on clustered storage systems, StorageSystem 1 and StorageSystem 2. StorageSystem 1 and StorageSystem 2 are configured as shown in the following table.

Storage System	Multimode vifs	Interfaces
StorageSystem 1	vif_name1	if1 if2
	vif_name2	if3 if4
StorageSystem 2	vif_name3	if5 if6
	vif_name4	if7 if8

Creating a second-level vif in a cluster

To create a second-level vif in a cluster, complete the following steps.

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

Step	Action
1	Enter the following commands on StorageSystem 1 to create two multimode vifs: <pre>vif create multi -b {rr mac ip} vif_name1 if1 if2 vif create multi -b {rr mac ip} vif_name2 if3 if4</pre> -b specifies the type of load-balancing method. ◆ rr—Round robin ◆ mac—MAC-address based ◆ ip—IP-address based (default) Note— You must ensure that all interfaces to be included in the vif are configured to be down. You can use the <code>ifconfig</code> command to configure an interface down.
2	Enter the following command on StorageSystem 1 to create a second-level interface from the multimode vifs: <pre>vif create single secondlev1 vif_name1 vif_name2</pre>
3	Enter the following commands on StorageSystem 2 to create two multimode vifs: <pre>vif create multi -b {rr mac ip} vif_name3 if5 if6 vif create multi -b {rr mac ip} vif_name4 if7 if8</pre>
4	Enter the following command on StorageSystem 2 to create a second-level interface from the multimode vifs: <pre>vif create single secondlev2 vif_name3 vif_name4</pre>

Step	Action
5	Enter the following command on StorageSystem 1 to configure the second-level vifs for takeover: <pre>ifconfig secondlev1 partner secondlev2</pre> Note_____ In this command, <i>secondlev1</i> and <i>secondlev2</i> (arguments to the partner option) should be the interface names and not the interface IP addresses.
6	Enter the following command on StorageSystem 2 to configure the second-level vifs for takeover: <pre>ifconfig secondlev2 partner secondlev1</pre> Note_____ In this command, <i>secondlev1</i> and <i>secondlev2</i> (arguments to the partner option) should be the interface names and not the interface IP addresses.

Example of creating a second-level vif in a cluster

The following commands create the second-level vif in the cluster shown in “[Normal cluster operation with second-level vifs](#)” on page 179. In this example, IP-based load balancing is used for the multimode vifs.

On StorageSystem 1:

```
vif create multi Firstlev1 e1 e2
vif create multi Firstlev2 e3 e4
vif create single Secondlev1 Firstlev1 Firstlev2
```

On StorageSystem 2:

```
vif create multi Firstlev3 e5 e6
vif create multi Firstlev4 e7 e8
vif create single Secondlev2 Firstlev3 Firstlev4
```

On StorageSystem 1:

```
ifconfig Secondlev1 partner Secondlev2
```

On StorageSystem 2:

```
ifconfig Secondlev2 partner Secondlev1
```


About this chapter

This chapter explains how to set up and manage the Internet Protocol Security (IPsec) suite of protocols to secure information within a network. IPsec enables authentication and encryption of data in transition between your storage system and its Solaris and Windows 2000 or higher clients, or between two storage systems.

Topics in this chapter

This chapter discusses the following topics:

- ◆ [“Understanding IPsec”](#) on page 186
- ◆ [“Setting up IPsec”](#) on page 190
- ◆ [“Managing security policies”](#) on page 192
- ◆ [“Viewing security associations”](#) on page 198

Understanding IPsec

What IPsec is

IPsec is a security protocol suite that protects data from unauthorized disclosure when it is being transmitted between storage systems and clients. Using IPsec, you can add policies on your storage system that do both of these things:

- ◆ Configure encryption and authentication algorithms between your storage system and client.
Policies can be configured from your storage system to the client and from the client to your storage system over a range of IP addresses and ports.
- ◆ Negotiate a security association (SA) between the two end-stations (systems that initiate and receive secure communications). The SA is used for secure data exchanges between your storage system and the client.

Note

Maintaining security is processor intensive. However, storage system security processing can be offloaded from the CPU to an IPsec accelerator adapter. For more information, see “[L Larval \(uninitiated\)](#)” on page 198.

About security associations

A security association (SA) is an authenticated simplex (uni-directional) data connection between two end-stations. Security configurations are typically configured in pairs. An SA has all of the following:

- ◆ A unique Security Parameter Index (SPI) number
- ◆ An IP destination address
- ◆ An IPsec security protocol

The IPsec security protocol must be either of the following:

- ❖ Authentication Header (AH)

The AH protocol inserts an authentication header into each packet before the data payload. The authentication header includes a checksum created with a cryptographic hash algorithm, either Message Digest function 95 (MD5 - 128 bit key) or Secure Hash Algorithm (SHA - 160 bit key). The AH protocol does not alter the packet's data payload.

- ❖ Encapsulating Security Payload (ESP)

The ESP protocol inserts a header before the data payload and a trailer after it. When you specify an encryption algorithm, either Data Encryption Standard (DES) or triple DES, ESP alters the data payload by encrypting it. Alternatively, you can specify packet authentication

using the same MD5 or SHA-1 algorithms that are available with the AH protocol. If you use the ESP security protocol, you need to specify either authentication or encryption, or both.

Note

When you specify the AH protocol, only packet authentication (providing data integrity) is enabled. When you specify the ESP protocol, both packet authentication and packet encryption (providing data privacy) can be enabled.

At least two security associations, inbound and outbound, are required between end-stations. Security associations are stored in the Security Association Database (SAD) when IPsec is enabled on an end-station.

Security associations are created from security policies.

About security policies

Security associations are created based on information collected in security policies, which determine how security is handled in a transfer of information. Security policies can include any of the following types of specifications:

- ◆ The source and destination addresses (or ranges of addresses) of the end-stations (storage system and client)
- ◆ Packet authentication methods
- ◆ Packet encryption methods
- ◆ Restrictions on ports and services
- ◆ Whether inbound and outbound SAs are mirrored
- ◆ Strictness of policy application

Security policies are stored in the Security Policy Database (SPD) when IPsec is enabled on an end-station. Matching security policies must be configured on your storage system and clients.

About key exchange

An IPsec SA is negotiated by means of the key management protocol IKE (Internet Key Exchange). Phase 1 of an IKE key exchange authenticates the identity of the end-stations, which allows the establishment of an IPsec SA in Phase 2.

Two key exchange mechanisms using IKE are supported between storage systems and clients: Kerberos and preshared keys. However, only preshared keys can be used between two storage systems.

- ◆ Kerberos is a network authentication system in which end-stations prove their identity by obtaining identical secret keys from a Key Distribution Center (KDC), the Kerberos security server. For Windows 2000 and later, the KDC is located on the Windows domain controller, which processes IKE authentication requests for storage systems and Windows clients in the domain.

Kerberos authentication is enabled automatically when CIFS is licensed and configured on your storage system.

- ◆ Preshared keys are identical ASCII text strings entered manually on each end-station. Authentication is validated when IKE successfully compares the hash value of the two keys. Preshared key configuration is simple, but it requires manual management on each end-station. Also, preshared keys are static and persistent, therefore vulnerable unless changed frequently.

Note

The authentication of end-station identity provided by the key exchange protocol IKE is different from the packet integrity authentication provided by the IPsec protocols AH and ESP.

About the Data ONTAP IPsec implementation

The IPsec implementation for Data ONTAP conforms to the Internet Engineering Task Force (IETF) Security Architecture for the Internet Protocol (RFC 2401) and related protocols. The following restrictions apply:

- ◆ By default, storage systems obey all IPsec parameters that are configured on clients.

The only exception is Perfect Forward Secrecy (PFS), which is not supported on storage systems.

- ◆ Only transport mode is supported on storage systems; tunnel mode is not supported.

Consequently, IPsec is supported for security associations between storage systems and clients, but it is not supported for security associations between storage systems and security gateways.

- ◆ Only clients running Solaris or Windows 2000 or later are supported for IPsec connections.
- ◆ The following authentication mechanisms are supported:
 - ❖ For Solaris—only preshared keys authentication
 - ❖ For Windows—preshared keys and Kerberos authentication; however, Kerberos authentication is available only for Windows Domains, not Windows Workgroups
 - ❖ Between storage systems—only preshared keys authentication

- ◆ Data ONTAP supports preshared keys and Kerberos key exchange mechanisms, but it cannot be configured to use a specific mechanism. Instead, Data ONTAP relies on the client to specify which key exchange mechanism to use.
- ◆ You cannot configure parameters associated with SA, for example, how long the SA is valid, how many bytes of data can pass through the SA, in Data ONTAP. Instead, Data ONTAP uses the parameters that the client provides.

For more information about implementation and standards, see the `na_ipsec(1)` man page.

IPsec in a cluster configuration

The IPsec protocol, by its nature, does not work well in a failover environment, that is, an environment in which one storage system in a cluster configuration must take over the other storage system. This is because security policies, but not security associations, are taken over from the failed storage system. Clients will continue to send packets to the failed client for the remainder of the client security association lifetime, after which a new security association must be renegotiated and dropped packets resent.

For this reason, you are advised to reduce the security association lifetime to a minimum value to optimize IPsec operation in a cluster configuration. This minimizes the time clients use to destroy their security associations and negotiate new ones with the storage system that took over.

Note

You set the value of the security association's lifetime on clients rather than on your storage system.

IPsec in a vFiler unit configuration

IPsec can be enabled on a per-vFiler unit basis, with distinct security policies for each vFiler unit. IPsec configuration is preserved when vFiler units are moved from one hosting storage system to another, unless the vFiler unit's IP address is changed.

IPsec configuration can be set within the context of a vFiler unit or at your storage system command line by using the `vfiler run` command.

Note

Policies and configurations discussed in this chapter must be set individually for each vFiler unit.

Setting up IPsec

Preparing to use IPsec

Before you can use IPsec, you must take both of these actions:

- ◆ Select and configure one of the following key-exchange mechanisms.
 - ❖ Kerberos
 - ❖ Preshared keys
- ◆ Enable IPsec functionality on your storage system.

Then do these things:

- ◆ Create security policies as described in “[Managing security policies](#)” on page 192.
- ◆ View security associations as described in “[Viewing security associations](#)” on page 198.

Configuring Kerberos

Kerberos support is enabled by default on storage systems when CIFS is licensed and configured for Windows domain authentication.

Kerberos support for Windows clients requires all of the following:

- ◆ A Windows 2000 or greater client that is a member of a domain
- ◆ Kerberos selected in the client’s Authentication Methods list
- ◆ A functioning Key Distribution Center (KDC) on an accessible domain controller

Configuring preshared keys

To configure preshared keys, you must create an ASCII text string and store it on your storage system and the client that will be sharing the secure connection.

To create and store the preshared key on your storage system, complete the following steps.

Step	Action
1	Create a file named psk.txt file in the /etc directory.
2	Decide upon an ASCII text key that you will use for authenticating client and storage system.

Step	Action
3	In the <code>psk.txt</code> file, enter a line using the following format: <pre><i>ip_address key</i></pre> <i>ip_address</i> is the IP address of the client. <i>key</i> is the preshared key you decided upon. Example: <code>172.25.102.81 ag8key</code> See the <code>na_psk.txt(5)</code> man page for more information.

The same preshared key must be entered on the client when you configure a policy using the Windows user interface.

Enabling or disabling IPsec

To enable or disable IPsec on your storage system, complete the following step.

Step	Action
1	At your storage system command line, enter the following command: <pre>options ip.ipsec.enable on off</pre> <code>on</code> enables IPsec. <code>off</code> disables IPsec.

Managing security policies

About the ipsec command

Security policies in the SPD can be added, modified, displayed, deleted, and monitored using the `ipsec` command. For more information, see the `na_ipsec(1)` man page.

Selecting security policy options

When you create security policies, you must select from the following required and optional parameters on your storage system. Corresponding values must also be selected on any Windows clients served by the storage system.

Parameter	Options	Description
source and destination address	-s and -t	Required. Addresses can have any of the following forms: ◆ A single IP address ◆ A range of addresses ◆ An IP address at a specific port ◆ A range of addresses at a specific port
security protocol	-p	Required. Must be either Authentication Header (AH) or Encapsulated Security Payload (ESP); see “ About security associations ” on page 186.
encryption	-e	Optional. If the ESP protocol is selected, DES, triple DES, or no encryption can be specified.
authentication	-a	Required for AH protocol, optional for ESP protocol. SHA-1, MD5, or no authentication can be specified.
direction	-d	Required. Specifies an inbound or outbound connection relative to your storage system. By default, a mirrored policy (with the same parameters, except direction) is created unless mirroring is turned off.

Parameter	Options	Description
protocol	-f	Optional. Specifies an upper-layer protocol by number.
permission level	-l	Optional. Traffic can be restricted or permitted if a valid SA is not available.
index	-i	Specifies an index in the Security Policy Database. The index is obtained by the <code>ipsec policy show</code> command.

Creating a security policy

To create a security policy, complete the following step.

Step	Action
1	Enter the following command: <pre>ipsec policy add [-s src_ip/prefixlen[port]] [-t dst_ip/prefixlen[port]] -p {esp ah none} [-e {des 3des null} -a {sha1 md5 null}] -d {in out} [-m] [-f ip_protocol] [-l {restrict permit}]</pre> The add options are described in “Selecting security policy options” on page 192. Additionally, see the <code>na_ipsec(1)</code> man page for details of these options.

Note

Ensure that policies match on the storage system and client (or group of clients) that are negotiating the secure connection.

Example: `ipsec policy add -s 10.56.18.5 -t 10.56.19.172/24[139] -p esp -e des -a ah -d in -l restrict`

For more information about policy options, see the `na_ipsec(1)` man page.

Displaying existing security policies

You can use the `ipsec policy show` command to display the contents of the Security Policies Database (SPD), either in its entirety or by combinations of these parameters:

- ◆ Source and destination addresses
- ◆ Security protocol (AH or ESP)
- ◆ Direction (relative to your storage system)
- ◆ Specifications of upper-level protocols

To display security policies, complete the following step.

Step	Action
1	At your storage system command line, enter the following command: <code>ipsec policy show [-s <i>src_ip</i>] [-t <i>dst_ip</i>] [-f <i>ip_protocol</i>] [-d {in out}] [-p {esp ah}]</code> The <code>show</code> options are described in “Selecting security policy options” on page 192. Additionally, see the <code>na_ipsec(1)</code> man page for details of these options.

Example:

The following example displays security policy information for the device that has a source IP address (`-s`) of 10.56.19.172:

```
ipsec policy show -s 10.56.19.172

Index IPAddress      /prefix/port/protocol Dir/Policy Alg/SecLevel
-----
1      10.56.19.172 / 0/    [any  ]/any -> in /IPSEC  esp/Default
```

Deleting a security policy

You can remove entries from the SPD by deleting any of the following:

- ◆ All entries
- ◆ Individual entries identified by SPD index number (displayed by the `ipsec policy show` command)
- ◆ Groups of entries identified by any of the following:
 - ❖ Source and destination addresses
 - ❖ Direction (relative to your storage system)
 - ❖ Mirror policy

To delete a security policy from your storage system, complete the following step.

Step	Action
1	At your storage system command line, enter the following command: <pre>ipsec policy delete all -i index [[-s src_ip -t dst_ip] -d {in out} [-m]]</pre> The delete options are described in “Selecting security policy options” on page 192. Additionally, see the <code>na_ipsec(1)</code> man page for details of these options.

You must delete the same policies from corresponding clients.

How to display IPsec statistics

You can use the `ipsec stats` command to verify IPsec configuration, monitor protocol processing, and display IPsec violations. The command displays the following statistics:

- ◆ Total number of IPsec packets processed inbound and outbound
- ◆ Total number of AH and ESP packets processed
- ◆ Total number of AH and ESP processing failures
- ◆ Total number of failures and successes of AH and ESP replay windows

The anti-replay service window protects against replay attacks.
- ◆ Transmit and receive violations, which might be any of the following:
 - ❖ Improper or missing policies
 - ❖ Improper or missing security associations
 - ❖ Successful and failed IKE exchanges

To display statistics about how IPsec is working, complete the following steps.

Step	Action
1	At your storage system command line, enter the following command: <pre>priv set advanced</pre> For more information about advanced privilege level, see the <code>na_priv(1)</code> man page.

Step	Action
2	Enter the following command: ipsec stats [-z] -z resets the statistics counter.
3	When you are finished viewing statistics, be sure to return to the normal administrative privilege level by entering the following command: priv set admin

Example:

The following example shows the statistics provided by the `ipsec stats` command in `priv set advanced` mode.

```
system1*> ipsec stats
ipsec:
    148460138 inbound packets processed successfully
    0 inbound packets violated process security policy
    983 inbound packets with no SA available
    0 invalid inbound packets
    0 inbound packets failed due to insufficient memory
    0 inbound packets failed getting SPI
    0 inbound packets failed on AH replay check
    0 inbound packets failed on ESP replay check
    143929988 inbound packets considered authentic
    0 inbound packets failed on authentication
    ESP input packets
        des      : 3886739
        3des     : 140043249
    AH input packets
        md5      : 4530150
    134002232 outbound packets processed successfully
    0 outbound packets violated process security policy
    0 outbound packets with no SP available
    11 outbound packets with no SA available
    0 invalid outbound packets
    0 outbound packets failed due to insufficient memory
    0 outbound packets with no route
    ESP output packets
        des      : 4571170
        3des     : 124667606
    AH output packets
```

```
md5      : 4763456
ike:
  IKE input packets
    Identity Protection      : 107
    Informational             : 3682
    Quick                     : 7310
  IKE output packets
    Identity Protection      : 108
    Informational             : 10
    Quick                     : 3663
```

Viewing security associations

Displaying security associations

You can use the `ipsec sa show` command to display any of the following:

- ◆ The entire contents of the Security Associations Database (SAD)
- ◆ An individual entry in the SAD identified by the Security Parameter Index (SPI)

To learn the SPI for a database entry, you must first display the entire contents of the SAD.

- ◆ A group of entries that include all of the following:
 - ❖ Source and destination addresses
 - ❖ Security protocol (AH or ESP)
 - ❖ Direction (relative to your storage system)
 - ❖ Upper-level protocols specified

To view the currently active security associations on your storage system, complete the following step.

Step	Action
1	At your storage system command line, enter the following command: <code>ipsec sa show [<i>spi</i> <i>options</i>]</code>

Example:

The following example displays security association information for the device that has a source IP address of 10.56.19.172:

```
ipsec sa show 1 -s 10.56.19.172 -p esp

Alg/State/Spi Current Bytes/CreateTime SrcIPAddr->DstIPAddr
-----
esp/M/0001388 0/20 Aug 2002 17:28:19 10.56.19.172->10.56.19.173
```

The values for state are:

```
M Mature and active
D Dead
d Dying
L Larval (uninitiated)
```

About this appendix This appendix describes the statistics displayed by the `ifstat` command for the network interfaces supported by Data ONTAP.

Topics in this appendix

This appendix discusses statistics for the following interfaces:

- ◆ “[Fast Ethernet interface statistics](#)” on page 200
- ◆ “[Gigabit Ethernet and Ethernet Controller IV interface statistics](#)” on page 204
- ◆ “[Statistics for FAS250/ FAS270 network interfaces](#)” on page 209
- ◆ “[FDDI interface statistics](#)” on page 213
- ◆ “[ATM interface statistics](#)” on page 217

Fast Ethernet interface statistics

RECEIVE section statistics

The following table describes the statistics in the RECEIVE section of the `ifstat` command output when you use the command on a Fast Ethernet interface, such as an X1001C or X1012C card.

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.
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”, “List overflows”, and “Bus overruns” statistics.
Multi/broadcast	Total number of multicast or broadcast packets received.
CRC errors	Number of Cyclic Redundancy Check (CRC) errors that occurred on the received packets due mainly to duplex mismatches.
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 a cluster configuration.
No buffers	Number of times the driver was unable to get a buffer from its buffer pool because the pool was empty.

Statistic	Meaning
Tag drop	Number of tagged frames dropped on an interface that is not configured to support VLAN tagging and receives tagged frames.
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.
List overflow	Number of frames dropped due to the unavailability of receive resources.
Bus overruns	Number of frames lost due to receive First In First Out (FIFO) overflows.
Runt frames	Number of runt frames received.
Long frames	Number of long frames received that exceeded the maximum Ethernet-specified size of 1,518 bytes.
Flow controls	Number of flow control frames received.

TRANSMIT section statistics

The following table describes the statistics in the TRANSMIT section of the ifstat command output when you use the command on a Fast 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.

Statistic	Meaning
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.
Queue overflows	Total number of frames dropped due to software queue overflow.
Max collisions	Total number of frames that were not transmitted because they encountered the maximum number of allowed collisions.
No buffers	Number of times the driver failed to allocate a buffer for the transmit packet.
Late collisions	Number of frames that were not transmitted because they encountered a collision outside the collision window.
Bus underruns	Number of times the transmitter aborted the frame to be transmitted because data arrived late from memory. These packets are retransmitted later.
Lost carriers	Number of frames that were transmitted by the device despite the deassertion of CRS during transmission.
Deferred	Number of frames that were deferred before transmission due to activity on the link.
Single collision	Number of transmitted frames that encountered one and only one collision.
Multiple collision	Number of transmitted frames that encountered more than one collision, but fewer than the maximum allowed collisions.
Flow controls	Number of flow control frames transmitted.

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 a Fast Ethernet interface.

Statistic	Meaning
Current state	The state of the link. It can be up, down, or enabling.
Up to downs	Number of times the link toggled between up (LINK_UP) and down (LINK_DOWN) states.
Speed	Current negotiated speed.
Duplex	Duplex of the link negotiated or set.
Flow control	Negotiated value of flow control if the interface is autonegotiable; otherwise, it is the configured setting.

Gigabit Ethernet and Ethernet Controller IV interface statistics

RECEIVE section statistics

The following table describes the statistics in the RECEIVE section of the `ifstat` command output when you use the command on a Gigabit Ethernet interface supported on the storage system or the onboard 10Base-T/100Base-TX Ethernet Controller IV.

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.
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 a cluster configuration.
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.

Statistic	Meaning
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.
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.
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 when you use the command on a Gigabit Ethernet interface supported on the storage system or the onboard 10Base-T/100Base-TX Ethernet Controller IV.

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.

Statistic	Meaning
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.
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 packet size when jumbo frames are enabled.

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 a Gigabit Ethernet interface supported on the storage system or the onboard 10Base-T/100Base-TX Ethernet Controller IV.

Statistic	Meaning
Current state	Current state of the interface: ◆ up or down—The state of the link. ◆ cfg_down—The interface is configured down. ◆ enabling—The interface is coming up.
Up to downs	Number of times the link toggled between up and down.
Auto	Operational state of autonegotiation: ◆ on—Autonegotiation is enabled and succeeded. ◆ off—Autonegotiation failed. This happens when the device to which the interface is connected has disabled autonegotiation or is incompatible with the interface. This may also indicate that the interface is down.
Speed	Speed of link negotiated or set.
Duplex	Duplex of the link negotiated or set.
Flow control	The operational flow control setting. For information on how the operational flow control setting is determined, see Chapter 1, “ Network Interface Configuration ,” on page 1.

Statistics for FAS250/ FAS270 network interfaces

RECEIVE section statistics

The following table describes the statistics in the RECEIVE section of the `ifstat` command output when you use the command on a FAS250/FAS270 network 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.
Multi/broadcast	Total number of multicast or broadcast packets received.
Total discards	Total number of “No buffers” packets that were discarded even though no errors were detected.
No buffers	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.
Non-primary u/c	Number of Ethernet frames received for the partner’s MAC address after a failover in a cluster configuration.
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.

Statistic	Meaning
Vlan untag drop	Number of untagged frames dropped on an interface that is configured to be part of a VLAN.
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.
CRC errors	Number of packets received with bad CRC.
Length errors	Number of frames received by the MAC where the actual number of bytes received did not match the length given in the Ethernet header.
Code errors	The number of frames received by the MAC that had a code error signaled by the Physical (PHY) layer.
Dribble errors	The number of frames received by the MAC with an alignment error. This is not used for 1000Mb/s operation.

TRANSMIT section statistics

The following table describes the statistics in the TRANSMIT section of the ifstat command output when you use the command on a FAS250/FAS270 network 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.

Statistic	Meaning
Multi/broadcast	Total number of multicast or broadcast packets transmitted.
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 overflow” statistics.
Queue overflow	Number of outgoing packets dropped because the driver’s queue was full. It might indicate a system problem.
No buffers	Number of times the driver fail to allocate a buffer for the transmit packet.
CRC errors	Number of packets transmitted by the MAC with CRC errors. This can happen only when the MAC is not appending the CRC to the transmitted packets.
Abort errors	Number of packets aborted during transmission. This could be because of a FIFO underrun.
Runt frames	Number of packets smaller than the minimum frame size (64 bytes) transmitted by the MAC.
Long frames	Number of packets larger than the maximum frame size transmitted by the MAC.
Single collision	Number of frames that encountered exactly one collision. Only valid in half-duplex mode.
Late collisions	Number of collisions that occurred outside the collision window. Only valid in half-duplex mode.
Deferred	Number of times a packet was aborted by the MAC due to excessive collisions during transmission. If 16 consecutive collisions occur during transmission of a packet, the transmission is deferred and MAC aborts the packet.

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 a FAS250/FAS270 network interface.

Statistic	Meaning
Current state	Current state of the interface: ◆ <code>up</code> or <code>down</code>—The state of the link. ◆ <code>cfg_down</code>—The interface is configured down. ◆ <code>enabling</code>—The interface is coming up.
Up to downs	Number of times the link toggled between up and down.
Speed	Speed of the link negotiated or set.
Duplex	Duplex of the link negotiated or set.
Flow Control	The operational flow control setting. For information on how the operational flow control setting is determined, see Chapter 1, “ Network Interface Configuration ,” on page 1.

FDDI interface statistics

RECEIVE section statistics

The following table describes the statistics in the RECEIVE section of the ifstat command output when you use the command on an FDDI interface.

Statistic	Meaning
Packets	Number of packets received on the interface.
Bytes	Number of bytes received on the interface.
Total Errors	Number of errors during reception, including all kinds of receive errors.
Queue full	Number of received packets dropped due to the unavailability of buffers.
H/w faults	Number of hardware faults.
SA matches	Number of Source Address (SA) matches.
DA nonmatches	Number of Destination Address (DA) nonmatches.
PDQ aborts	Number of PDQ aborts.
Invalid lengths	Number of invalid lengths.
Fragments	Number of fragments.
Format errs	Number of frames received that contained an odd number of symbols; that is, the number of nonintegral octets.
MAC resets	Number of MAC resets.
CRC Errors	Number of frames received with a bad frame check sequence.
Status Errors	Number of frames received that had the E (error detected) indicator set or a missing E indicator.
Long Frames	Number of frames received that were larger than the maximum legal size.
Short Frames	Number of short frames received.
Copied buffers	Number of copied buffers.

Statistic	Meaning
No buffers	Number of times there were insufficient buffers.
LLC errors	Number of LLC errors.
Blk check errs	Number of block check errors.
Length errors	Number of length errors.
Overruns	Number of times that the interface received a frame loss indication from the hardware on the interface.

TRANSMIT section statistics

The following table describes the statistics in the TRANSMIT section of the ifstat command output when you use the command on an FDDI interface.

Statistic	Meaning
Packets	Number of packets attempted to be transmitted.
Bytes	Number of bytes attempted to be transmitted.
Errors	Number of hardware errors encountered while attempting to transmit.
Queue full	Number of received packets dropped due to unavailability of buffers.
No buffers	Number of times buffers were not available.
No buf resets	Number of times the process was reset due to the unavailability of buffers.
No descriptors	Number of times descriptors were unavailable.
I/f down	Number of times the interface was down.
Underruns	Number of times an underrun occurred because the interface's transmit FIFO became empty due to insufficient memory. That is, the device tried to transmit but could not get data into memory fast enough.
Failures	Number of frames that were transmitted with an error.

DEVICE section statistics

The following table describes the statistics in the DEVICE section of the `ifstat` command output when you use the command on an FDDI interface.

Statistic	Meaning
Interrupts	Number of times the FDDI device interrupted the host.
Spurious ints	Number of spurious interrupts received from the device.
Resets	Number of times the device was reset due to unrecoverable hardware faults.
Non-existing mem	Number of times the device was asked to address nonexistent memory.
Parity errs	Number of parity errors detected on the interface.
Host parity errs	Number of parity errors detected on the host.
Spurious Unsols	Number of unknown unsolicited error events reported by this link entity.
Ring inits in	Number of ring claim processes initiated by this link entity.
Ring beacon out	Number of ring beacon processes initiated by this link entity.
Ring beacon in	Number of ring beacon processes recognized by this link entity.
Dupl. address	Number of times the duplicate address test failed.
Dupl. token	Number of times this link entity detected a duplicate token.
Ring purge errs	Number of times the MAC PurgeError Event bit was set.
FCI strip errs	Number of times a Frame Content Independent Strip operation was terminated.
Trace init	Number of PC traces initiated by this station.
Selftest failure	Number of failed self-tests during device resets.
PC trace rcvd	Number of PC traces received by this station.

Statistic	Meaning
Config changes	Number of configuration changes.
LEM rejects	Number of Link Error Monitor (LEM) rejects.
Elastic buf errs	Number of Elasticity Buffer and Link Monitor (ELM) Chip EBUFF_ERROR events.
Link test reject	Number of link test rejects.

ATM interface statistics

RECEIVE section statistics

The following table describes the statistics in the RECEIVE section of the `ifstat` command output when you use the command on an ATM interface.

Statistic	Meaning
Packets	Number of packets received on the interface.
Bytes	Number of bytes received on the interface.
Errors	Number of errors during reception, including all kinds of receive errors.
Queue full	Number of packets dropped because they could not be put in the transmit queue.
Collisions	Ignore this field.

TRANSMIT section statistics

The following table describes the statistics in the TRANSMIT section of the `ifstat` command output when you use the command on an ATM interface.

Statistic	Meaning
Packets	Number of packets attempted to be transmitted.
Bytes	Number of bytes attempted to be transmitted.
Errors	Number of hardware errors encountered while attempting to transmit.

About this appendix This appendix describes configuration procedures that might improve your storage system's performance.

Balance NFS traffic on network interfaces Attach multiple interfaces on your storage system to the same physical network to balance network traffic among different interfaces. For example, if two Ethernet interfaces on a storage system named toaster are attached to the same network where four NFS clients reside, specify in the `/etc/fstab` file on client1 and client2 that these clients mount from toaster-0:/home. Specify in the `/etc/fstab` file on client3 and client4 that these clients mount from toaster-1:/home. This scheme can balance the traffic among interfaces if all clients generate about the same amount of traffic.

Your storage system always responds to an NFS request by sending its reply on the interface on which the request was received.

Correct 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. 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 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.

Upgrade to a faster network interface You can increase storage system performance by upgrading to a faster network interface. The following lists network interfaces from the fastest to the slowest:

- ◆ Gigabit Ethernet interfaces
- ◆ ATM OC-12 interfaces
- ◆ ATM OC-3 interfaces
- ◆ Fast Ethernet 100Base-T interfaces
- ◆ FDDI interfaces

About this appendix This appendix describes the Data ONTAP services file that is available in the /etc directory. The /etc/services file is in the same format as its corresponding UNIX systems /etc/services file. Although this file is not used by Data ONTAP, it is provided in this appendix as information useful to system administrators.

Host identification Although some port scanners are able to identify storage systems as storage systems, others port scanners report storage systems as unknown types, UNIX systems because of their NFS support, or Windows systems because of their CIFS support. There are several services that are not currently listed in the /etc/services file.

Below is an example of a complete list of the file contents.

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
http	80/tcp	# HTTP license, FilerView, SecureAdmin
kerberos	88/udp	# Kerberos 5 - outbound only

Service	Port/ Protocol	Description
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
ftp-data	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
shell	514/tcp	# rsh, insecure remote command execution.
syslog	514/udp	# outbound only
route	520/udp	# for RIP routing protocol
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
ttcp	5001/udp	# unused, shouldn't be listed here.
ttcp	5001/tcp	# unused, shouldn't be listed here.
ndmp	10000/tcp	# for network backups
snapmirro	10566/tcp	# also SnapVault

Service	Port/ Protocol	Description
ndmp-local	32243/tcp	# Internal connection inside your storage system

/etc/services NNTP and TTCP ports

The `nntp` and `ttcp` ports are unused by your storage system and should never be detected by a port scanner.

Ports found in a block starting around 600

The following ports are found on the storage system with NFS enabled:

UDP	602	NFS mount daemon (<code>mountd</code>)
TCP	603	NFS mount daemon (<code>mountd</code>)
UDP	604	NFS status daemon (<code>statd</code> , <code>statmon</code>)
TCP	605	NFS status daemon (<code>statd</code> , <code>statmon</code>)
UDP	606	NFS lock manager (<code>lockd</code> , <code>nlockmgr</code>)
TCP	607	NFS lock manager (<code>lockd</code> , <code>nlockmgr</code>)
UDP	608	NFS quota daemon (<code>quotad</code> , <code>rquotad</code>)

On other systems, the ports appear as follows:

UDP	611	NFS mount daemon (<code>mountd</code>)
TCP	612	NFS mount daemon (<code>mountd</code>)
UDP	613	NFS status daemon (<code>statd</code> , <code>statmon</code>)
TCP	614	NFS status daemon (<code>statd</code> , <code>statmon</code>)
UDP	615	NFS lock manager (<code>lockd</code> , <code>nlockmgr</code>)
TCP	616	NFS lock manager (<code>lockd</code> , <code>nlockmgr</code>)

UDP 617 NFS quota daemon (quotad,
rquotad)

Enter the following command on UNIX systems to obtain the correct information by querying the port mapper on port 111:

toaster# rpcinfo -p storage.system.name.or.ip.address

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. Storage systems can have these services running on other port numbers. Because the system selects these port numbers at random when it boots, they are not listed in the /etc/services file.

Other ports not listed in /etc/services

The following ports appear in a port scan but are not listed in /etc/services file:

Protocol	Port	Service
TCP	22	SSH (SecureAdmin)
TCP	443	SSL (SecureAdmin)
TCP	3260	iSCSI-Target
UDP	xxxx	Legato ClientPack for your storage system runs on random UDP ports and is now deprecated. It is recommended that NDMP be used to back up your storage system using Legato Networker.

Note

Disable open ports that you do not need.

FTP

- ◆ ftp-data
- ◆ ftp

File transfer protocol (FTP) uses TCP ports 20 and 21. For a detailed description of the FTP support for your storage system, see the *Data ONTAP File Access and Protocols Management Guide*. 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.

SSH

- ◆ ssh

Secure Shell (SSH) protocol is a secure replacement for RSH and runs on TCP port 22. This only appears in a port scan if the SecureAdmin™ software is installed on your storage system.

There are three commonly deployed versions of the SSH protocol:

- ◆ SSH version 1—is much more 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, use the following CLI command:

```
secureadmin disable ssh
```

Telnet

- ◆ 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 not secure because:

- ◆ When users log into 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 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 the administrators use the same passwords on your storage system as they do on other network devices, the attacker can use these passwords to access those resources as well.

Note

To reduce the potential for attack, establish and enforce policies preventing administrators from using the same passwords on your storage system that they use for access to 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, but because a packet sniffing attack is easier, TCP session attacks are less common.

To disable Telnet, set options `telnet.enable` to off.

SMTP

- ◆ `smtp`

The 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

- ◆ `time`
- ◆ `ntp`

Your storage system supports two different time service protocols:

- ◆ TIME protocol (also known as rdate) is specified in the RFC 868 standard. This standard allows for time services to be provided on TCP or UDP port 37. Your storage system uses only UDP port 37.
- ◆ Simple network time protocol (NTP) is specified in the RFC 2030 standard and is provided only on UDP port 123.

When your storage system has option `timed.enable` set to On and a remote protocol (`rdate` or `ntp`) is specified, the storage system synchronizes to a network time server.

If the `timed.enable` option is set to Off, your storage system is unable to synchronize with the network time server using NTP. The `rdate` time protocol can still be used by manually issuing the `rdate` command from your storage system console.

You should set the `timed.enable` option to On in a cluster configuration.

DNS

- ◆ `domain`

The Domain Name Service (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. Your 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

- ◆ `dhcps`

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

- ◆ `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. Disable TFTP using the following option:

```
options tftpd.enable off
```

HTTP

- ◆ `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

- ◆ Files when the HTTP protocol 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.

The HTTP protocol 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 only authentication methods defined by the HTTP protocol 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

In versions of Data ONTAP earlier than 7.0, your storage system listens for new connections (by default, set to TCP port 80) even when the HTTP protocol is not licensed and FilerView is disabled. However, starting with Data ONTAP 7.0, 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 continue to listen for new connections.

Kerberos

- ◆ `kerberos`
- ◆ `kerberos-sec`

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

- ◆ portmap
- ◆ nfsd

The 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, use the `nfs off` command.

CIFS

- ◆ netbios-name
- ◆ netbios-dg
- ◆ netbios-ssn
- ◆ cifs-tcp

The 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 also must 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

◆ ssl

The Secure Sockets Layer (SSL) protocol provides encryption and authentication of TCP connections.

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 through a loopback connection to pass the traffic to FilerView, running on TCP port 80. This loopback connection is contained within your storage system and no unencrypted data is transmitted over the network.

TCP port 443 can be disabled using FilerView or with the following command:

```
secureadmin disable ssl
```

SNMP

◆ 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

- ◆ 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 the type of information that is available. These differences are not important from a security point of view. This version of SNMP is currently used on your storage systems.
- ❖ 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. Your storage system's SNMP implementation 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.

Set 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

◆ shell

Remote shell protocol (RSH) is used for remote command execution and is the only protocol supported on your storage system. It is even 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. If possible, RSH should be disabled 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

- ◆ `syslog`

Your storage system sends messages to hosts specified by the user in the `/etc/syslog.conf` file using the syslog protocol on UDP port 514. It does not listen on this port, nor does it act as a syslog server.

Routed

- ◆ `routed`

The route daemon, `routed`, listens on UDP port 520. It receives broadcast messages from routers or other hosts using the 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 sift 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

- ◆ `ndmp`
- ◆ `ndmp-local`

Network Data Management Protocol (NDMP) runs on TCP port 10000 and is used primarily for backup of network-attached storage (NAS) devices, such as your 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 systems support both the TEXT and MD5 authentication methods. Most NDMP-enabled backup software uses MD5 by default.

To entirely disable the TEXT authentication method, set the `ndmpd.authtype` option to `challenge`.

To restrict NDMP commands to certain authorized backup hosts, use the `ndmp.access` option.

Regardless of the authentication method used, NDMP sends backup data in unencrypted 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, set the `ndmp.enable` option to `off`.

SnapMirror and SnapVault

- ◆ `snapmirror`

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, configure the `snapmirror.access` or `snapvault.access` option. Once 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, set the `snapmirror.checkip.enable` option to `On`.

To disable SnapMirror, set the `snapmirror.enable` option to `Off`. To disable SnapVault, set the `snapvault.enable` option to `Off`.

About this appendix This appendix presents network error codes that are generated by the `netdiag` command.

Only a small fraction of the possible network error messages are presented in this appendix. If you receive any problem code not listed in this chapter, contact your technical support.

**Error code
descriptions and
recommended
actions**

The following table lists some network error codes, describes problems that the error codes point to, and suggests actions that you can take to fix the problems.

Error code	Description	Recommended actions
201	Link not detected.	Complete the following steps until you detect a link: 1. Ensure that a 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 up, 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.

Error code	Description	Recommended actions
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 up.	Configure the interface up.
205	Duplex mismatch.	Change the interface or peer switch port duplex setting so they match.
206	Link capacity problem.	Upgrade to a faster interface.
207	The interface is not transmitting or receiving.	Complete the following steps: 1. Pull the network cable out from the network interface card. 2. Reinsert the cable. 3. Use <code>ifstat</code> to display statistics and see Appendix A, “Network Interface Statistics,” on page 199 to determine the type of error. ❖ Link errors, such as CRC, are caused by a faulty switch port, cable, or NIC; replace them one-by-one to locate the fault. ❖ Out-of-resource errors are caused by heavy loads. 4. If the problem persists, contact your technical support.
208	Excessive I/O errors.	Complete the following steps: 1. Reseat the interface card. 2. Check the cables. 3. If the problem persists, contact your technical support.

Error code	Description	Recommended actions
209	Excessive unsupported protocol packets are being sent to your storage system.	The problem is not with your storage system. 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 using the <code>ifconfig</code> 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.
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.

Error code	Description	Recommended actions
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.
409	There are excessive embryonic TCP connection drops because the system is under security attack or because a client has a bug.	Contact your network administrator to resolve the problem. A packet trace might assist in locating the problem.
410	Excessive TCP checksum errors. These errors can be caused bad hardware on the client, in the network infrastructure (e.g., 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. ◆ Replace hardware components until the problem goes away. ◆ Contact your network administrator to resolve the problem.

Error code	Description	Recommended actions
411	There are packets because of a client. Your system might be under a security attack.	The problem is not with your storage system. ◆ Check your client system for bugs. ◆ 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.

Index

Symbols

/etc/dgateways file, deprecated 67
/etc/hosts file
 creation of 81
 resolving host names with 79
 updating of 80, 81
/etc/nsswitch.conf file 100, 101
/etc/rc file, default route 67
/etc/resolv.conf file, creating 85
/etc/services file 221

Numerics

100tx, mediatype 8
100tx-fd, mediatype 8

A

adapter failover, automatic 29
address
 IP address, configuring 12, 14
AH (Authentication Header), IPsec 186
aliases, configuring for an interface (ifconfig) 18
ATM and LANs, bridging between 27
ATM commands
 atm adinfo (verifies adapter operation) 33
 atm adstat (verifies connection works) 34
 atm atmarp (deletes incoming FORE/IP PVCs) 62
 atm atmarp (deletes outgoing FORE/IP PVCs) 62
 atm atmarp (displays FORE/IP PVC address resolution) 58
 atm atmarp (establishes incoming FORE/IP PVCs) 57
 atm atmarp (establishes outgoing FORE/IP PVCs) 57
 atm atmconfig (changes ATM AAL) 61
 atm atmconfig (changes SPANS AAL) 61
 atm atmconfig (displays configuration data) 59
 atm elconfig add (adds emulated LAN) 41

atm elconfig delete (deletes emulated LAN from adapter) 44
atm elconfig set (configures LANE configuration server) 37
atm elconfig show (verifies adapter configurations) 47
atm elconfig show (verifies elements of emulated LAN) 48
atm uniconfig set failover (modifies load balancing groups) 51
atm uniconfig show (verifies UNI operation) 35

ATM ELAN interface, frame size 5

ATM interface, statistics 217

ATM protocol

 automatic adapter failover of 29
 bridging between ATM and LANs 27
 BUS, description of 29
 cause codes 27
 cells, description of 26
 checking UNI operation (atm uniconfig show) 35
 configuring logical Ethernet interface (ifconfig) 43
 deleting emulated LAN from adapter (atm elconfig delete) 44
 description of 26
 differences between LANs and 27
 emulated LANs
 adding (atm elconfig add) 41
 components of 28
 deleting from adapter (atm elconfig delete) 44
 description of 28
 frame size 5
 saving host and IP address in 53
 verifying adapter configurations (atm elconfig show) 47
 verifying communications (ping) 46
 verifying elements of (atm elconfig show) 48
 establishing incoming FORE/IP PVCs (atm

- atmarp) 57
- establishing outgoing FORE/IP PVCs (atm atmarp) 57
- FORE/IP
 - changing ATM AAL (atm atmconfig) 61
 - deleting incoming PVCs (atm atmarp) 62
 - deleting outgoing PVCs (atm atmarp) 62
 - description of 26
 - displaying configuration data (atm atmconfig) 59
 - over SPANS, description of 54
 - PVCs, description of 56
 - PVCs, displaying address resolution (atm atmarp) 58
- LANE
 - Clients, description of 28
 - configuration server, configuring (atm elconfig set) 37
 - configuration server, description of 29
 - description of 26, 27, 28
 - handling addressing and resolution of 30
 - preparing ATM adapter to use 31
 - Server, description of 29
 - standards supported 30
- load balancing 29, 50
 - description of 50
 - modifying load balancing groups (atm uniconfig set failover) 51
- UNI 29
- PVCs
 - and SVCs, description of 54
 - description of 55
- saving configuration commands 52
- saving host and IP address 53
- SPANS
 - changing the SPANS AAL (atm atmconfig) 61
- UNI (User-Network Interface), description of 29
- VCCs (Virtual Channel Connections),
 - component of emulated LANs 28
- verifying a connection works (atm adstat) 34
- verifying adapter configurations in (atm elconfig show) 47
- verifying adapter operation (atm adinfo) 33

- verifying elements of emulated LAN (atm elconfig show) 48
- ways to use 26
- Authentication Header (AH), IPsec 186
- auto, mediatype 8
- automatic adapter failover 29

B

- boot
 - from diskette 84, 95, 116, 122
- bridging between ATM and LANs 27
- BUS, within an emulated LAN 29

C

- cause codes, ATM 27
- cf.takeover.on_network_ interface_failure option 17
- clusters
 - IPsec in 189
 - routing in 75
 - second-level vifs in 179, 182
 - SNMP in 114
 - with DNS name caching 86
- command, netdiag 23
- commands. *See*
 - NIS commands
 - vifs commands
 - VLAN commands
- configuration
 - of aliases 19
 - of IP addresses 12, 14
 - of Kerberos for IPsec 190
 - of LANE configuration server 37
 - of logical Ethernet interface 43
 - of network interfaces 12
 - of preshared keys 190
- configuring a switch for multimode vifs 156
- custom MIB 109
- MIB 109

D

- default route 67
- DELETE 14

deleting an interface in a vif 168

DEVICE statistics

on FDDI card 215

DHCP 228

DNS

about 228

changing domain name (options
dns.domainname) 85

configuring 85

dynamic updates

about 89, 90

changing the TTL of 91

enabling 91

enabling and disabling (options dns.enable) 86

managing with FilerView 83

name caching 86

DNS commands

dns flush 87

options dns.cache.enable 87

options dns.domainname (changes domain
name) 85

options dns.enable (enables and disables DNS)
86

Domain Name Service (DNS). *See* DNS, DNS
commands

domain names, changing of 85

duplex settings, correcting mismatches 219

Dynamic Host Configuration Protocol (DHCP)
228

E

Emulated LANs

adding (atm elconfig add) 41

and a LANE Client 29

ATM BUS, description of 29

components of 28

configuring frame size of 5

deleting from adapter 44

description of 28

saving host and IP address in 53

verifying communications (ping) 46

verifying elements of 48

Encapsulating Security Payload (ESP), IPsec 186

error codes, netdiag 23

error codes, network 235

error messages

network error codes 235

serious 235

ESP (Encapsulating Security Payload), IPsec 186

EtherChannel. *See* vifs

Ethernet interfaces, media types 8

F

failover

modifying load-balancing groups 51

of adapter 29

FAS250/FAS270 interfaces, statistics 209

fast path mechanism, description of 65

FDDI card, statistics 213

FilerView management

of /etc/hosts file 81

of DNS 83

of host name search order 101

of network interfaces 13

of NIS 95

of routing 71

of SNMP 115

firewall security 16

flags, in routing table 73

flow control on Gigabit Ethernet 10

FORE/IP

changing ATM AAL (atm atmconfig) 61

displaying configuration data (atm atmconfig)
59

over SPANS, description of 54

PVCs

deleting incoming (atm atmarp) 62

deleting outgoing (atm atmarp) 62

displaying address resolution (atm
atmarp) 58

establishing incoming (atm atmarp) 57

establishing outgoing (atm atmarp) 57

establishment of 56

frame size

ATM ELAN interface 5

default 5

definition of 5

FDDI interface 5

Gigabit Ethernet interface 5

G

Gigabit Ethernet

flow control, description of 10

Gigabit Ethernet interface, statistics 204

H

host names

changing search order for 100

for interfaces, description of 4

resolving 79

hosts.byaddr map 93

hosts.byname map 93

HTTP 229

Hypertext Transport Protocol (HTTP) 229

I

IEEE 802.3ad 155

ifconfig command

changing interface status 20

configuring aliases for an interface 19

configuring an IP address using 14

configuring logical Ethernet interfaces 43

negotiated failover option (nfo option) 17

network mask, configuring 15

nfo option 17

untrusted interface, configuring 16

ifstat command 217

IKE 187

interface

negotiated failover (nfo option) 17

trusted, setting 16

untrusted, setting 16

interfaces

alias

configuring (ifconfig) 19

description of 18

balancing NFS traffic 219

configuration 12

description of 2

Gigabit Ethernet flow control 10

host name creation, description of 4

jumbo frames

and MTU size 5

client-size recommendations 6

description of 5

ways to set up 6

managing with FilerView 13

media types on Ethernet 8

multiple ports, description of 3

naming conventions 3

numbering of 2

physical, adding (vif add) 167

selecting active vif 163

statistics for FAS250/FAS270 209

statistics for N3700 209

status of, changing 20

types of 2

VLAN supported in 139

Internet Key Exchange. *See* IKE

Internet Protocol Security. *See* IPsec

IP address, configuring 12, 14

IP ports 221

IP-address based load balancing 158

IPsec

Authentication Header (AH) 186

cluster configuration 189

description of 186, 188

disabling 191

enabling 191

Encapsulating Security Payload (ESP) 186

IKE 187

Kerberos 188

key exchange 187

Perfect Forward Secrecy (PFS) 188

preshared keys 188

Security Association (SA) 186

security policies 187, 192

setup 190

statistics 195

transport mode 188

tunnel mode, not supported 188

vFiler unit configuration 189

IPsec commands

ipsec 192

ipsec policy add 193

ipsec policy delete 194

- ipsec policy show 194
- ipsec sa show 198
- ipsec stats 195
- options ip.ipsec.enable 191

J

- jumbo frames
 - client configuration for 6
 - description of 5
 - setup 6
 - using for vifs 155, 163

K

- Kerberos
 - configuring for IPsec 190
 - key exchange, description of 188

L

- LACP 155

- LANE

- and Emulated LAN configuration information 28

- Client, description of 28
 - configuration server, configuring 37
 - configuration server, description of 29
 - description of 27
 - description of service 28
 - handling addressing and resolution of 30
 - preparing ATM adapter to use 31
 - Server, description of 29
 - service, description of 27
 - standards supported 30

- LANs, bridging between ATM and 27

- lifetime, Security Association (SA) 189

- Link Aggregation Control Protocol (LACP) 155

- Link aggregation. *See* vifs

- LINK INFO statistics

- on FAS250/FAS270 interfaces 212
 - on Fast Ethernet card 203
 - on Gigabit Ethernet interface 208
 - on N3700 interfaces 212

- load balancing methods 158

M

- MAC-address based load balancing 158

- media type, autonegotiate 8

- media types, Ethernet 8

- MTU size, definition of 5

- multimode vifs, creating (vif create multi) 166, 177

- multiple ports on interfaces, description of 3

- MultiStore. *See* vFiler units

N

- N3700 interfaces, statistics 209

- name caching, DNS

- description of 86

- enabling 87

- flushing 87

- in clusters 86

- name resolution, NIS and DNS configuration files 77

- negotiated failover, specifying 17

- netdiag, command 23

- netstat command

- output flags 73

- network error codes 235

- network interfaces

- configuring logical Ethernet 43

- IP address, configuring 14

- negotiated failover (nfo option) 17

- network mask, configuring 15

- statistics, displaying (ifstat) 21, 22

- storage system supported 2

- virtual 2

- network mask, configuring 15

- network time protocol (NTP) 227

- network, VLAN 134

- nfo option 17

- NFS protocol

- balancing traffic 219

- over-UDP routing, description of 65

- NIS

- changing NIS domain names (options nis.domainname) 96

- displaying information (nis info) 87, 98

- displaying server name (ypwhich) 99

managing with FilerView 95

slave

guidelines for using 93

nis.slave.enable option (to enable NIS slave) 94

selection of a master 93

using for name resolution 92

specifying servers to bind to (options nis.servers) 96

NIS commands

nis info (displays NIS information) 87, 98

nis.slave.enable option (to enable NIS slave) 94

options nis.domainname (changes NIS domain name) 96

options nis.servers (binds NIS servers) 96

ypwhich (displays NIS server name) 99

nis.slave.enable option (to enable NIS slave) 94

NTP 227

O

options

cf.takeover.on_network_interface_failure option 17

nis.slave.enable (to enable NIS slave) 94

P

packets, jumbo frames 5

PAGP 155

Perfect Forward Secrecy (PFS) 188

performance, improving storage system 219

physical interfaces, adding (vif add) 167

ping command 23

ping6 command 23

Port Aggregation Protocol (PAGP) 155

ports, IP 221

prshared keys

configuring for IPsec 190

description of 188

PVCs

deleting incoming FORE/IP (atm atmarp) 62

deleting outgoing FORE/IP (atm atmarp) 62

description of 55

displaying address resolution (atm atmarp) 58

establishing incoming FORE/IP (atm atmarp) 57

establishing outgoing FORE/IP (atm atmarp) 57

R

rameters 126

RECEIVE statistics

on ATM card 217

on FAS250/FAS270 interfaces 209

on Fast Ethernet card 200

on FDDI card 213

on Gigabit Ethernet interface 204

on N3700 interfaces 209

round robin load balancing 158

route, static (adding) 75

routed, command 72

routing

default route 67

description of 64

fast path mechanism 65

in clusters 75

managing with FilerView 71

NFS-over-UDP, description of 65

routed daemon 64

table

description of 72

displaying (netstat) 72

managing 67

modification of 75

modifying (route) 75

TCP, description of 65

turning on or off (routed) 70, 71, 115

vFiler units 68

routing commands

netstat (displays routing table) 72

route (modifies routing table) 75

routed 72

routed (turns routing on or off) 70, 71, 115

routing table flags 73

routing table output 73

S

SA (Security Association) 186, 189

- search order, changing (nsswitch.conf file) 101
- second-level vifs
 - creating in a cluster 182
 - creating on a single storage system 178
 - in a cluster, description of 179
 - in single storage system, description of 176
- Secure Shell (SSH) 226
- security
 - trusted interface 16
 - untrusted interface 16
- Security Association (SA)
 - description of 186
 - displaying 198
 - lifetime 189
- security policies, IPsec
 - about 187
 - creating 193
 - deleting 194
 - displaying 194
 - managing 192
- services file 221
- set port channel command (for Cisco switch configuration) 156
- setting, IP addresses 12, 14
- single-mode vifs, creating (vif create single) 161
- slave, NIS 92
- SNMP commands
 - commands for traps 123
 - snmp configuration 116
- SNMP protocol
 - agent and groups supported 104
 - cluster configuration 114
 - configuration commands 116
 - custom MIB, description of 109
 - Data ONTAP implementation, description of 104
 - managing with FilerView 115
 - MIB specifications implemented 104
 - traps
 - commands 123
 - description of 120
 - parameters supported 126
 - types of 104
- SPANS, changing the AAL (atm atmconfig) 61
- static route, adding to routing table 75

- statistics
 - displaying interface (ifstat) 21, 22
 - ifstat command, description of 21
 - IPsec 195
 - on ATM card 217
 - on FAS250/FAS270 interfaces 209
 - on FDDI card 213
 - on Gigabit Ethernet interface 204
 - on N3700 interfaces 209
- stats commands
 - ifstat (displays interface statistics) 21, 22
 - IPsec stats (displays IPsec statistics) 195
 - vlan stat (displays VLAN statistics) 148
- subnet mask, configuring 15
- SVCs and PVCs, description of 54

T

- TCP transport
 - routing over 65
- TFTP 228
- time service 227
- Time-to-live (TTL), changing for dynamic DNS entries 91
- tp, mediatype 8
- tp-fd, mediatype 8
- TRANSMIT statistics
 - on ATM card 217
 - on FAS250/FAS270 interfaces 210
 - on Fast Ethernet card 201
 - on FDDI card 214
 - on Gigabit Ethernet card 206
 - on N3700 interfaces 210
- transport mode, IPsec 188
- traps, SNMP
 - commands 123
 - description of 120
 - parameters supported 126
 - types of 104
- Trivial File Transfer Protocol (TFTP) 228
- troubleshooting, network problems 23
- trunks. *See* vifs
- trusted, ifconfig option 16
- tunnel mode, not supported in IPsec 188

U

UDP transport

- configuring MTU size on UDP clients 7
- routing with NFS 65

UNI (User-Network Interface)

- description of 29
- verifying 35

untrusted, ifconfig option 16

user authentication, NIS and DNS configuration files 77

User Datagram Protocol (UDP). *See* UDP transport

V

vFiler units

- configuration with IPsec 189
- routing with 68

vif command 159

vif status command output, description of 169

vifs

- adding interface to (vif add) 167
- advantages of 152
- commands
 - active interface, selection of 163
 - persistence of 160
 - set port channel (for Cisco switch configuration) 156
 - vif (command syntax) 159
 - vif add (adds an interface to a virtual interface) 167
 - vif create (creates a virtual interface) 161, 182
 - vif create multi (creates multimode interface) 166
 - vif delete (deletes a virtual interface) 168
 - vif destroy (destroys a virtual interface) 174
 - vif favor (specifies preferred interface) 163
 - vif nofavor (specifies a non-preferred interface) 164
 - vif stat (displays statistics of a virtual interface) 173
 - vif status (displays status of a virtual interface) 169

creating, guidelines for 159

deleting an interface from 168

described 153

destroying 174

displaying statistics of virtual interface (vif stat) 173

displaying status of virtual interface (vif status) 169

Gigabit Ethernet interfaces in 159

IEEE 802.3ad 155

jumbo frames in 159

kinds of 154

Link Aggregation Control Protocol (LACP) 155

load-balancing methods in 158

management of (vif command) 159

maximum number of interfaces in 159

multimode vifs

configuring a (Cisco) switch for 156

creating (vif create multi) 166, 177

creating second-level vifs 175

default load balancing method 165

example of 155

IP-address based load balancing 158

load balancing methods 158

MAC-address based load balancing 158

operation of 155

prerequisites for creating 165

round robin load balancing 158

not favored interface, designating 164

Port Aggregation Protocol (PAgP) 155

preferred interface, specifying (vif favor) 163

second-level vifs

(on a single storage system), example of 176

creating in a cluster 182

creating on a single storage system 178

description of 175

example of 183

in a cluster, described 179

in single storage system, described 176

prerequisites for creating 177

single-mode vifs

active interface, selecting 163

creating 161

- operation of 154
 - preferred interface in 163
 - prerequisites for creating 161
- types of 154
- vif stat command output, description of 173
- VLAN interfaces in 159
- virtual aggregation. *See* vifs
- virtual interfaces. *See* vifs
- virtual local area network. *See* VLAN
- VLAN
 - adding an interface to 144
 - advantages of 136
 - ATM interface 139
 - configuring on a storage system 142
 - considerations for reverting Data ONTAP
 - version 139
 - creating on a storage system 141
 - definition 134
 - deleting on a storage system 145
 - display statistics of 148
 - FDDI interface 139
 - guidelines for setting up 138
 - how tagging works 136
 - ifconfig command 142
 - members, communication between 134
 - membership 134
 - persistence across reboots 138
 - port-based 134
 - setup requirements 137
 - statistics, viewing 148
 - tag 136
 - vlan command 140
- VLAN commands
 - persistence of 140
 - syntax of 140
 - vlan add 144
 - vlan create 141
 - vlan delete 145
 - vlan stat 148
- VLANs
 - interfaces in vifs 159

