

IBM System Storage N series Data ONTAP 7.3 Network Management Guide

GC52-1280-05 NA 210-04777_A0

Contents

Copyright information	11
Trademark information	13
About this guide	15
Audience	15
Supported features	16
Getting information, help, and services	16
Before you call	16
Using the documentation	17
Web sites	17
Accessing online technical support	17
Hardware service and support	17
Supported servers and operating systems	17
Firmware updates	
Accessing Data ONTAP man pages	
Terminology	19
Where to enter commands	
Keyboard and formatting conventions	20
Special messages	
How to send your comments	22
Network interfaces on your storage system	23
Network interface naming	23
Maximum number of network interfaces	25
The eOM interface	
How to use the RLM or BMC to manage Data ONTAP remotely	27
Ways to configure the RLM	27
Ways to configure the BMC	
Standards and characteristics of Ethernet frames	29
What jumbo frames are	29
Network interface requirements for jumbo frames	30
Guidelines to configure clients for jumbo frames	30
Flow control	30
Support for IPv6	31

Ways to configure IPv6 addresses	31
IPv6 address types	31
IPv6 address scopes	
IPv6 address states	32
How to transition from IPv4 to IPv6	33
Enabling or disabling IPv6	33
Types of address autoconfiguration	34
What stateless address autoconfiguration is	
Enabling or disabling router-advertised messages	35
What Neighbor Discovery does	35
ND message types	
How DAD works with Data ONTAP	
Network interface configuration	
Configuring network interfaces	39
Configuring an IP address for a network interface	40
Specifying a subnet mask for a network interface	41
Specifying the prefix length for a network interface	42
Specifying a broadcast address	42
Specifying a media type for a network interface	43
Specifying an MTU size for a network interface	43
Specifying the flow control type for a network interface	44
Specifying whether a network interface is trusted	44
Specifying the partner IP address in an active/active configuration	45
Specifying the partner interface in an active/active configuration	46
Enabling or disabling automatic takeover for a network interface	46
Removing a primary IP address from a network interface	47
Specifying the number of DAD attempts	
Viewing network interface settings	49
Creating or removing aliases	49
Changing the status of an interface	50
Viewing or modifying interface settings with FilerView	50
Blocking or unblocking protocols from network interfaces	51
Network interface information you can view	52
Viewing statistics of all active TCP connections	53
Viewing or clearing network interface statistics	54
Viewing network interface information with FilerView	57

How routing in Data ONTAP works	59
What fast path is	59
Similarities and differences between fast path over IPv4 and IPv6	60
How to manage the routing table	61
What the routed daemon does	61
When the routed daemon should be turned off	62
How dynamic routing works for IPv6	62
Routing tables in a vFiler unit environment	62
Circumstances that might alter the routing table	63
Specifying the default route	63
How to enable or disable routing mechanisms	64
Enabling or disabling fast path	64
Enabling or disabling the routed daemon from the command-line interface	64
Enabling or disabling the routed daemon with FilerView	65
How to view the routing table and default route information	65
Viewing the routing table from the command-line interface	66
Viewing the default route information from the command-line interface	67
Viewing the routing table and routing information by using FilerView	68
Modifying the routing table	68
How to maintain host-name information	71
How the /etc/hosts file works	71
Adding a host name in the /etc/hosts file	72
Hard limits for the /etc/hosts file	73
Editing the /etc/hosts file with FilerView	73
Changing the host name of a storage system	73
How to configure DNS to maintain host information	74
Configuring DNS from the command-line interface	75
How DNS resolves host names	76
DNS name caching	77
DNS information you can view	77
How to use dynamic DNS to update host information	78
How dynamic DNS updates work in Data ONTAP	79
Support for dynamic DNS updates in Data ONTAP	79
Enabling or disabling dynamic DNS updates	80
Disabling the transmission of DNS updates for an IP address	80

	Changing the time-to-live setting for DNS entries	. 81
	How to use NIS to maintain host information	. 81
	How using NIS slaves can improve performance	. 82
	How an NIS master is selected	. 83
	Creating /etc/hosts from the NIS master	83
	Guidelines for using NIS slaves	. 83
	NIS administrative commands	. 84
	How to configure NIS with Data ONTAP interfaces	. 85
	Enabling or disabling NIS using the command-line interface	. 85
	Specifying the NIS domain name	. 86
	Specifying NIS servers to bind to your storage system	. 86
	Enabling an NIS slave on your storage system	. 87
	What NIS information you can view	. 88
	Viewing NIS performance statistics	. 88
	Configuring DNS and NIS with FilerView	. 89
	How to change the host-name search order	. 90
	Changing the host-name search order with FilerView	. 91
	Changing the host-name search order	. 91
How	VLANs work	93
	VLAN membership	. 93
	VLAN membership How VLAN membership affects communication	. 93 . 94
	VLAN membership How VLAN membership affects communication GARP VLAN Registration Protocol	. 93 . 94 . 95
	VLAN membership How VLAN membership affects communication GARP VLAN Registration Protocol GVRP configuration for VLAN interfaces	. 93 . 94 . 95 . 95
	VLAN membership How VLAN membership affects communication GARP VLAN Registration Protocol GVRP configuration for VLAN interfaces VLAN tags	. 93 . 94 . 95 . 95 . 95
	VLAN membership How VLAN membership affects communication GARP VLAN Registration Protocol GVRP configuration for VLAN interfaces VLAN tags Advantages of VLANs	. 93 . 94 . 95 . 95 . 95 . 96
	VLAN membership How VLAN membership affects communication GARP VLAN Registration Protocol GVRP configuration for VLAN interfaces VLAN tags Advantages of VLANs Prerequisites for setting up VLANs	93 94 95 95 95 95 96 97
	VLAN membership How VLAN membership affects communication GARP VLAN Registration Protocol GVRP configuration for VLAN interfaces VLAN tags Advantages of VLANs Prerequisites for setting up VLANs Guidelines for setting up VLANs in Data ONTAP	93 94 95 95 95 95 96 97
	VLAN membership How VLAN membership affects communication GARP VLAN Registration Protocol GVRP configuration for VLAN interfaces VLAN tags Advantages of VLANs Prerequisites for setting up VLANs Guidelines for setting up VLANs in Data ONTAP The vlan command syntax	. 93 . 94 . 95 . 95 . 95 . 95 . 95 . 97 . 97
	VLAN membership How VLAN membership affects communication GARP VLAN Registration Protocol GVRP configuration for VLAN interfaces VLAN tags Advantages of VLANs Prerequisites for setting up VLANs Guidelines for setting up VLANs in Data ONTAP The vlan command syntax Creating a VLAN	. 93 . 94 . 95 . 95 . 95 . 95 . 95 . 95 . 97 . 97 . 98
	VLAN membership How VLAN membership affects communication GARP VLAN Registration Protocol GVRP configuration for VLAN interfaces VLAN tags Advantages of VLANs Prerequisites for setting up VLANs Guidelines for setting up VLANs in Data ONTAP The vlan command syntax Creating a VLAN	. 93 . 94 . 95 . 95 . 95 . 95 . 96 . 97 . 98 . 98 . 98
	VLAN membership How VLAN membership affects communication GARP VLAN Registration Protocol GVRP configuration for VLAN interfaces VLAN tags Advantages of VLANs Prerequisites for setting up VLANs Guidelines for setting up VLANs in Data ONTAP The vlan command syntax Creating a VLAN IPv6 link-local addresses for VLANs	. 93 . 94 . 95 . 95 . 95 . 95 . 95 . 95 . 95 . 97 . 97 . 97 . 98 . 98 100
	VLAN membership How VLAN membership affects communication GARP VLAN Registration Protocol GVRP configuration for VLAN interfaces VLAN tags Advantages of VLANs Prerequisites for setting up VLANs Guidelines for setting up VLANs in Data ONTAP The vlan command syntax Creating a VLAN IPv6 link-local addresses for VLANs Adding an interface to a VLAN	. 93 . 94 . 95 . 95 . 95 . 95 . 95 . 95 . 95 . 97 . 98 . 98 . 98 . 98 100 101
	VLAN membership How VLAN membership affects communication GARP VLAN Registration Protocol GVRP configuration for VLAN interfaces VLAN tags Advantages of VLANs Prerequisites for setting up VLANs Guidelines for setting up VLANs in Data ONTAP The vlan command syntax Creating a VLAN Configuring a VLAN IPv6 link-local addresses for VLANs Adding an interface to a VLAN Deleting VLANs	. 93 . 94 . 95 . 95 . 95 . 96 . 97 . 97 . 97 . 98 . 98 . 98 100 101 101
	VLAN membership How VLAN membership affects communication	. 93 . 94 . 95 . 95 . 95 . 96 . 97 . 97 . 98 . 98 100 101 101 101
	VLAN membership How VLAN membership affects communication GARP VLAN Registration Protocol GVRP configuration for VLAN interfaces VLAN tags Advantages of VLANs Prerequisites for setting up VLANs Guidelines for setting up VLANs in Data ONTAP The vlan command syntax Creating a VLAN Configuring a VLAN IPv6 link-local addresses for VLANs Adding an interface to a VLAN Deleting VLANs Enabling or disabling GVRP on your VLAN interface Viewing VLAN statistics	. 93 . 94 . 95 . 95 . 95 . 97 . 97 . 98 . 98 . 98 . 98 . 98 . 100 101 101 102 103
	VLAN membership How VLAN membership affects communication	. 93 . 94 . 95 . 95 . 95 . 97 . 97 . 97 . 98 . 98 100 101 101 101 102 103 104

How vifs work in Data ONTAP	107
Types of vifs	108
Single-mode vif	109
Static multimode vif	109
Dynamic multimode vif	110
Load balancing in multimode vifs	112
IP address and MAC address load balancing	112
Round-robin load balancing	112
Port-based load balancing	112
Guidelines for configuring vifs	113
The vif command	113
Creating a single-mode vif	114
Selecting an active interface in a single-mode vif	116
Designating a nonfavored interface in a single-mode vif	117
Failure scenarios for a single-mode vif	117
Creating a static multimode vif	118
Creating a dynamic multimode vif	119
Adding interfaces to a vif	121
Deleting interfaces from a vif	121
Viewing vif status	122
What the vif status information table contains	123
Viewing vif statistics	124
Destroying a vif	125
Second-level vifs	126
Guidelines for creating a second-level vif	126
Creating a second-level vif	126
Enabling failover in a second-level vif	127
Second-level vifs in an active/active configuration	128
Creating a second-level vif in an active/active configuration	129
How CDP works with Data ONTAP	
Data ONTAP support for CDP	133
Enabling or disabling CDP on your storage system	
Configuring hold time for CDP messages	134
Setting the intervals for sending CDP advertisements	135
Viewing or clearing CDP statistics	135
Viewing neighbor information by using CDP	137

How to monitor your storage system with SNMP	. 139
Types of SNMP traps in Data ONTAP	139
What a MIB is	140
What the SNMP agent does	140
How to configure the SNMP agent	140
Enabling or disabling SNMP using the command-line interface	142
Configuring SNMPv3 users	142
Setting SNMP access privileges	143
Viewing or modifying your SNMP configuration from the command-	
line interface	143
Modifying your SNMP configuration from FilerView	144
SNMP command syntax	144
SNMP security parameters	145
Example: SNMP commands	146
User-defined SNMP traps	148
How SNMP traps work	148
How to define or modify a trap	149
Viewing or modifying trap values from the command-line interface	149
Viewing or modifying trap values from FilerView	149
Defining traps in a configuration file	150
Example: Trap definitions	151
Command syntax for SNMP trap parameters	151
SNMP trap parameters	152
Internet Protocol Security	. 157
What security associations are	157
What security policies include	158
Key exchanges	158
IPsec implementation in Data ONTAP	159
IPsec in an active/active configuration	160
IPsec in a vFiler unit configuration	160
How to set up IPsec	161
Configuring certificate authentication	161
Requesting a signed certificate from a Windows 2000 certificate	
authority	162
Installing a certificate signed by a Windows 2000 certificate authority or	I
a Windows client	163

Requesting a signed certificate from a non-Windows 2000 certificate
authority164
Installing a certificate signed by a non-Windows 2000 certificate
authority on a Windows client
Installing a signed certificate on a storage system 166
Installing root certificates on a storage system 166
Specifying the subset of root certificates that Data ONTAP uses for
certificate authentication167
Viewing the subset of root certificates that Data ONTAP uses for
certificate authentication 167
Installing root certificates on a Windows client
Enabling the IPsec certificate authentication mechanism on a storage
system
Enabling the IPsec certificate authentication mechanism on a Windows
client 168
Kerberos support 169
Configuring preshared keys 169
Enabling or disabling IPsec 170
Security policies and IPsec
Creating a security policy 170
Security policy options 171
Displaying existing security policies
Deleting a security policy 172
Viewing IPsec statistics
Viewing security associations 175
How to diagnose network problems 177
Diagnosing transport layer problems 178
Viewing diagnostic results 179
How to diagnose ping problems
Increasing the ping throttling threshold value
Checking the ping throttling threshold status
Disabling ping throttling
Protecting your storage system from forged ICMP redirect attacks
Network interface statistics 183
Statistics for Gigabit Ethernet controller IV - VI and G20 interfaces 183

interfaces
Statistics for the N3700 network interfaces
Statistics for the BGE 10/100/1000 Ethernet interface
Ways to improve your storage system's performance 19'
IP port usage on a storage system 199
Host identification
/etc/services NNTP and TTCP ports 202
NFS-enabled ports 20
Ports not listed in /etc/services 20
FTP
SSH
Telnet
SMTP
Time service
DNS
DHCP
TFTP
HTTP
Kerberos
NFS
CIFS
SSL
SNMP
RSH
Syslog
The routed daemon
NDMP
SnapMirror and SnapVault
Error codes for the netdiag command 213
Index

Copyright and trademark information

Copyright	Copyright ©1994 - 2010 Network Appliance, Inc. All rights reserved. Printed in the U.S.A.
information	Portions copyright © 2010 IBM Corporation. All rights reserved.
	US Government Users Restricted Rights - Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.
	No part of this document covered by copyright may be reproduced in any form or by any means—graphic,electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrievalsystem—without prior written permission of the copyright owner.
	References in this documentation to IBM products, programs, or services do not imply that IBM intendsto make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM's product, program, or service maybe used. Any functionally equivalent product, program, or service that does not infringe any of IBM'sor NetApp's intellectual property rights may be used instead of the IBM or NetApp product, program, or service. Evaluation and verification of operation in conjunction with other products, except thoseexpressly designated by IBM and NetApp, are the user's responsibility.
	No part of this document covered by copyright may be reproduced in any form or by any means—graphic,electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrievalsystem—without prior written permission of the copyright owner.
	Software derived from copyrighted NetApp material is subject to the following license and disclaimer:
	THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS ORIMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIESOF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBYDISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUTNOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANYTHEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT(INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OFTHIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.
	NetApp reserves the right to change any products described herein at any time, and without notice.NetApp assumes no responsibility or liability arising from the use of products described herein, exceptas expressly agreed to in writing by NetApp. The use or purchase of this product does not convey alicense under any patent rights, trademark rights, or any other intellectual property rights of NetApp.
	The product described in this manual may be protected by one or more U.S.A. patents, foreign patents, or pending applications.
	RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph $(c)(1)(i)$ of the Rights in Technical Data and Computer

Softwareclause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business information

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

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

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

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

NetApp; the NetApp logo; the Network Appliance logo; Cryptainer; Cryptoshred; DataFabric; Data ONTAP; Decru; Decru DataFort; FAServer; FilerView; FlexCache; FlexClone; FlexShare; FlexVol; FPolicy; gFiler; Go further, faster; Manage ONTAP; MultiStore; NearStore; NetCache; NOW (NetApp on the Web); ONTAPI; RAID-DP; SANscreen; SecureShare; Simulate ONTAP; SnapCopy; SnapDrive; SnapLock; SnapManager; SnapMirror; SnapMover; SnapRestore; SnapValidator; SnapVault; Spinnaker Networks; Spinnaker Networks logo; SpinAccess; SpinCluster; SpinFlex; SpinFS: SpinHA; SpinMove; SpinServer; SpinStor; StoreVault; SyncMirror; Topio; vFiler; VFM; and WAFL are registered trademarks of NetApp, Inc. in the U.S.A. and/or other countries. Network Appliance, Snapshot, and The evolution of storage are trademarks of NetApp, Inc. in the U.S.A. and/or other countries and registered trademarks in some other countries. The StoreVault logo, ApplianceWatch, ApplianceWatch PRO, ASUP, AutoSupport, ComplianceClock, DataFort, Data Motion, FlexScale, FlexSuite, Lifetime Key Management, LockVault, NOW, MetroCluster, OpenKey, ReplicatorX, SecureAdmin, Shadow Tape, SnapDirector, SnapFilter, SnapMigrator, SnapSuite, Tech OnTap, Virtual File Manager, VPolicy, and Web Filer are trademarks of NetApp, Inc. in the U.S.A. and other countries. Get Successful and Select are service marks of NetApp, Inc. in the U.S.A.

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.

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

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

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

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

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

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

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

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

Any references in this information to non-IBM web sites are provided for convenience only and do not in any manner serve as an endorsement of those web sites. The materials at those web sites are not part of the materials for this IBM product and use of those web sites is at your own risk. IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

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

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

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

About this guide

You can use your product more effectively when you understand this document's intended audience and the conventions that this document uses to present information.

This guide describes how to configure and manage network interfaces, virtual interfaces (vifs), virtual LANs (VLANs), routing, IPsec, and SNMP on storage systems that run Data ONTAP. This guide also describes host-name resolution and SNMP.

This guide describes all storage system models; however, some models do not support all the networking interfaces. See the hardware guide for your storage system to identify which interfaces are supported on your system.

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

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

Next topics

Audience on page 15 Supported features on page 16 Getting information, help, and services on page 16 Accessing Data ONTAP man pages on page 18 Terminology on page 19 Where to enter commands on page 20 Keyboard and formatting conventions on page 20 Special messages on page 21 How to send your comments on page 22

Audience

This document is written with certain assumptions about your technical knowledge and experience.

This document is for systems administrators who are familiar with operating systems that run on storage system clients such as UNIX, MAC OSX, and Windows. It also assumes that you are familiar

with how Network File System (NFS), Common Internet File System (CIFS), and HyperText Transfer Protocol (HTTP) are used for file sharing or transfers.

Supported features

IBM[®] System Storage[™] N series storage systems are driven by NetApp[®] Data ONTAP[®] software. Some features described in the product software documentation are neither offered nor supported by IBM. Please contact your local IBM representative or reseller for further details. Information about supported features can also be found at the following Web site:

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

A listing of currently available N series products and features can be found at the following Web site:

www.ibm.com/storage/nas/

Getting information, help, and services

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

Next topics

Before you call on page 16 *Using the documentation* on page 17 *Web sites* on page 17 *Accessing online technical support* on page 17 *Hardware service and support* on page 17 *Supported servers and operating systems* on page 17 *Firmware updates* on page 18

Before you call

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

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

Using the documentation

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

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

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

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

Web sites

IBM maintains pages on the World Wide Web where you can get the latest technical information and download device drivers and updates.

- For NAS product information, go to the following Web site: www.ibm.com/storage/nas/
- For NAS support information, go to the following Web site: www.ibm.com/storage/support/nas/
- For AutoSupport information, go to the following Web site: www.ibm.com/storage/support/nas/
- For the latest version of publications, go to the following Web site: www.ibm.com/storage/support/nas/

Accessing online technical support

For online Technical Support for your IBM N series product, visit the following Web site:

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

Hardware service and support

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

www.ibm.com/planetwide/

Supported servers and operating systems

IBM N series products attach to many servers and many operating systems. To determine the latest supported attachments, follow the link to the Interoperability Matrices from the following Web site:

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

Firmware updates

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

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

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

Accessing Data ONTAP man pages

You can use the Data ONTAP manual (man) pages to access technical information.

About this task

Data ONTAP manual pages are available for the following types of information. They 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

Step

- 1. View man pages in the following ways:
 - Enter the following command at the storage system command line:

man command_or_file_name

- Click the manual pages button on the main Data ONTAP navigational page in the FilerView user interface.
- Use the Commands: Manual Page Reference, Volumes 1 and 2

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 distinguish storage system 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.

Terminology

To understand the concepts in this document, you might need to know how certain terms are used.

Storage terms

storage	Refers to the component of a storage system that runs the Data ONTAP operating
controller	system and controls its disk subsystem. Storage controllers are also sometimes called
	<i>controllers, storage appliances, appliances, storage engines, heads, CPU modules,</i> or <i>controller modules.</i>

storageRefers to the hardware device running Data ONTAP that receives data from andsystemsends data to native disk shelves, third-party storage, or both. Storage systems that
run Data ONTAP are sometimes referred to as *filers, appliances, storage appliances,*
gateways, or *systems.*

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

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

vif Refers to a single virtual interface that is created by grouping together multiple physical interfaces.

Cluster and high-availability terms

active/active configuration	In the Data ONTAP 7.2 and 7.3 release families, refers to a pair of storage systems (sometimes called <i>nodes</i>) configured to serve data for each other if one of the two systems stops functioning. Also sometimes referred to as <i>active/active pairs</i> . In the Data ONTAP 7.1 release family, this functionality is referred to as a <i>cluster</i> .
cluster	In the Data ONTAP 7.1 release family, refers to a pair of storage systems

(sometimes called *nodes*) configured to serve data for each other if one of the two systems stops functioning. In the Data ONTAP 7.3 and 7.2 release families, this functionality is referred to as an *active/active configuration*.

Where to enter commands

You can use your product more effectively when you understand how this document uses command conventions to present information.

You can perform common administrator tasks in one or more of the following ways:

- You can enter commands either at the system console or from any client computer that can obtain access to the storage system using a Telnet or Secure Shell (SSH) session.
 In examples that illustrate command execution, the command syntax and output shown might differ from what you enter or see displayed, depending on your version of the operating system.
- You can use the FilerView graphical user interface. For information about accessing your system with FilerView, see the *Data ONTAP System Administration Guide.*
- You can enter Windows, ESX, HP-UX, AIX, Linux, and Solaris commands at the applicable client console.

In examples that illustrate command execution, the command syntax and output shown might differ from what you enter or see displayed, depending on your version of the operating system.

- You can use the client graphical user interface. Your product documentation provides details about how to use the graphical user interface.
- You can enter commands either at the switch console or from any client that can obtain access to the switch using a Telnet session.

In examples that illustrate command execution, the command syntax and output shown might differ from what you enter or see displayed, depending on your version of the operating system.

Keyboard and formatting conventions

You can use your product more effectively when you understand how this document uses keyboard and formatting conventions to present information.

Convention	What it means
The IBM NAS support site	Refers to http://www.ibm.com/storage/support/nas/.

Keyboard conventions

Convention	What it means
Enter, enter	 Used to refer to the key that generates a carriage return; the key is named Return on some keyboards. Used to mean pressing one or more keys on the keyboard and then pressing the Enter key, or clicking in a field in a graphical interface and then typing information into the field.
hyphen (-)	Used to separate individual keys. For example, Ctrl-D means holding down the Ctrl key while pressing the D key.
type	Used to mean pressing one or more keys on the keyboard.

Formatting conventions

Convention	What it means
<i>Italic</i> font	 Words or characters that require special attention. Placeholders for information that you must supply. For example, if the guide says to enter the arp -d hostname command, you enter the characters "arp -d" followed by the actual name of the host. Book titles in cross-references.
Monospaced font	 Command names, option names, keywords, and daemon names. Information displayed on the system console or other computer monitors. Contents of files. File, path, and directory names.
Bold monospaced font	Words or characters you type. What you type is always shown in lowercase letters, unless your program is case-sensitive and uppercase letters are necessary for it to work properly.

Special messages

This document might contain the following types of messages to alert you to conditions that you need to be aware of.

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

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

How to send your comments

Your feedback is important in helping us provide the most accurate and high-quality information. If you have comments or suggestions for improving this document, send us your comments by e-mail to *starpubs@us.ibm.com*. Be sure to include the following:

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

Network interfaces on your storage system

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

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

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

In addition, some storage system models include a physical network interface named e0M. The e0M interface is used only for Data ONTAP management activities, such as for running a Telnet, SSH, or RSH session.

Next topics

Network interface naming on page 23 *Maximum number of network interfaces* on page 25 *The eOM interface* on page 26 *How to use the RLM or BMC to manage Data ONTAP remotely* on page 27

Related concepts

Network interface configuration on page 39 *How vifs work in Data ONTAP* on page 107 *How VLANs work* on page 93

Network interface naming

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

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

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

Vif names are user specified. A vif's name should meet the following criteria:

- It must begin with a letter.
- It must not contain any spaces.
- It must not contain more than 15 characters.
- It must not already be in use for a vif.

VLAN interface names are in the following format:

- <physical_interface_name>-<vlan_ID>
- <vif_name>-<vlan_ID>

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

Interface type	Interface name format	Examples of names
Physical interface on a single-port adapter or slot	e <slot_number></slot_number>	e0 e1
Physical interface on a multiple-port adapter or slot	e <slot_number><port_letter></port_letter></slot_number>	e0a e0b e0c e0d e1a e1b
Vif	Any user-specified string that meets certain criteria	web_vif vif1
VLAN	<physical_interface_name>-<vlan-id> or <vif_name>-<vlan_id></vlan_id></vif_name></vlan-id></physical_interface_name>	e8-2 vif1-3

Host names

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

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

Interface type	Host name
Single-port Ethernet interface in slot 0	toaster-e0

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

Maximum number of network interfaces

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

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

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

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

Storage system memory	Maximum number of network interfaces
2 GB or less	128
2 GB or less in an active/active configuration	256
6 GB or less	256
6 GB or less in an active/active configuration	512
More than 6 GB	512
More than 6 GB in an active/active configuration	1,024

Related references

Network interface statistics on page 183

The e0M interface

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

On a storage system that includes the e0M interface, the Ethernet port that is indicated by a wrench icon on the rear of the chassis connects to an internal Ethernet switch. The internal Ethernet switch then provides connectivity to the e0M interface and the Remote LAN Module (RLM). The following diagram illustrates the connections.



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

The e0M interface does not support vifs, VLANs, and jumbo frames.

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

- Telnet
- RSH
- HTTP or HTTPS
- SSH

• SNMP

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

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

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

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

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

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

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

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

Next topics

Ways to configure the RLM on page 27 *Ways to configure the BMC* on page 28

Ways to configure the RLM

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

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

• Initializing a storage system that has the RLM pre-installed

When the storage system setup process is complete, the rlm setup command runs automatically. For more information about the entire setup process, see the *Data ONTAP Software Setup Guide*.

- Running the Data ONTAP setup script The setup script ends by initiating the rlm setup command.
- Running the Data ONTAP rlm setup command For information about using the rlm setup command to configure the RLM, see the *Data ONTAP System Administration Guide*.

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

Ways to configure the BMC

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

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

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

For information about using the bmc setup command to configure the BMC, see the *Data ONTAP System Administration Guide*.

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

Standards and characteristics of Ethernet frames

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

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

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

Next topics

What jumbo frames are on page 29 *Flow control* on page 30

What jumbo frames are

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

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

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

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

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

Next topics

Network interface requirements for jumbo frames on page 30 *Guidelines to configure clients for jumbo frames* on page 30

Network interface requirements for jumbo frames

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

Guidelines to configure clients for jumbo frames

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

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

- Configure jumbo frames on the client and on your storage system. Find how to configure jumbo frames on your client by checking the network adapter documentation for your client.
- Enlarge the client's TCP window size.

The minimum value for the client's window size should be two times the MTU size, minus 40, and the maximum value can be the highest value your system allows. Typically, the maximum value you can set for your client's TCP window is 65,535. If your storage system is configured to support jumbo frames and the client is not, the communication between the storage system and the client occurs at the client's frame size.

- Configure the storage system and UDP clients to have the same MTU size. User Datagram Protocol (UDP) systems do not negotiate the MTU size. If your storage system and clients do not have the same MTU size, the storage system might send packets that the clients cannot receive.
- Check the MTU size of any intermediate subnets if your storage system and the client are on different subnets.

If the storage system and the client (both configured to use jumbo frames) are on different subnets and an intermediate subnet does not support jumbo frames, the intermediate router fragments the IP packets and the advantages of using jumbo frames are lost.

Related tasks

Specifying an MTU size for a network interface on page 43

Flow control

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

To achieve flow control, you can specify a flow control option that causes packets called Pause frames to be used as needed. For example, link-partner A sends a Pause On frame to link-partner B when its receive buffers are nearly full. Link-partner B suspends transmission until it receives a Pause Off frame from link-partner A or a specified timeout threshold is reached.

Support for IPv6

Starting with Data ONTAP 7.3.1, Internet Protocol version 6 (IPv6) is supported on your storage system's network. IPv6 increases the IP address size from 32 bits (in IPv4) to 128 bits. This larger address space provides expanded routing and addressing capabilities.

Data ONTAP 7.3 and earlier used IPv4 for all the addressing and networking requirements. However, IPv4 has many limitations, such as limited address space and security. To address these limitations, the Internet Engineering Task Force (IETF) developed a new version of IP, called IPv6.

You can enable the IPv6 option and configure IPv6 addresses on the network interfaces of the storage system. IPv6 addresses can also be automatically configured.

Next topics

Ways to configure IPv6 addresses on page 31 *How to transition from IPv4 to IPv6* on page 33 *Enabling or disabling IPv6* on page 33 *Types of address autoconfiguration* on page 34 *What Neighbor Discovery does* on page 35 *How DAD works with Data ONTAP* on page 36

Ways to configure IPv6 addresses

IPv6 addresses can be configured on the network interfaces of your storage system, either manually or automatically. The configuration of an IPv6 address depends on the type and scope of the address.

Next topics

IPv6 address types on page 31 *IPv6 address scopes* on page 32 *IPv6 address states* on page 32

IPv6 address types

There are three types of IPv6 addresses: unicast, anycast, and multicast.

Unicast address	This address identifies a single interface. A data packet sent to a unicast address is delivered only to the interface that is identified by that address.
Anycast address	This address identifies a set of interfaces. A data packet sent to an anycast address is delivered to the nearest interface (according to the routing protocols' measure of distance) that is identified by that address.

Note: Anycast address is not supported in Data ONTAP.

MulticastThis address identifies a set of interfaces. A data packet sent to a multicast addressaddressis delivered to all the interfaces that are identified by that address.

Note: In IPv6, multicast addresses replace broadcast addresses.

IPv6 address scopes

IPv6 addresses fall under three scopes: global, link-local, and unique local.

Global address	This address has an unlimited scope.
Link-local	This address has a link-only scope that can be used to reach neighboring nodes that are attached to the same link. This address is automatically assigned to a network interface.
Unique local address	The address scope is limited to a local site or local set of sites. These addresses cannot be routed on the global Internet.

IPv6 address states

Before and after an IPv6 address is assigned, it goes through various states, such as tentative address, duplicate address, preferred address, and so on. These address states are applicable to both manually and automatically configured addresses.

An IPv6 address can have one or more of the following states:

Tentative address	An address whose uniqueness on a link is being verified. When an address is configured on a network interface (either manually or automatically), the address is initially in the tentative state. Such an address is not considered to be assigned to an interface. An interface discards received packets addressed to a tentative address, but accepts Neighbor Discovery packets related to Duplicate Address Detection (DAD) for the tentative address.
Duplicate address	If DAD finds that an address is not unique, it is moved to the duplicate state. Such an address cannot be used for sending and receiving data.
Preferred address	An address used to send and receive data packets from and to a network interface without any restriction on the upper layer protocols.
Deprecated address	A preferred address becomes deprecated when its preferred lifetime expires. The use of this address is discouraged, but not prohibited.
Valid address	A uniquely verified address that you can assign to a network interface for sending and receiving data. A valid address can be a preferred or deprecated address.
Invalid address	A network interface address that can no longer send or receive data packets. A valid address becomes invalid when its valid lifetime expires. An invalid address is removed from the network interface.

How to transition from IPv4 to IPv6

A transition mechanism enables IPv6 hosts and routers to be compatible with IPv4 hosts and routers. Starting with Data ONTAP 7.3.1, a dual stack mechanism is used for transitioning from IPv4 to IPv6.

In the dual stack mechanism, the following modes are supported:

- Only IPv4 mode: In this mode, IPv6 is not enabled.
- Only IPv6 mode: In this mode, IPv6 is enabled and IPv4 addresses are not configured on any interface.
- IPv6/IPv4 mode: In this mode, IPv6 is enabled and both IPv4 and IPv6 addresses are configured on the network interfaces.

Attention: In the "Only IPv6 mode," address lookup can return both IPv4 and IPv6 addresses. If you use an IPv4 address to set up communication in the "Only IPv6 mode," the communication fails. Therefore, you should have at least one IPv4 address configured in a network interface and then use the "IPv6/IPv4 mode."

Data ONTAP does not support the following IPv6 transition mechanisms (defined in RFC 2893):

- Configured tunneling of IPv6 over IPv4
- IPv4-mapped IPv6 addresses
- Automatic tunneling of IPv6 over IPv4

Enabling or disabling IPv6

You can enable IPv6 on all the interfaces of your storage system either during setup or when the storage system is in operation. You can disable IPv6 on your storage system if you want to revert to IPv4 addressing.

About this task

- You can enable IPv6 during initial system configuration when the setup command is run for the first time. If you want to enable IPv6 later, you can rerun the setup command or configure IPv6 manually. For more information about the setup command, see the *Data ONTAP Software Setup Guide*.
- You can enable IPv6 only for the entire storage system, but not for a network interface or a vFiler unit.

Step

1. To enable or disable IPv6 when the storage system is in operation (not during setup), enter the following command:

options ip.v6.enable {on|off} on—Enables IPv6 off—Disables IPv6

After you finish

If you have enabled IPv6 when the storage system is in operation, you must manually restart all server applications, except CIFS, FTP, and HTTP, to run over IPv6. For CIFS, FTP, and HTTP to work over IPv6, you must enable their individual IPv6 options. For more information about the protocols supported over IPv6, see the *Data ONTAP File Access and Protocols Management Guide*.

Note: If the applications are running only over IPv4, you do not need to restart the applications.

Types of address autoconfiguration

IPv6 defines both a stateful and a stateless address autoconfiguration mechanism. Data ONTAP 7.3.1 and later supports IPv6 stateless address autoconfiguration.

The Neighbor Discovery protocol is one of the protocols that facilitates address autoconfiguration.

Next topics

What stateless address autoconfiguration is on page 34 *Enabling or disabling router-advertised messages* on page 35

Related concepts

What Neighbor Discovery does on page 35

What stateless address autoconfiguration is

The stateless address autoconfiguration mechanism allows a host to generate its own addresses by using a combination of locally available information and router-advertised information. The stateless address autoconfiguration requires minimal manual configuration of hosts and routers.

Data ONTAP supports the following two types of autoconfigured IPv6 addresses:

- Autoconfigured address based on the router-advertised prefix: This address is a combination of the network prefix, which is router-advertised, and the network interface identifier.
- Autoconfigured link-local address: In the absence of routers, a host can generate only link-local addresses. Link-local addresses allow communication between hosts and routers that are on the same link.

RFC 2462 describes address autoconfiguration.

Enabling or disabling router-advertised messages

RA messages help in autoconfiguring addresses that have global scope and in learning routes and prefixes. If, due to security reasons, you do not want the MAC address of the network interfaces to be viewed by any external network, you can disable RA address autoconfiguration.

About this task

You can use the ip.v6.ra_enable option to enable or disable router-advertised (RA) messages.

- By default, the ip.v6.ra_enable option is set to on.
- You can enable the RA option only for the entire storage system; you cannot enable it for a network interface.
- Disabling the RA option does not remove the existing autoconfigured addresses and the routes learned.
- When the RA option is disabled, the RA message is dropped. Therefore, no default route is learned, the default router failover is disabled, and link MTU updates are stopped.

Step

1. To enable or disable RA address autoconfiguration, enter the following command:

```
options ip.v6.ra_enable {on|off}
```

on-Enables RA address autoconfiguration

off-Disables RA address autoconfiguration

What Neighbor Discovery does

The Neighbor Discovery (ND) protocol enables hosts and routers to discover the presence of neighboring IPv6 hosts and routers. The ND protocol also helps in identifying the link-layer address of hosts and routers and in performing Duplicate Address Detection (DAD).

The ND protocol replaces the IPv4 protocols, such as Address Resolution Protocol (ARP), Internet Control Message Protocol (ICMP) Router Discovery, and ICMP Redirect.

The various ND mechanisms for enabling interaction between nodes that are on the same link, as described in RFC 2461, are described below:.

Router discovery	How hosts find routers that reside on an attached link.
Prefix discovery	How hosts discover the set of address prefixes that define which destinations are on-link for an attached link.
Parameter discovery	How hosts discover operating parameters such as link MTU and default hop limit for outgoing packets.

Address autoconfiguration	How hosts and routers automatically configure an address for an interface.
Address resolution	How hosts and routers determine the link-local address of a neighbor by using the IPv6 address of the neighbor.
Next-hop determination	How hosts and routers determine the IPv6 address of a neighbor to which a packet should be sent, by using the destination address. The next hop can be either the destination address or a router address.
Neighbor Unreachability Detection	How hosts and routers determine that a neighbor is no longer reachable.
Duplicate Address Detection	How hosts and routers determine that an address considered for use is not already in use by a neighbor.
Redirect	How a router informs a host of a better first-hop router to reach a particular destination.

ND message types

There are five types of ND messages: Router Solicitation, Router Advertisement, Neighbor Solicitation, Neighbor Advertisement, and Redirect. You can specify various ND options in an ND message.

Router Advertisement and Router Solicitation messages facilitate host-router discovery functions. Neighbor Solicitation and Neighbor Advertisement messages facilitate exchange of information between neighboring hosts on the same network. The Redirect message is used to inform a host of a better route for sending data packets to a particular destination. All the ND message types use the Internet Control Message Protocol version 6 (ICMPv6) message structure.

The ND options that can be specified in an ND message are the source link-layer address, target linklayer address, prefix information, MTU, and redirected header. These ND options provide additional information such as MAC addresses, on-link network prefixes, on-link MTU information, and redirection data.

Note: Data ONTAP supports a maximum of 10 options in an ND message.

How DAD works with Data ONTAP

Before assigning unicast addresses to an interface, Duplicate Address Detection (DAD) is performed on the addresses to ensure that the addresses configured on a link are unique. DAD is performed on all unicast addresses (both manually and automatically configured). When the DAD procedure fails for an address, the address is not configured.

DAD prevents multiple nodes from using the same address simultaneously. DAD is performed on all unicast addresses of a network interface, provided the value of the dad_attempts option for that interface is greater than zero.
To check the uniqueness of an address, a node sends Neighbor Solicitation messages, each separated by an interval of 1 second. The number of Neighbor Solicitation messages sent is equal to the value of the dad_attempts option for the network interface.

An address on which the DAD procedure is applied remains in the tentative state until the procedure has been successfully completed. The target address of the Neighbor Solicitation message is set to the address that is being checked and remains in the tentative state. If the node receives a valid Neighbor Advertisement message with the tentative address as target, the tentative address is not unique. The tentative address is marked duplicated and cannot be used for any data communication.

If DAD fails for a link-local address, the network interface is configured to the down status.

If a node does not receive a Neighbor Advertisement message after sending the Neighbor Solicitation messages for a tentative address, the address is considered unique. When an address is determined to be unique, it is assigned to the network interface.

Example: Duplicated unicast address

The following example shows DAD failure for a unicast address, where the address state changes from tentative to duplicated.

```
system1> ifconfig e0b 2001:0db8::99
system1> ifconfig e0b
e0b: flags=0x2508866<BROADCAST,RUNNING,MULTICAST,TCPCKSUM> mtu 1500
inet6 2001:0db8::99 prefixlen 64 tentative
ether 00:a0:98:08:64:07 (auto-1000t-fd-cfg_down) flowcontrol full
system1> Wed Aug 6 09:24:44 GMT [system1:netif.linkUp:info]: Ethernet
e0b: Link up.
Wed Aug 6 09:24:44 GMT [system1:netinet6.nbr.dad.dtcDupAdr:error]:
eOb: DAD detected duplicate IPv6
address 2001:0db8::99: %d NS, 0 NA.
Wed Aug 6 09:24:44 GMT [system1:netinet6.nbr.dad.complete:error]:
eOb: DAD complete for 2001:0db8::99- duplicate found.
Wed Aug 6 09:24:44 GMT [system1:netinet6.nbr.manl.intvtnReg:error]:
e0b: Manual intervention required.
Wed Aug 6 09:24:45 GMT [systeml:netinet6.nbr.dadStrc.notFndl:error]:
nd6_dad_timer: DAD structure is not found.
system1> ifconfig e0b
e0b: flags=0x2d48867<UP, BROADCAST, RUNNING, MULTICAST, TCPCKSUM, LINK_UP>
mtu 1500
inet6 2001:0db8::99 prefixlen 64 duplicated
inet6 fe80::2a0:98ff:fe08:6407 prefixlen 64 scopeid 0x2 autoconf
ether 00:a0:98:08:64:07 (auto-1000t-fd-up) flowcontrol full
```

Example: Duplicated link-local address

The following example shows DAD failure for a link-local address, where the network interface is configured to the down status.

system1> ifconfig e0b up system1> Tue Jul 22 16:46:38 GMT [system1: netif.linkUp:info]:

Ethernet e0b: Link up. Tue Jul 22 16:46:39 GMT [system1: netinet6.nbr.dad.dtcDupAdr:error]: eOb: DAD detected duplicate IPv6 address fe80:0002::02a0:98ff:fe08:6407: %d NS, 0 NA. Tue Jul 22 16:46:39 GMT [system1: netinet6.nbr.dad.complete:error]: eOb: DAD complete for fe80:0002::02a0:98ff:fe08:6407 - duplicate found. Tue Jul 22 16:46:39 GMT [system1: netinet6.nbr.manl.intvtnReq:error]: e0b: Manual intervention required. Tue Jul 22 16:46:39 GMT [system1: netif.linkInfo:info]: Ethernet e0b: Link configured down. Tue Jul 22 16:46:40 GMT [system1: netinet6.nbr.dadStrc.notFnd1:error]: nd6_dad_timer: DAD structure is not found. system1> ifconfig -a e0a: flags=0x2508866<BROADCAST,RUNNING,MULTICAST,TCPCKSUM> mtu 1500 ether 00:a0:98:08:64:06 (auto-1000t-fd-cfg_down) flowcontrol full e0b: flags=0x2508866<BROADCAST,RUNNING,MULTICAST,TCPCKSUM> mtu 1500 ether 00:a0:98:08:64:07 (auto-1000t-fd-cfg_down) flowcontrol full e0c: flags=0x2508866<BROADCAST,RUNNING,MULTICAST,TCPCKSUM> mtu 1500 ether 00:a0:98:08:64:08 (auto-1000t-fd-cfg_down) flowcontrol full e0d: flags=0x2508866<BROADCAST,RUNNING,MULTICAST,TCPCKSUM> mtu 1500 ether 00:a0:98:08:64:09 (auto-1000t-fd-cfg_down) flowcontrol full e0e: flags=0x2508866<BROADCAST,RUNNING,MULTICAST,TCPCKSUM> mtu 1500 ether 00:a0:98:08:64:0a (auto-1000t-fd-cfg_down) flowcontrol full e0f: flags=0x2508866<BROADCAST,RUNNING,MULTICAST,TCPCKSUM> mtu 1500 ether 00:a0:98:08:64:0b (auto-1000t-fd-cfg_down) flowcontrol full 10: flags=0x1948049<UP,LOOPBACK,RUNNING,MULTICAST,TCPCKSUM,LINK_UP,UDPCKSU M> mtu 8160 inet 127.0.0.1 netmask 0xff000000 broadcast 127.0.0.1 inet6 fe80::1 prefixlen 64 scopeid 0x7 autoconf inet6 ::1 prefixlen 128 ether 00:00:00:00:00 (VIA Provider)

Related tasks

Specifying the number of DAD attempts on page 48

Network interface configuration

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

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

- Assign an IP address to a network interface.
- Set parameters such as network mask, broadcast address, and prefix length.

Note: If IPv6 is enabled on your storage system, you can set only the prefix length. IPv6 does not have a network mask and does not support broadcast addresses.

- Set hardware-dependent values such as media type, MTU size, and flow control.
- Specify whether the interface should be attached to a network with firewall security protection.
- Specify whether the network interface must be registered with Windows Internet Name Services (WINS), if CIFS is running and at least one WINS server has been configured.
- Specify the IP address of an interface or specify the interface name on an active/active configuration partner for takeover mode.

Note: When using IPv6 in an active/active configuration, you can specify only the partner interface name (and not the IP address) on the active/active configuration for takeover mode.

• View the current configuration of a specific interface or all interfaces that exist on your storage system.

Next topics

Configuring network interfaces on page 39 Creating or removing aliases on page 49 Changing the status of an interface on page 50 Viewing or modifying interface settings with FilerView on page 50 Blocking or unblocking protocols from network interfaces on page 51 Network interface information you can view on page 52

Related concepts

Network interfaces on your storage system on page 23

Configuring network interfaces

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

During system setup, you can configure the IP addresses for the network interfaces. An ifconfig command is included in the /etc/rc file of the root volume for each network interface that you

configured during the system setup. After your storage system has been set up, the ifconfig commands in the /etc/rc file are used to configure the network interfaces on subsequent storage system reboots.

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

Next topics

Configuring an IP address for a network interface on page 40 Specifying a subnet mask for a network interface on page 41 Specifying the prefix length for a network interface on page 42 Specifying a broadcast address on page 42 Specifying a media type for a network interface on page 43 Specifying an MTU size for a network interface on page 43 Specifying the flow control type for a network interface on page 44 Specifying whether a network interface is trusted on page 44 Specifying the partner IP address in an active/active configuration on page 45 Specifying the partner interface in an active/active configuration on page 46 Enabling or disabling automatic takeover for a network interface on page 47 Specifying the number of DAD attempts on page 48 Viewing network interface settings on page 49

Configuring an IP address for a network interface

You can configure IP addresses for your network interface during system setup. To configure the IP addresses later, you should use the *ifconfig* command. You can configure both IPv4 and IPv6 addresses for a network interface.

About this task

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

Step

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

ifconfig interface_name IP_address

interface_name is the name of the network interface.

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

Example

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

ifconfig e3a 192.0.2.10

To configure a quad-port Ethernet interface e3a to use the IPv6 address 2001:0db8:35ab:0:8a2e: 0:0370:85, enter the following command:

ifconfig e3a 2001:0db8:35ab:0:8a2e:0:0370:85

Related tasks

Specifying a subnet mask for a network interface on page 41

Specifying a subnet mask for a network interface

You must specify a subnet mask if you have created subnets that do not match the class boundary of the IPv4 address of the network interface. You can specify a subnet mask for a network interface by using the ifconfig command. IPv6 does not support subnet masks.

About this task

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

Step

1. To specify a subnet mask, enter the following command:

ifconfig interface_name netmask mask

interface_name is the name of the network interface.

mask is the subnet mask.

Example

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

ifconfig e3a netmask 255.255.255.0

Related tasks

Configuring an IP address for a network interface on page 40

Specifying the prefix length for a network interface

Prefix length specifies the number of bits in the IP address that are to be used as the subnet mask. You can specify the prefix length for a network interface by using the *ifconfig* command.

About this task

For an IPv4 address, the prefix length must be less than or equal to 32 bits. For an IPv6 address, the prefix length must be less than or equal to 128 bits. The default value of the prefix length for an IPv6 address is 64 bits.

Step

1. To specify the prefix length, enter the following command:

ifconfig interface_name ip_address prefixlen length

ip_address is the IP address assigned to the network interface.

length is the prefix length for the network interface.

Example

To configure a prefix length of 24 bits, enter the following command:

ifconfig e0a 192.0.2.16 prefixlen 24

To configure a prefix length of 64 bits for an IPv6 address, enter the following command:

ifconfig e3a 2001:0db8:35ab:0:8a2e:0:0370:85 prefixlen 64

Specifying a broadcast address

You can use a broadcast address to send a message to all the machines on a subnet. You can specify a broadcast address by using the ifconfig command.

About this task

IPv6 does not support broadcast addresses.

Step

1. To specify a broadcast address, enter the following command:

ifconfig interface_name broadcast address

interface_name is the name of the network interface.

address is the broadcast address.

Example

To set a broadcast address of 192.0.2.25 for the network 192.0.2.10 with subnet mask 255.255.255.0, enter the following command:

ifconfig e3a broadcast 192.0.2.25

Specifying a media type for a network interface

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

Step

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

ifconfig interface_name mediatype type

interface_name is the name of the network interface.

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

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

Example

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

ifconfig e2a mediatype 100tx-fd

Specifying an MTU size for a network interface

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

Step

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

ifconfig interface_name mtusize size

interface_name is the name of the network interface.

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

Example

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

ifconfig e8 mtusize 9000

Related concepts

Standards and characteristics of Ethernet frames on page 29 What jumbo frames are on page 29 Guidelines to configure clients for jumbo frames on page 30

Specifying the flow control type for a network interface

You can specify the flow control type for a network interface to manage the flow of frames between two directly connected link-partners by using the *ifconfig* command. You can configure flow control on interfaces operating at or above 1,000 Mbps.

About this task

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

You can also use the ifstat command to view the operational flow control setting.

Step

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

ifconfig interface_name flowcontrol value

interface_name is the name of the network interface.

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

none	No flow control
receive	Able to receive flow control frames
send	Able to send flow control frames
full	Able to send and receive flow control frames

The default flow control type is full.

Example

To turn off flow control on interface e8, enter the following command:

ifconfig e8 flowcontrol none

Related concepts

Flow control on page 30

Specifying whether a network interface is trusted

You can specify whether a network interface is trustworthy or untrustworthy. When you specify an interface as untrusted (untrustworthy), any packets received on the interface are likely to be dropped.

For example, if you run a ping command on an untrusted interface, the interface drops any ICMP response packet received.

Step

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

ifconfig interface_name {trusted|untrusted}

interface_name is the name of the network interface.

trusted specifies that the network interface is trustworthy.

untrusted specifies that the network interface is untrustworthy.

Example

To specify that the network attached to interface e8 is not trustworthy for firewall security, enter the following command:

ifconfig e8 untrusted

Specifying the partner IP address in an active/active configuration

In an active/active configuration, you can assign a partner IP address to a network interface. The network interface takes over this IP address when a failover occurs. You can use the *ifconfig* command to specify the partner IP address.

About this task

You can specify only an IPv4 address for takeover in an active/active configuration.

Step

1. To assign the partner IP address, enter the following command:

ifconfig interface_name partner address

interface_name is the name of the network interface.

address is the partner IP address.

Example

To specify the IP address on the partner interface that takes over the interface e8 in case of a failover, enter the following command:

ifconfig e8 partner 192.0.2.10

Related tasks

Specifying the partner interface in an active/active configuration on page 46

Specifying the partner interface in an active/active configuration

In an active/active configuration, you can assign the name of a partner interface. The partner interface takes over the network interface when a failover occurs. You can specify the partner interface by using the ifconfig command.

About this task

When using IPv6, you must specify the partner interface, and not an IP address.

Step

1. To specify a partner interface name, enter the following command:

ifconfig interface_name partner partner_interface

interface_name is the name of the network interface.

partner_interface is the name of the partner network interface.

Example

To specify e3 as the interface for the active/active configuration partner that takes over the interface e8 when e8 fails, enter the following command:

ifconfig e8 partner e3

Related tasks

Specifying the partner IP address in an active/active configuration on page 45

Enabling or disabling automatic takeover for a network interface

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

Before you begin

You must enable takeover on interface failures by entering the following command:

options cf.takeover.on_network_interface_failure enable

About this task

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

Step

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

```
ifconfig interface_name {nfo|-nfo}
```

interface_name is the name of the network interface.

nfo-Enables negotiated failover

-nfo-Disables negotiated failover

Example

To enable negotiated failover on the interface e8 of an active/active configuration, enter the following command:

ifconfig e8 nfo

Removing a primary IP address from a network interface

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

Before you begin

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

Step

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

ifconfig interface_name 0

interface_name is the name of the network interface.

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

ifconfig interface_name 0.0.0.0

Example

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

ifconfig e3 0

Note: To remove a primary IPv6 address, you can use either of these commands:

- ifconfig interface_name 0::0
- ifconfig interface_name inet6 0

Related tasks

Creating or removing aliases on page 49

Specifying the number of DAD attempts

To check the uniqueness of an IPv6 address, a node sends Neighbor Solicitation messages. You can use the ifconfig command to specify the number of consecutive Neighbor Solicitation messages (dad_attempts) to be sent while performing DAD on a tentative address before it can be configured.

Before you begin

IPv6 must be enabled on the storage system.

About this task

A value of zero for the dad_attempts option indicates that DAD is not performed on the tentative addresses. A value of one for the dad_attempts option indicates a single transmission with no follow-up retransmission and so on.

Step

1. Enter the following command:

ifconfig interface_name dad_attempts value

interface_name is the name of the interface

value is the total number of consecutive Neighbor Solicitation messages sent while performing DAD on a tentative address. The default value is 2.

You can set the dad_attempts value from 0 to 15 for physical interfaces and from 0 to 7 for vifs and VLANs.

Note: A dad_attempts value that is greater than 13 does not work in certain scenarios. Therefore, it is best to set the dad_attempts value to less than 13.

Example

You can configure the interface e0a for sending four consecutive Neighbor Solicitation messages by using the following command:

ifconfig e0a dad_attempts 4

The following is the output of the ifconfig command:

```
ifconfig e0a
e0a: flags=0x2d48867<UP,BROADCAST,RUNNING,MULTICAST,TCPCKSUM> mtu 1500
dad_attempts 4
inet6 fe80::2a0:98ff:fe06:c8f6 prefixlen 64 scopeid 0x3 autoconf
tentative
ether 00:a0:98:06:c8:f6 (auto-1000t-fd-up) flowcontrol full
```

Related concepts

How DAD works with Data ONTAP on page 36

Viewing network interface settings

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

Step

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

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

Related tasks

Viewing or modifying interface settings with FilerView on page 50

Creating or removing aliases

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

About this task

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

Note: For IPv6 addresses, the link-local and autoconfigured addresses are automatically added as alias addresses even without a primary address configured for an interface.

Step

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

ifconfig interface_name [-]alias address

Example

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

```
ifconfig e0 alias 192.0.2.30
```

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

ifconfig e0 -alias 192.0.2.30

Changing the status of an interface

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

About this task

If you have enabled IPv6 on your storage system, you can bring up the interface without a valid address configured because a link-local address is created automatically for the interface. However, if the /etc/rc file contains an entry to keep a network interface in down status, a link-local address is not created.

Step

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

ifconfig interface {up|down}

up-makes the interface active

down-makes the interface inactive

Viewing or modifying interface settings with FilerView

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

Steps

- 1. From the list on the left pane, click Network > Manage Interfaces.
- **2.** Depending on whether you want to view or modify the configuration settings, perform the following step:

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

Blocking or unblocking protocols from network interfaces

You can use the interface.blocked.protocol option to block specified network protocols, including CIFS, iSCSI, FTP, or NFS, on selected network interfaces. You can also unblock a protocol from a network interface.

Step

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

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

Example

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

options interface.blocked.cifs e9

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

options interface.blocked.cifs e0a,e0b

To block NFS, CIFS, and FTP from the interface e0a, enter the following commands:

options interface.blocked.nfs e0a

options interface.blocked.cifs e0a

options interface.blocked.ftpd e0a

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

```
options interface.blocked.cifs ""
```

Network interface information you can view

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

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

Command	Information displayed
ifconfig -a	Interface status (up or down)Configuration parameters
ifstat	 Packets sent and received Collisions and other errors Negotiated media type settings between storage system interfaces and link partners
netstat	 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 the man pages for these commands.

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

Next topics

Viewing statistics of all active TCP connections on page 53 Viewing or clearing network interface statistics on page 54 Viewing network interface information with FilerView on page 57

Viewing statistics of all active TCP connections

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

Step

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

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

Example

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

syste	eml> ne	tstat -aM									
Activ	ve TCP	connection	s (inc	luding servers)							
Ctx	Local	Address		Remote Address		Swind	Sen	d-Q	Rwind	Recv-Q	State
lg	*.443			*.*		0		0	0	0	LISTEN
lg	*.22			*.*		0		0	0	0	LISTEN
lg	*.1056	8		*.*		0		0	0	0	LISTEN
lg	*.1056	9		*.*		0		0	0	0	LISTEN
lg	*.1056	7		*.*		0		0	0	0	LISTEN
lg	*.1057	1		*.*		0		0	0	0	LISTEN
lg	*.8514			*.*		0		0	0	0	LISTEN
lg	*.514			*.*		0		0	0	0	LISTEN
lg	*.23			*.*		0		0	0	0	LISTEN
lg	*.8023			*.*		0		0	0	0	LISTEN
lq	*.4047			*.*		0		0	0	0	LISTEN
lq	*.4045			*.*		0		0	0	0	LISTEN
lq	*.4046			*.*		0		0	0	0	LISTEN
lq	*.2049			*.*		0		0	0	0	LISTEN
la	*.111			* *		0		0	0	0	LISTEN
la	*.2807	3		* *		0		0	0	0	LISTEN
lq	*.3224	3		* *		Ō		Ō	0	0	LISTEN
la	*.2289	9		* *		0		0	0	0	LISTEN
1	192 16	8 1 72 204	9	192 168 1 36 800		33952		328	26280	0	ESTABLISHED
10	* 2049	0.11, 2.201	-	* *		00000		0_0	0	0	LISTEN
-9	.2019			•		0		0	0	0	D1010K
Activ		sockets (i	ncludi	ng servers)							
Loca	l Addre	ad Docucep (I	Remot	e Address	Send-	O Recv	-0				
* 10	570	55	* *	ie maarebb	bena	0	ñ				
* 69	570		• * *			0	0				
* 16	1		* *			0	0				
* 404	10		• * *			0	0				
* 40	17		• * *			0	0				
* 40	1 .		* *			0	0				
* 40	10		* *			0	0				
* 20	±0 40		* *			0	0				
* 11	1		* *			0	0				
* 011			··· • •			0	0				
* . ZI:	000		^.^ + +			0	0				
^.5Z	J		^ • ^			0	U				
The	followi	ng exampl	e shov	ws the output of the	ne net	tstat	-aE	s co	mmand	:	
nets	tat -aB										
Activ	ve TCP	connection	s (inc	luding servers)							
Loca	l Addre	SS	Remot	e Address	Swir	d Send	-0 1	Rwin	d Recv-	0 State	e
Sent	Rcvd						~			-	

54 | Data ONTAP 7.3 Network Management Guide

localhost-20.1023 0 0	localhost-10.671	65535	0	8760	0 ESTABLISHED	
localhost-20.8514 1 44	localhost-10.626	66608	1	8760	0 ESTABLISHED	
localhost-20.18576 9284 606K	localhost-10.7951	66608	0	8760	0 ESTABLISHED	
*.10568 0 0	*.*	0	0	0	0 LISTEN	
*.10569 0 0	*.*	0	0	0	0 LISTEN	
*.10567 0 0	*.*	0	0	0	0 LISTEN	
*.22 0 0	*.*	0	0	0	0 LISTEN	
*.443 0 0	*.*	0	0	0	0 LISTEN	
*.8514 0 0	*.*	0	0	0	0 LISTEN	
*.514 0 0	*.*	0	0	0	0 LISTEN	
*.23 0 0	*.*	0	0	0	0 LISTEN	
*.8023 0 0	*.*	0	0	0	0 LISTEN	
*.32243 0 0	*.*	0	0	0	0 LISTEN	
*.22899 0 0	* . *	0	0	0	0 LISTEN	
Active UDP sockets (including servers)						
Local Address *.10570	Remote Address *.*	Send-Q Rec 0	v-Q 0			
*.69 *.161	*.*	0 0	0 0			

Viewing or clearing network interface statistics

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

About this task

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

Because the statistics displayed by the ifstat command are cumulative, a giveback does not cause the statistics to zero out.

Step

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

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

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

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

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

ifstat e0a

An output similar to the following is displayed.

system1> ifstat e0a

interface e0a	(8 days, 20 hours, 10 minu	tes, 27 seconds)
RECEIVE		
Frames/second:	13 Bytes/second:	800 Errors/
Discards/minute: 62415k	0 Total frames:	897k Total bytes:
Total errors:	0 Total discards:	0 Multi/
broadcast: 734k		
No buffers:	0 Non-primary u/c:	0 Tag
drop: 0		
Vian tag drop:	0 Vlan untag drop:	U CRC
Runt frames:	0 Fragment:	0 J. Long
frames: 0		
Jabber:	0 Alignment errors:	0 Bus
overruns: 0)	
Queue overflows:	0 Xon:	0
Xoff:	0	
Jumbo:	0 Reset:	0
Reset1:	0	
Reset2:	0	
TRANSMIT		
Frames/second:	2 Bytes/second:	110 Errors/
minute: 0		
Discards/minute:	U TOTAL TRAMES:	153K TOTAL DYTES:
Total errorg:	0 Total discards:	$0 \mid Multi/$
broadcast: 9478	o iocai discaids:	o Marci/

Queue overflows:	0	No buffers:	0	Max
collisions: 0				
Single collision:	0	Multi collisions:	0	Late
collisions: 0				
Timeout:	0	Xon:	0	
Xoff:	0			
Jumbo:	0			
LINK_INFO				
Current state:	up	Up to downs:	0	
Auto:	on			
Speed:	1000m	Duplex:	full	
Flowcontrol:	none			

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

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

system1> ifstat e0a

interface e0a	(0 ho	urs, 0 minutes, 8 seco	nds)	
RECEIVE				
Frames/second:	1	Bytes/second:	32	Errors/
minute: 0 Discards/minute:	0	Total frames:	7	Total
bytes: 448	U		,	IOCUI
Total errors:	0	Total discards:	0	Multi/
broadcast: 0	0		0	
No bullers:	0	Non-primary u/c:	0	Tag
Vlan tag drop:	0	Vlan untag drop:	0	CRC
errors: 0				
Runt frames:	0	Fragment:	0	Long
Irames: U Jabber:	0	Alignment errors:	0	Bug
overruns: 0	0		0	Dub
Queue overflows:	0	Xon:	0	
Xoff:	0			
Jumbo:	0	Reset:	0	
Reset1:	0			
TRANSMIT	0			
Frames/second:	1	Bytes/second:	17	Errors/
minute: 0				· ·
Discards/minute:	0	Total frames:	4	Total
bytes: 361	0	Total diggarda.	0	Multi/
broadcast: 0	0	IOCAI discards.	0	MUICI/
Queue overflows:	0	No buffers:	0	Max
collisions: 0		'		
Single collision:	0	Multi collisions:	0	Late
collisions: 0	0	Verst	0	1
Yoff:	0		0	I
Jumbo:	0			

LINK_INFO		
Current state:	up Up to downs:	0
Auto:	on	
Speed:	1000m Duplex:	full
Flowcontrol:	none	

Related references

Statistics for Gigabit Ethernet controller IV - VI and G20 interfaces on page 183 Statistics for the BGE 10/100/1000 Ethernet interface on page 193

Viewing network interface information with FilerView

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

About this task

The Network Report in FilerView provides the same information that you get by running the netstat -1, routed status, and netstat -rn commands.

Step

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

The Network Report displays the interface statistics and routing tables.

How routing in Data ONTAP works

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

Data ONTAP uses two routing mechanisms:

- Fast pathData ONTAP uses this mechanism to route NFS packets over UDP and to route all
TCP traffic.
- **Routing table** To route IP traffic that does not use fast path, Data ONTAP uses the information available in the local routing table. The routing table contains the routes that have been established and are currently in use, as well as the default route specification.

Next topics

What fast path is on page 59 How to manage the routing table on page 61 Specifying the default route on page 63 How to enable or disable routing mechanisms on page 64 How to view the routing table and default route information on page 65 Modifying the routing table on page 68

What fast path is

Fast path is an alternative routing mechanism to the routing table, in which the responses to incoming network traffic are sent back by using the same interface as the incoming traffic. It provides advantages such as load balancing between multiple network interfaces and improved storage system performance.

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

Note: Fast path is supported over IPv6.

Using fast path provides the following advantages:

- Load balancing between multiple network interfaces on the same subnet.
 Load balancing is achieved by sending responses on the same interface of your storage system that receives the incoming requests.
- Increased storage system performance by skipping routing table lookups.

How fast path works with NFS-over-UDP

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

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

How fast path works with TCP

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

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

How fast path affects Telnet sessions and the ping utility

If fast path is enabled and the default router stops working, you cannot use the ping utility to communicate with your storage system. However, the Telnet sessions to your storage system can still be established from a non-local subnet. This happens because the ping utility uses routing table lookups.

Fast path not compatible with asymmetric routing

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

Related tasks

Enabling or disabling fast path on page 64

Similarities and differences between fast path over IPv4 and IPv6

Starting with Data ONTAP 7.3.3, fast path is supported over IPv6. Fast path over IPv4 and IPv6 provide improved storage system performance. However, fast path over IPv6 does not provide load balancing between multiple interfaces like IPv4 does.

Similarities between fast path over IPv4 and IPv6

Fast path over IPv4 and IPv6 provide improved system performance because of the following reasons:

• When fast path is enabled, TCP checksum computation is automatically offloaded to the network interfaces.

Note: Only specific NICs support this functionality.

• Route lookup to the final destination is skipped when fast path is enabled.

Differences between fast path over IPv4 and IPv6

Fast path over IPv4 provides load balancing between multiple network interfaces on the same subnet because responses are sent on the same network interface that receives the incoming requests. IPv4 uses the same source IPv4 address and the source MAC address of the incoming packet in the destination packet.

Fast path over IPv6 does not provide load balancing because it uses the default gateway of the incoming interface as the destination. Fast path over IPv6 always performs an NDP lookup to find the MAC address of the next hop. Therefore, the responses might not be sent on the same interface that receives the request.

How to manage the routing table

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

Next topics

What the routed daemon does on page 61 When the routed daemon should be turned off on page 62 How dynamic routing works for IPv6 on page 62 Routing tables in a vFiler unit environment on page 62 Circumstances that might alter the routing table on page 63

What the routed daemon does

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

The routed daemon performs the following functions by default:

- Deletes redirected routes after a specified period
- Performs router discovery with ICMP Router Discovery Protocol (IRDP) This is useful only if there is no static default route.
- Listens for Routing Information Protocol (RIP) packets
- Migrates routes to alternate interfaces when multiple interfaces are available on the same subnet

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

- Control RIP and IRDP behavior
- Generate RIP response messages that update a host route on your storage system
- Recognize distant gateways identified in the /etc/gateways file

Note: The routed daemon supports only IPv4.

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

When the routed daemon should be turned off

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

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

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

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

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

Related tasks

Enabling or disabling the routed daemon from the command-line interface on page 64 *Enabling or disabling the routed daemon with FilerView* on page 65

How dynamic routing works for IPv6

IPv6 routing table entries are created by default when you enable IPv6. Additional entries are added dynamically in the routing table on receiving Router Advertisement and ICMP redirect messages.

Storage systems populate a default router list and a prefix list, based on the information in the Router Advertisement messages. The default router list is used to select a router for off-link destinations, and the prefix list is used to determine whether a destination address is on-link.

Related tasks

Enabling or disabling IPv6 on page 33

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

Circumstances that might alter the routing table

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

The routing table might be modified in the following circumstances:

- A new interface is configured with the ifconfig command and there are no existing entries for the new network number in the routing table.
- You use the route add command to add an entry to the routing table.
- Your storage system receives an ICMP/ICMPv6 redirect packet, which notifies the storage system of a better first-hop router for a particular destination.

Note: Your storage system ignores ICMP/ICMPv6 redirect packets if the ip.icmp_ignore_redirect.enable option is turned on.

- Your storage system is rebooted after the default route in the /etc/rc file is modified.
- The default route is added to the routing table on receiving an IPv6 Router Advertisement message.

Specifying the default route

The default route entry routes to destinations that are not listed in the routing table. You can specify the default route in Data ONTAP either during initial setup or by modifying the /etc/rc file.

About this task

If IPv6 is enabled on your storage system, the default route is automatically generated.

Steps

- 1. Open the /etc/rc file in the root volume by using a text editor.
- 2. Add the following command to the /etc/rc file:

```
route add default route_IP
```

route_IP is the IP address of the default route

Example

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

```
hostname sys1
ifconfig e0 192.0.2.21 netmask 255.255.255.0 mediatype 100tx-fd
route add default 192.0.2.1 1
routed on
```

How to enable or disable routing mechanisms

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

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

Next topics

Enabling or disabling fast path on page 64 *Enabling or disabling the routed daemon from the command-line interface* on page 64 *Enabling or disabling the routed daemon with FilerView* on page 65

Enabling or disabling fast path

Fast path provides advantages such as load balancing and improved storage system performance. You can enable or disable fast path by using the options ip.fastpath.enable command.

Step

1. Enter the following command from the command-line interface:

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

on-Enables fast path

off—Disables fast path

Note: You can use the -x option with the netstat command to check if the fast path mechanism is enabled.

Related concepts

What fast path is on page 59

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

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

About this task

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

Step

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

routed {on|off}

on-Turns on the routed daemon

off—Turns off the routed daemon

Related concepts

What the routed daemon does on page 61 *When the routed daemon should be turned off* on page 62

Related references

The routed daemon on page 211

Enabling or disabling the routed daemon with FilerView

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

Steps

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

Related concepts

What the routed daemon does on page 61 When the routed daemon should be turned off on page 62

Related references

The routed daemon on page 211

How to view the routing table and default route information

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

Next topics

Viewing the routing table from the command-line interface on page 66 Viewing the default route information from the command-line interface on page 67 Viewing the routing table and routing information by using FilerView on page 68

Viewing the routing table from the command-line interface

You can view information such as default route and the routes for specific destination addresses. If you have enabled the IPv6 option, the routing table displays both the IPv4 and IPv6 information.

Step

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

- netstat -rn
- route -s

Example for interpreting the routing table The output of the netstat -rn command is as follows: netstat -rn Routing tables Internet: Destination Gateway Flags Refs Use Interface default 192.0.2.1 UGS 3 21397 e0a 127.0.0.1 127.0.0.1 UH 0 lo 0 192.0.2/24 link#11 0 UC 0 e0a 192.0.2.1 0:d0:d3:0:30:0 UHL 1 e0a 0 192.0.2.23 0:1:30:b8:30:c0 UHL 0 \cap e0a 192.0.2.24 0:1:30:b8:2e:c0 UHL 0 0 e0a Internet v6: Destination Flags Use Interface Refs Gateway default fe80::21b:2bff:fed7:ec00%e1a UG 0 0 e1a ::1 ::1 0 lo UH 0 2001:0db8::/64 link#3 UC 0 0 ela 2001:0db8:b255:4213::/64 link#3 UC 0 0 ela

UHL

2001:0db8:b255:4213::1 link#3 0 0 e1a

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

The IPv6 routing table also has the same network parameters except that the network mask is replaced by the prefix length for that network.

Routing table flags

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

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

Viewing the default route information from the command-line interface

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

Step

1. Enter the following command:

routed status

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

Example

The output of the routed status command is as follows:

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

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

Viewing the routing table and routing information by using FilerView

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

Step

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

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

Modifying the routing table

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

Step

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

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

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

Example

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

route add 192.0.2.25 gateway.com 1

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

```
systeml> netstat -rn
Routing tables
```

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

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

Related tasks

Changing the status of an interface on page 50

Related references

Routing table flags on page 67

How to maintain host-name information

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

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

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

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

Next topics

How the /etc/hosts file works on page 71 How to configure DNS to maintain host information on page 74 How to use dynamic DNS to update host information on page 78 How to use NIS to maintain host information on page 81 How to configure NIS with Data ONTAP interfaces on page 85 What NIS information you can view on page 88 Configuring DNS and NIS with FilerView on page 89 How to change the host-name search order on page 90

How the /etc/hosts file works

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

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

- localhost
- All interfaces on your storage system

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

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

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

• Locally—You can add entries by using the command-line interface or FilerView.

• Remotely—If the file has many entries and you have access to an NIS makefile master, you can use the NIS makefile master to create the /etc/hosts file. This method prevents errors that might be caused by editing the file manually.

Next topics

Adding a host name in the /etc/hosts file on page 72 Hard limits for the /etc/hosts file on page 73 Editing the /etc/hosts file with FilerView on page 73 Changing the host name of a storage system on page 73

Adding a host name in the /etc/hosts file

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

About this task

During setup, if you enable IPv6 on the storage system and configure IPv6 addresses for your network interfaces, these IPv6 addresses are also added to the /etc/hosts file.

Step

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

IP_address host_name aliases

IP_address is the IP address of the host.

host_name is the name of the host.

aliases are the alias names for the host.

Example

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

192.0.2.16 myhost newhost myhost-e0a

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

The following is a sample /etc/hosts file:

```
#Auto-generated by setup Tue Apr 21 17:41:40 IST 2009
127.0.0.1 localhost
192.0.2.16 myhost myhost-e0a
# 0.0.0.0 myhost-e0b
# 0.0.0.0 myhost-e0c
# 0.0.0.0 myhost-e0d
```
The following is a sample /etc/hosts file in which an IPv6 address is also configured for the interface e0a:

```
#Auto-generated by setup Tue Apr 21 17:41:40 IST 2009
127.0.0.1 localhost
192.0.2.16 myhost myhost-e0a
2001:0db8::95 myhost myhost-e0a
# 0.0.0.0 myhost-e0b
# 0.0.0.0 myhost-e0c
# 0.0.0.0 myhost-e0d
```

Hard limits for the /etc/hosts file

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

The hard limits are as follows:

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

Note: There is no limit on file size.

Editing the /etc/hosts file with FilerView

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

Steps

- 1. In FilerView, click **Network** in the list on the left pane.
- 2. In the list under Network, click Manage Hosts File.
- 3. Click in the hosts window, then click **Insert**.
- 4. Complete the fields in the Create a New /etc/hosts Line window for each host you want to add and click OK.
- 5. Click Apply in the Manage Hosts File window.

Changing the host name of a storage system

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

Steps

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

2. Enter the following command to specify a new name for the host:

hostname new_name

new_name is the new host name of the storage system.

3. Reboot the storage system.

Attention: Ensure that you complete both steps before rebooting the storage system. If you skip Step 2 and then reboot the storage system, any manual or scheduled SnapMirror operations might fail. Use the hostname command to specify the correct name before any SnapMirror operations are initiated.

How to configure DNS to maintain host information

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

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

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

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

Note: You can configure IPv4 and IPv6 addresses as DNS server addresses.

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

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

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

- Using FilerView
- Using the command-line interface

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

Correct host-name resolution depends on correctly configuring 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(1) and na_dns(8) man pages.

Next topics

Configuring DNS from the command-line interface on page 75 How DNS resolves host names on page 76 DNS name caching on page 77 DNS information you can view on page 77

Related concepts

How the /etc/hosts file works on page 71

Configuring DNS from the command-line interface

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

Steps

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

If you are	Then
Creating the /etc/ resolv.conf file	By using a text editor, create the /etc/resolv.conf file in the root volume. The file can consist of up to three lines, each specifying a name server host in the following format.
	nameserver ip_address
	$ip_address$ is the IP address of the DNS name server. The IP address can be an IPv4 or an IPv6 address.
	Note: If an IPv6 link-local address is specified as a DNS name server, the address must be appended with <i>%interface_name</i> . The appended <i>interface_name</i> is the name of the interface on the storage system that is connected to the same link as the specified DNS name server. For example:
	nameserver 2001:0db8:85a3:0:0:8a2e:0370:99
	e0a is the interface on the storage system that is connected to the same link as the DNS name server with the IPv6 address 2001:0db8:85a3:0:0:8a2e: 0370:99.

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

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

options dns.domainname domain

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

3. Enter the following command to enable DNS:

options dns.enable {on|off} on—Enables DNS

off—Disables DNS

Hard limits for the /etc/resolv.conf file

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

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

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

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

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

Note: There is no limit on file size.

How DNS resolves host names

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

For resolving IPv4 addresses, DNS uses the A record. The A record can store a 32-bit address and can resolve IPv4 addresses. To resolve IPv6 addresses, DNS uses the AAAA record. The AAAA record can store a 128-bit address and can resolve IPv6 addresses.

IPv4 reverse DNS lookups use the in-addr.arpa domain. An IPv4 address is represented in the inaddr.arpa domain by a sequence of bytes, represented as decimal numbers, in reverse order. The numbers are separated by dots and end with the suffix .in-addr.arpa. IPv6 reverse DNS lookups use the ip6.arpa domain. An IPv6 address is represented as a name in the ip6.arpa domain by a sequence of nibbles, represented as hexadecimal digits, in reverse order. These nibbles are separated by dots and end with the suffix .ip6.arpa.

The following table shows sample IPv4 and IPv6 addresses and their reverse DNS lookups:

IP address	Reverse lookup domain name
192.0.2.10	10.2.0.192.in-addr.arpa
2001:0db8:85a3:0:0:8a2e:0370:99	9.9.0.0.0.7.3.0.e.2.a.8.0.0.0.0.0.0.0.3.a.5.8.8.b.d. 0.1.0.0.2.ip6.arpa

DNS name caching

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

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

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

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

DNS information you can view

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

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

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

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

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

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

How to use dynamic DNS to update host information

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

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

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

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

Next topics

How dynamic DNS updates work in Data ONTAP on page 79 Support for dynamic DNS updates in Data ONTAP on page 79 Enabling or disabling dynamic DNS updates on page 80 Disabling the transmission of DNS updates for an IP address on page 80 Changing the time-to-live setting for DNS entries on page 81

How dynamic DNS updates work in Data ONTAP

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

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

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

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

Support for dynamic DNS updates in Data ONTAP

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

The following conditions apply to dynamic DNS updates:

- By default, dynamic DNS updates are disabled in Data ONTAP.
- Dynamic DNS updates are supported on UNIX and Windows systems.
- On Windows DNS servers, secure dynamic DNS updates can be used to prevent malicious updates on the DNS servers. Kerberos is used to authenticate updates.
 Even if secure dynamic DNS updates are enabled, your storage system initially tries sending updates in clear text. If the DNS server is configured to accept only secure updates, the updates sent in clear text are rejected. Upon rejection, the storage system sends secure DNS updates.
- For secure dynamic DNS updates, your storage system must have CIFS running and must be using Windows Domain authentication.
- Dynamic DNS updates can be sent for the following:
 - Physical interfaces
 - vif and VLAN interfaces
 - vFiler units
- You cannot set TTL values for individual vFiler units. All vFiler units inherit the TTL value that is set for vfiler0, which is the default vFiler unit and is the same as the physical storage system.
- DHCP addresses cannot be dynamically updated.
- In a takeover situation, the hosting storage system is responsible for sending DNS updates for IP addresses for which it is responding.

• For both manual and autoconfigured global IPv6 unicast addresses, the dynamic DNS update is sent after Duplicate Address Detection is performed. For IPv6 addresses of any other type and scope, your storage system does not send any dynamic DNS update.

Enabling or disabling dynamic DNS updates

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

Step

1. Enter the following command:

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

on-Enables dynamic DNS updates

off-Disables dynamic DNS updates

secure—Enables secure dynamic DNS updates

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

Disabling the transmission of DNS updates for an IP address

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

About this task

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

You can also disable dynamic DNS updates for an IPv6 address.

Step

1. Enter the following command:

ifconfig interface_name no_ddns IP_address

interface_name is the name of the interface.

IP_address is the IP address of the interface.

Example

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

ifconfig e0a no_ddns 192.0.2.30

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

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

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

Changing the time-to-live setting for DNS entries

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

Step

1. Enter the following command:

```
options dns.update.ttl time
```

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

Example

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

options dns.update.ttl 2h

Related concepts

How dynamic DNS updates work in Data ONTAP on page 79

How to use NIS to maintain host information

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

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

- Using the /etc/hosts file on the NIS server You can download the /etc/hosts file on the NIS server o your storage system's default volume for local host-name lookup.
- Using a hosts map that is maintained as a database on the NIS server The storage system uses the hosts map to query during a host lookup request across the network.
- Using the ipnodes map that is maintained as a database on the NIS server The ipnodes map is used for host lookup when IPv6 is enabled on your storage system.

Note: The ipnodes database is supported only on Solaris NIS servers. To resolve a host name to an address, your storage system (with IPv6 enabled) first looks in the ipnodes database. If

the IP address is not present in the ipnodes database, the application looks in the hosts database. However, if IPv6 is not enabled, then your storage system looks only in the hosts database and does not refer the ipnodes database.

Next topics

How using NIS slaves can improve performance on page 82 How an NIS master is selected on page 83 Creating /etc/hosts from the NIS master on page 83 Guidelines for using NIS slaves on page 83 NIS administrative commands on page 84

How using NIS slaves can improve performance

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

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

• Downloading the maps from the NIS master server

You can download the maps from the NIS master server to the NIS slave by running the yppush command from the NIS server. You can also download the maps by disabling and then enabling the NIS slave from your storage system. After the maps are downloaded, they are stored in the / etc/yp/nis_domain_name directory. The NIS slave then services all the NIS requests from your storage system by using these maps. The NIS slave checks the NIS master every 45 minutes for any changes to the maps. If there are changes, they are downloaded.

· Listening for updates from the NIS master

When the maps on the NIS master are changed, the NIS master administrator can optionally notify all slaves. Therefore, in addition to periodically checking for updates from the NIS master, the NIS slave also listens for updates from the master.

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

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

Related concepts

Guidelines for using NIS slaves on page 83

Related tasks

Enabling an NIS slave on your storage system on page 87

How an NIS master is selected

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

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

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

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

Creating /etc/hosts from the NIS master

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

Steps

- 1. On the NIS server, open the NIS Makefile with a text editor.
- 2. Locate the section for hosts.time.
- 3. Add the following lines at the end of the hosts.time section, replacing *dirname* with a directory name of your choice, and *toaster 1*, *toaster2*, and so on with names of the storage systems:

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

4. Save the NIS Makefile.

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

Related concepts

How the /etc/hosts file works on page 71

Guidelines for using NIS slaves

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

The following guidelines apply when using the NIS slave:

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

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

You can use the NIS slave for the following:

- Vifs and VLAN interfaces
- vFiler units
- Active/active configurations

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

Related concepts

How using NIS slaves can improve performance on page 82

NIS administrative commands

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

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

Command	Function
ypcat	Prints an entire NIS map.
урдгоир	Displays the NIS group cache entries.
ypmatch	Looks up specific entries in an NIS map.

Command	Function
ypwhich	Returns the name of the current NIS server.

How to configure NIS with Data ONTAP interfaces

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

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

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

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

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

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

Next topics

Enabling or disabling NIS using the command-line interface on page 85 Specifying the NIS domain name on page 86 Specifying NIS servers to bind to your storage system on page 86 Enabling an NIS slave on your storage system on page 87

Enabling or disabling NIS using the command-line interface

You can enable or disable NIS on your storage system for host-name resolution.

Step

1. Enter the following command:

options nis.enable {on|off}

on-Enables NIS

off—Disables NIS

Specifying the NIS domain name

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

Step

1. Enter the following command:

options nis.domainname domain

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

Specifying NIS servers to bind to your storage system

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

About this task

Keep the following in mind before performing the binding procedure:

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

Step

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

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

The asterisk (*) specifies that broadcast (for IPv4) and multicast (for IPv6) is used to bind to NIS servers if the servers in the list are not responding. The '*' is the default value. If you do not specify the broadcast or multicast option, and none of the listed servers is responding, NIS services are disrupted until one of the preferred servers responds.

Example

The following command lists two servers and uses the default broadcast (multicast for IPv6) option:

options nis.servers 192.0.2.1, nisserver-1,*

Your storage system first tries to bind to 192.0.2.1. If the binding fails, the storage system tries to bind to nisserver-1. If this binding also fails, the storage system binds to any server that

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

The following command lists an NIS server with an IPv6 address and uses the default multicast option:

```
options nis.servers 2001:0db8:85a3:0:0:8a2e:0370:99,*
```

Related concepts

How an NIS master is selected on page 83

Enabling an NIS slave on your storage system

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

About this task

If you enable IPv6 on your storage system, your storage system can have multiple addresses configured for it in the host-name database. These addresses appear in the host-name lookup, depending on the following conditions:

- If you disable the NIS slave, you can obtain all the addresses from either the hosts database or the ipnodes database in the NIS server.
- If you disable the NIS slave, your storage system reverts to the original configuration, in which it contacts an NIS server to resolve host names.
- If you enable the NIS slave, only the last address from the list of addresses available in the /etc/ hosts file is stored for a host name in the host database downloaded to your system.
- If you enable the NIS slave, a maximum of three addresses are stored for a host name in the ipnodes database downloaded to your system. At least one address from each address family is stored.

Step

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

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

Related concepts

How using NIS slaves can improve performance on page 82 *Guidelines for using NIS slaves* on page 83

What NIS information you can view

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

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

- NIS domain name
- Last time the local group cache was updated
- Information about each NIS server that was polled by your storage system:
 - IP address of the NIS server
 - Type of NIS server
 - State of the NIS server
 - Whether your storage system is bound to the NIS server
 - Time of polling
- Information about the NIS netgroup cache:
 - Status of the cache
 - Status of the "*.*" entry in the cache
 - Status of the "*.nisdomain" entry in the cache
- Whether an NIS slave is enabled
- NIS master server
- Last time the NIS map was checked by the NIS slave
- NIS performance statistics:
 - Number of YP lookup network retransmission
 - Total time spent in YP lookups
 - Number of network retransmission
 - Minimum time spent in a YP lookup
 - Maximum time spent in a YP lookup
 - Average time spent in a YP lookup
- Response statistics for the three most recent YP lookups

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

Viewing NIS performance statistics

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

Step

1. Enter the following command to view NIS information:

nis info

Example

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

```
system1*> nis info
NIS domain is lab.example.com
   NIS group cache has been disabled
     IP Address Type State Bound Last Polled
Client
calls Became Active
    192.0.2.12 PREF ALIVE YES Mon Jan 23 23:11:14
GMT 2008 0 Fri Jan 20 22:25:47 GMT 2008
                 NIS Performance Statistics:
              Number of YP Lookups: 153
               Total time spent in YP Lookups: 684 ms, 656 us
              Number of network re-transmissions: 0
              Minimum time spent in a YP Lookup: 0 ms, 1 us
              Maximum time spent in a YP Lookup: 469 ms, 991 us
              Average time spent in YP Lookups: 4 ms, 474 us
                  3 Most Recent Lookups:
              [0] Lookup time: 0 ms, 1 us Number of network re-
transmissions: 0
              [1] Lookup time: 5 ms, 993 us Number of network re-
transmissions: 0
              [2] Lookup time: 0 ms, 1 us Number of network re-
transmissions: 0
NIS netgroup (*.* and *.nisdomain) cache status:
uninitialized
         *.* eCode: 0
         *.nisdomain eCode: 0
  NIS Slave disabled
```

Configuring DNS and NIS with FilerView

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

Steps

- 1. Click **Network** in the list on the left pane.
- 2. In the list under Network, click Configure Host Name Resolution (DNS & NIS).

The Host Name Resolution Policy Wizard is displayed.

90 | Data ONTAP 7.3 Network Management Guide

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

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

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

Related concepts

How to configure DNS to maintain host information on page 74 DNS name caching on page 77 How to use dynamic DNS to update host information on page 78 How to configure NIS with Data ONTAP interfaces on page 85 Guidelines for using NIS slaves on page 83 How to change the host-name search order on page 90

How to change the host-name search order

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

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

hosts: files nis dns passwd: files nis ldap netgroup: files nis ldap group: files nis ldap shadow: files nis

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

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

- /etc/hosts file
- NIS
- DNS

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

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

Next topics

Changing the host-name search order with FilerView on page 91 *Changing the host-name search order* on page 91

Changing the host-name search order with FilerView

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

Steps

- 1. From the list on the left pane, click Network.
- 2. In the list under Network, click Manage DNS and NIS Name Service.
- 3. In the Name Service section, select the desired values from the Hosts drop-down list.

Changing the host-name search order

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

Steps

- 1. If the /etc/nsswitch.conf file does not exist in the root volume of the storage system, create it.
- 2. Edit the file, entering each line in the following format:

hosts: service

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

3. Save the file.

Example

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

hosts: nis

How VLANs work

VLANs provide logical segmentation of networks by creating separate broadcast domains. A VLAN can span multiple physical network segments. The end-stations belonging to a VLAN are related by function or application.

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

Next topics

VLAN membership on page 93 GARP VLAN Registration Protocol on page 95 VLAN tags on page 95 Advantages of VLANs on page 96 Prerequisites for setting up VLANs on page 97 Guidelines for setting up VLANs in Data ONTAP on page 97 The vlan command syntax on page 98 Creating a VLAN on page 98 Configuring a VLAN on page 100 Adding an interface to a VLAN on page 101 Deleting VLANs on page 102 Enabling or disabling GVRP on your VLAN interface on page 103 Viewing VLAN statistics on page 104 Viewing statistics for a specific VLAN on page 104

VLAN membership

An end-station must become a member of a VLAN before it can share the broadcast domain with other end-stations on that VLAN. The switch ports can be configured to belong to one or more VLANs (static registration), or end-stations can register their VLAN membership dynamically, with VLAN-aware switches.

VLAN membership can be based on one of the following:

- Switch ports
- End-station MAC addresses
- Protocol

In Data ONTAP, VLAN membership is based on switch ports. With port-based VLANs, ports on the same or different switches can be grouped to create a VLAN. As a result, multiple VLANs can exist on a single switch.

How VLAN membership affects communication

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

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



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

Related concepts

VLAN membership on page 93

GARP VLAN Registration Protocol

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

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

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

GVRP configuration for VLAN interfaces

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

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

Related tasks

Enabling or disabling GVRP on your VLAN interface on page 103

VLAN tags

A VLAN tag is a unique identifier that indicates the VLAN to which a frame belongs. Generally, a VLAN tag is included in the header of every frame sent by an end-station on a VLAN.

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



For example, in this figure, if a member of VLAN 10 on Floor 1 sends a frame for a member of VLAN 10 on Floor 2, Switch 1 inspects the frame header for the VLAN tag (to determine the VLAN) and the destination MAC address. The destination MAC address is not known to Switch 1. Therefore, the switch forwards the frame to all other ports that belong to VLAN 10, that is, port 4 of Switch 2 and Switch 3. Similarly, Switch 2 and Switch 3 inspect the frame header. If the destination MAC address on VLAN 10 is known to either switch, that switch forwards the frame to the destination. The end-station on Floor 2 then receives the frame.

Advantages of VLANs

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

VLANs provide the following advantages:

• Ease of administration

VLANs enable logical grouping of end-stations that are physically dispersed on a network. When users on a VLAN move to a new physical location but continue to perform the same job function, the end-stations of those users do not need to be reconfigured. Similarly, if users change their job

function, they need not physically move: changing the VLAN membership of the end-stations to that of the new team makes the users' end-stations local to the resources of the new team.

- Confinement of broadcast domains VLANs reduce the need to have routers deployed on a network to contain broadcast traffic. Flooding of a packet is limited to the switch ports that belong to a VLAN.
- Reduction in network traffic Confinement of broadcast domains on a network significantly reduces traffic.
- Enforcement of security policies
 By confining the broadcast domains, end-stations on a VLAN are prevented from listening to or
 receiving broadcasts not intended for them. Moreover, if a router is not connected between the
 VLANs, the end-stations of a VLAN cannot communicate with the end-stations of the other
 VLANs.

Prerequisites for setting up VLANs

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

The following are the prerequisites for setting up VLANs:

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

Guidelines for setting up VLANs in Data ONTAP

VLANs in Data ONTAP are implemented in compliance with the IEEE 802.1Q standard.

You should keep these guidelines in mind when setting up VLANs in Data ONTAP:

- You cannot set up VLANs using the setup procedure. You must use the command-line interface or FilerView to create, change, or delete VLANs.
- You must add the commands to create VLANs on the storage system to the /etc/rc file to make the VLANs persistent across reboots.
- You can create any number of VLANs on a NIC (supporting IEEE 802.1Q) on the storage system.

However, Data ONTAP imposes a limit on the number of interfaces (including physical, vif, VLAN, vh, and loopback interfaces) per storage system.

- You can create VLANs on physical interfaces and vifs.
- You can configure IPv4 and IPv6 addresses on a VLAN interface.

- You can use VLANs to support packets of different Maximum Transmission Unit (MTU) sizes on the same network interface.
 If a network interface is a member of multiple VLANs, you can specify different MTU sizes for individual VLANs.
- You can assign an identification number from 1 to 4094 to a VLAN.
- You must ensure that the interface on your storage system is also a member of its partner's VLANs in an active/active configuration.
- You cannot configure any parameters except mediatype for the physical network interface configured to handle VLANs.

Related concepts

Maximum number of network interfaces on page 25

The vlan command syntax

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

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

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

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

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

Creating a VLAN

You can create a VLAN for ease of administration, confinement of broadcast domains, reduced network traffic, and enforcement of security policies. You can use the vlan create command to

include an interface in one or more VLANs as specified by the VLAN identifier, enable VLAN tagging, and optionally enable GVRP.

About this task

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

Step

1. Enter the following command:

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

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

if name is the name of the network interface.

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

Example: Creating and naming of VLAN interfaces

Create VLANs with identifiers 10, 20, and 30 on the interface e4 of a storage system by using the following command:

```
vlan create e4 10 20 30
```

As a result, VLAN interfaces e4-10, e4-20, and e4-30 are created. The *ifconfig* command output displays e4 as a VLAN interface as follows:

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

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

vlan create reallylongname 100

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

After you finish

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

Related concepts

Prerequisites for setting up VLANs on page 97 Guidelines for setting up VLANs in Data ONTAP on page 97

Configuring a VLAN

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

About this task

You can configure the following parameters for a VLAN:

- IP address (IPv4 and IPv6)
- Network mask
- Prefix length
- Interface status
- Partner

Step

1. Enter the following command:

ifconfig ifname-vlanid IP_address netmask mask

ifname-vlanid is the VLAN interface name.

IP_address is the IP address for this interface.

mask is the network mask for this interface.

Example

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

vlan create e5a 1760

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

ifconfig e5a-1760 192.0.2.11 netmask 255.255.255.0

To configure the VLAN interface e5a-1760 with an IPv6 address, use the following command:

ifconfig e5a-1760 2001:0db8:85a3:0:0:8a2e:0370:99 prefixlen 64

Related concepts

Configuring network interfaces on page 39

IPv6 link-local addresses for VLANs

When IPv6 is enabled on your storage system, all VLANs share the same link-local address as the underlying network interface (physical or vif). When VLANs share the same link-local address, there are no address duplication (DAD) issues because link-local addresses cannot be routed and are confined to a LAN.

Related concepts

IPv6 address scopes on page 32

Related tasks

Enabling or disabling IPv6 on page 33

Adding an interface to a VLAN

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

About this task

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

Step

1. Enter the following command:

```
vlan add interface_name vlanid
```

interface_name is the name of the network interface.

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

Example

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

vlan add e4 40 50

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

After you finish

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

Related tasks

Configuring a VLAN on page 100 *Creating a VLAN* on page 98

Deleting VLANs

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

Step

1. Enter the following command:

vlan	delete	[-q]	interface_	name
------	--------	------	------------	------

If you want to	Enter the following command:	
Delete one or more specific VLANs	<pre>vlan delete [-q] interface_name vlanid Note: If you want to delete more than one specific VLAN, you can include a list of VLAN identifiers. For example, to delete the VLAN e4-30, enter the following command: vlan delete of 30</pre>	
Delete all VLANs	vian delete (-gl interfage name	
configured on a network interface	Vian defete [-d] interface_name	
	For example, to delete all VLANs configured on the interface e4, enter the following command:	
	vlan delete e4	

interface_name is the name of the network interface.

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

-q option invokes the quiet mode.

Result

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

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

Enabling or disabling GVRP on your VLAN interface

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

About this task

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

Step

1. Enter the following command:

vlan modify -g {on|off} adap_name

-g on enables GVRP.

-g off disables GVRP.

adap_name is the name of the network adapter.

Related concepts

GARP VLAN Registration Protocol on page 95 GVRP configuration for VLAN interfaces on page 95

Viewing VLAN statistics

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

Step

1. Enter the following command:

vlan stat interface_name

interface_name is the name of the network interface.

Example

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

```
vlan stat e4
Vlan Physical Interface: e4 (5 hours, 50 minutes, 38 seconds) --
Vlan IDs: 3,5
GVRP: enabled
RECEIVE STATISTICS
Total frames: 0 | Total bytes: 0 |Multi/broadcast: 0
Untag drops:0 | Vlan tag drops: 0
TRANSMIT STATISTICS
Total frames: 8 | Total bytes: 368
Vlan Interface: e4-3 (0 hours, 20 minutes, 45 seconds) --
ID: 3 MAC Address: 00:90:27:5c:58:14
```

Viewing statistics for a specific VLAN

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

Step

1. Enter the following command:

```
vlan stat interface_name vlanid
```

interface_name is the name of the network interface.

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

How vifs work in Data ONTAP

A virtual interface (vif) is 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 a vif is created, it is indistinguishable from a physical network interface.

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



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



Different vendors refer to vifs by the following terms:

- Virtual aggregations
- · Link aggregations
- Trunks
- EtherChannel

Vifs provide several advantages over individual network interfaces:

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

Next topics

Types of vifs on page 108Load balancing in multimode vifs on page 112Guidelines for configuring vifs on page 113The vif command on page 113Creating a single-mode vif on page 114Creating a static multimode vif on page 118Creating a dynamic multimode vif on page 119Adding interfaces to a vif on page 121Deleting interfaces from a vif on page 121Viewing vif status on page 122Viewing vif statistics on page 124Destroying a vif on page 125Second-level vifs in an active/active configuration on page 128

Types of vifs

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

Each vif provides different levels of fault tolerance. Multimode vifs provide methods for load balancing network traffic.

Starting with Data ONTAP 7.3.1, IPv6 supports both single-mode and multimode vifs.

Next topics

Single-mode vif on page 109 Static multimode vif on page 109 Dynamic multimode vif on page 110
Single-mode vif

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

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. Single-mode vifs also do not require a switch that supports link aggregation.

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



Static multimode vif

The static multimode vif implementation in Data ONTAP is in compliance with IEEE 802.3ad (static). Any switch that supports aggregates, but does not have control packet exchange for configuring an aggregate, can be used with static multimode vifs.

Static multimode vifs do not support IEEE 802.3ad (dynamic), also known as Link Aggregation Control Protocol (LACP). Port Aggregation Protocol (PAgP), the proprietary link aggregation protocol from Cisco, too is not supported.

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

Several load-balancing options are available to distribute traffic among the interfaces of a static multimode vif.

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

The following figure is an example of a static multimode vif. Interfaces e0, e1, e2, and e3 are part of the MultiTrunk1 multimode vif. All four interfaces in the MultiTrunk1 multimode vif are active.



Several technologies exist that enable traffic in a single aggregated link to be distributed across multiple physical switches. The technologies used to enable this capability vary among networking products. Static multimode vifs in Data ONTAP conform to the IEEE 802.3 standards. If a particular multiple switch link aggregation technology is stated to interoperate or conform to the IEEE 802.3 standards, it should operate with Data ONTAP.

Dynamic multimode vif

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

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

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

In a dynamic 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.

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

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

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

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

The following figure is an example of a dynamic multimode vif. Interfaces e0, e1, e2, and e3 are part of the MultiTrunk1 multimode vif. All four interfaces in the MultiTrunk1 dynamic multimode vif are active.



Several technologies exist that enable traffic in a single aggregated link to be distributed across multiple physical switches. The technologies used to enable this capability vary among networking products. Dynamic multimode vifs in Data ONTAP conform to the IEEE 802.3 standards. If a particular multiple switch link aggregation technology is stated to interoperate or conform to the IEEE 802.3 standards, it should operate with Data ONTAP.

Load balancing in multimode vifs

You can ensure that all interfaces of a multimode vif are equally utilized for outgoing traffic. You can use the IP address, MAC address, round-robin, or port based load-balancing methods to equalize traffic.

The load-balancing method 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.

Next topics

IP address and MAC address load balancing on page 112 *Round-robin load balancing* on page 112 *Port-based load balancing* on page 112

IP address and MAC address load balancing

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

These load-balancing methods use a fast hashing algorithm on the source and destination addresses (IP address and MAC address). If the result of the hashing algorithm maps to an interface that is not in the UP link-state, the next active interface is used.

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

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

Round-robin load balancing

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

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

Port-based load balancing

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

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

Guidelines for configuring vifs

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

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

• 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 can group up to 16 physical Ethernet interfaces on your storage system to obtain a vif.

- You cannot include a VLAN interface in a vif.
- The interfaces that form a vif must have the same Maximum Transmission Unit (MTU) size. If you attempt to create or add to a vif and the member interfaces have different MTU sizes, Data ONTAP automatically modifies the MTU size to be the same. To ensure that the desired MTU size is configured, you can use the ifconfig command to configure the MTU size of the vif after it is created. You need to configure the MTU size only if you are enabling jumbo frames on the interfaces.
- You can include any interface, except the eOM management interface that is present on some storage systems.
- 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

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

Command	Description
vif create [single multi lacp] vif_name-b [rr mac ip port] [interface_list]	Create a single-mode or multimode vif
vif {favor nofavor} interface_name	Designate a favored or nonfavored interface in a single-mode vif
vif add vif_name interface_list	Add network interfaces to a vif
vif deletevif_name interface_name	Delete a network interface from a vif

The following table gives the vif command syntax:

Command	Description
vif destroy vif_name	Destroy a vif
vif status[vif_name]	View the status of a vif
<pre>vif stat vif_name [interval]</pre>	View the statistics of data packets on the network interfaces of a vif

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

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

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

Creating a single-mode vif

You can 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. A single-mode vif increases the redundancy for providing more availability.

Before you begin

- 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.
- To make a specific interface active, you must specify that interface as preferred by using the vif favor command; otherwise, an interface is randomly selected to be the active interface.

Steps

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

ifconfig interface_list down

interface_list is a list of the interfaces you want as part of the vif.

Example

ifconfig e0a e0b down

2. To create a vif, enter the following command:

vif create single vif_name [interface_list]

vif_name is the name of the vif.

interface_list is a list of the interfaces you want as part of the vif.

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

3. To configure the vif, enter the following command:

ifconfig vif_name IP_address

vif_name is the name of the vif.

IP_address is the IP address for this interface.

Note: If you have enabled IPv6 on your storage system, you can create a vif and then configure the vif to the up status. After this, the vif has two IPv6 addresses automatically configured on it. Therefore, you need not manually configure the IP address for a vif.

Example: Creating a single-mode vif with an IPv4 address

1. To create a single-mode vif, enter the following command:

```
vif create single SingleTrunk1 e0 e1
```

2. To configure an IP address of 192.0.2.4 and a netmask of 255.255.255.0 on the single-mode vif SingleTrunk1, enter the following command:

ifconfig SingleTrunk1 192.0.2.4 netmask 255.255.255.0

Example: Creating a single-mode vif when IPv6 is enabled

1. To create a single-mode vif, enter the following command:

```
vif create single SingleTrunk1 e0 e1
```

- 2. Configure the vif by using one of the following methods:
 - To automatically configure the vif, configure the interface to the up status by entering the following command:

ifconfig SingleTrunk1 up

The vif now has two automatically configured addresses, as shown below:

inet6 2001:0db8:a0:98ff:fe08:618a prefixlen 64 scopeid 0x9 autoconf inet6 2001:0db8:a0:98ff:fe08:618a prefixlen 64 autoconf ether 02:a0:98:08:61:8a (Enabled virtual interface)

• To manually configure an IPv6 address of 2001:0db8:85a3:0:0:8a2e:0370:99 for the vif, enter the following command:

ifconfig SingleTrunk1 2001:0db8:85a3:0:0:8a2e:0370:99

Next topics

Selecting an active interface in a single-mode vif on page 116 Designating a nonfavored interface in a single-mode vif on page 117 Failure scenarios for a single-mode vif on page 117

Related concepts

Single-mode vif on page 109

Related tasks

Changing the status of an interface on page 50

Selecting an active interface in a single-mode vif

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

Step

1. Enter the following command:

vif favor interface_name

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

Example

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

vif favor el

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

Related concepts

Single-mode vif on page 109

Related tasks

Designating a nonfavored interface in a single-mode vif on page 117

Designating a nonfavored interface in a single-mode vif

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

About this task

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

Step

1. Enter the following command:

vif nofavor interface_name

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

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

Example

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

vif nofavor e2

Related concepts

Single-mode vif on page 109

Related tasks

Selecting an active interface in a single-mode vif on page 116

Failure scenarios for a single-mode vif

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

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

Link-monitoring ARP packets are sent over the ports of a single-mode vif to detect whether the ports are in the same broadcast domain. If these ARP packets do not reach any of the interfaces in the vif, the vif is configured to the down status. To avoid this problem, you must ensure that all the interfaces of a single-mode vif are in the same broadcast domain (for example, a LAN or a VLAN).

Related concepts

Single-mode vif on page 109

Related tasks

Viewing vif status on page 122

Creating a static multimode vif

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

Before you begin

You must 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 that you want to select as part of the vif.
- Configure all interfaces that are to be included in the vif to the down status, by using the ifconfig command.

About this task

You can improve throughput by creating a static multimode vif. With 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.

Steps

1. To create the vif, enter the following command:

```
vif create multi vif_name -b {rr|mac|ip|port} [interface_list]
```

-ь describes the load-balancing method.

rr specifies round-robin load balancing.

mac specifies MAC address load balancing.

Note: Do not select the MAC address load-balancing method when creating vifs on a storage system that connects directly to a router. In such a setup, for every outgoing IP frame, the

destination MAC address is the MAC address of the router. As a result, only one interface of the vif is used.

ip indicates IP address load balancing (default).

port indicates port-based load balancing.

vif_name is the name of a previously created vif.

interface_list is a list of the interfaces you want to add to the vif.

Example

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

vif create multi MultiTrunk1 -b mac e0 e1 e2 e3

2. To configure the vif, enter the following command:

ifconfig vif_name IP_address netmask mask

Related concepts

Static multimode vif on page 109 *Load balancing in multimode vifs* on page 112

Related tasks

Changing the status of an interface on page 50

Creating a dynamic multimode vif

With a dynamic 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. Dynamic multimode vifs can detect both loss of link and loss of data flow.

Before you begin

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

- Identify or install a switch that supports LACP over multiple port connections in your network, configured according to your switch vendor's instructions.
- Decide on a case-sensitive name for the 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 to select as part of the vif.

• Configure all interfaces that are to be included in the vif to the down status, by using the ifconfig command.

About this task

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

Steps

1. To create a dynamic multimode vif, enter the following command:

```
vif create lacp vif_name -b {rr|mac|ip|port} [interface_list]
```

-b specifies the load-balancing method.

rr specifies round-robin load balancing.

mac specifies MAC address load balancing.

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

ip specifies IP address based load balancing (default).

port indicates port-based load balancing.

vif_name is the name of a previously created vif.

interface_list is a list of the interfaces that form the vif.

Example

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

vif create lacp MultiTrunk1 -b ip e0 e1 e2 e3

2. To configure the dynamic multimode vif, enter the following command:

ifconfig vif_name IP_address netmask mask

Related concepts

Dynamic multimode vif on page 110 *Load balancing in multimode vifs* on page 112

Related tasks

Changing the status of an interface on page 50

Adding interfaces to a vif

You can add one or more interfaces to a vif to expand your network. You can add physical interfaces to a vif any time after you create it by using the vif add command.

Before you begin

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

Step

1. Enter the following command:

```
vif add vif_name interface_list
```

vif_name is the name of a previously configured vif.

interface_list is a list of the interfaces you want to add to the vif.

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

Example

To add the interface e4 to the multimode vif MultiTrunk1, enter with the following command:

```
vif add MultiTrunk1 e4
```

Related tasks

Changing the status of an interface on page 50

Deleting interfaces from a vif

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

Before you begin

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

ifconfig vif_name down

vif_name is the name of the vif that you want to configure to the down status.

About this task

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

Step

1. Enter the following command:

vif delete vif_name interface

vif_name is the name of the vif.

interface is the interface of the vif you want to delete.

Example

To delete the interface e4 from a multimode vif MultiTrunk1, enter the following commands:

ifconfig MultiTrunk1 down

vif delete MultiTrunk1 e4

Related tasks

Changing the status of an interface on page 50

Viewing vif status

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

Step

1. Enter the following command:

```
vif status [vif_name]
```

vif_name is the name of the vif whose status you want to display.

If you do not specify the vif name, the status of all vifs is displayed.

Example

The following example displays the status of the vif vif1:

```
mediatype: auto-1000t-fd-up
        flags: enabled
        input packets 324193, input bytes 468036576
        output packets 161472, output bytes 13983580
        up indications 2, broken indications 0
       drops (if) 0, drops (link) 0
        strike count: 0 of 10
        indication: up at 30Jan2009 14:23:24
                consecutive 604, transitions 5
eOc: state up, since 30Jan2009 14:23:24 (00:04:40)
       mediatype: auto-1000t-fd-up
        flags: enabled
        input packets 526276, input bytes 762227558
        output packets 262355, output bytes 22321102
        up indications 2, broken indications 0
        drops (if) 0, drops (link) 0
        strike count: 0 of 10
        indication: up at 30Jan2009 14:23:24
                consecutive 606, transitions 5
```

What the vif status information table contains

You can view the status information of a vif by using the vif status command.

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

Field	Subfield	Description
default		Indicates the default values for fields such as transmit, VIF Type, and fail. These values apply if you do not specify any values for these fields when creating a vif.
	transmit	Indicates the default load-balancing method.
	VIF Type	Indicates the default vif type.
	fail	Indicates the default location where the errors are logged.
vif_name		Indicates that the data that follows this field pertains to the vif, vif_name.
	transmit	Indicates the load-balancing method used.
	VIF Type	Indicates the type of vif. Possible values are single-mode, multi_mode, or lacp.
	fail	Indicates the location where errors are logged for the vif.
	VIF Status	Indicates the current status of the vif, vif_name.
	Addr_set	Indicates that a MAC address has been configured for the vif, vif_name, and all its interfaces.

Field	Subfield	Description
	state	Indicates the current link-state of the interface. Possible values are up or down.
	since	Indicates the date, time, and number of hours since the interface has been up.
	mediatype	Indicates the media type that defines the speed and duplex for that interface.
	flags	Indicates whether the interface is enabled or disabled to send and receive data.
	strike count	Indicates the number of attempts for link-monitoring. When an underlying link of a vif does not receive any packets (including ARP packets that are used for link-monitoring), the strike count gets incremented once in 5 seconds. If this strike count reaches 10, the underlying link is brought "down."
	consecutive	Indicates the number of consecutively received "up" or "broken" indications from the switch and link interaction.
	transitions	Indicates the number of indications received that caused a state transition from "up" to "broken" or "down" to "up".

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

Viewing vif statistics

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

Step

1. Enter the following command:

```
vif stat [vif_name] [interval]
```

vif_name is the name of the vif. If you do not specify a vif, the status of all vifs is displayed.

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

Example

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

```
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. The following rows show the total number of packets received and sent per second thereafter.

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

```
vif stat vif1
Virtual interface (trunk) vif1
e0c e0b
Pkts In Pkts Out Pkts In Pkts Out
82 208k 796k 208k
1 27342 104774 27326
2 26522 102088 26560
8 20332 77275 20335
5 27198 103529 27186
```

Destroying 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 destroy the vif, the interfaces in the vif act individually rather than as an aggregate.

Steps

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

ifconfig vif_name down

vif_name is the name of the vif you want to configure to the down status.

2. Enter the following command:

vif destroy vif_name

vif_name is the name of the vif you want to destroy.

Second-level vifs

If you have more than one multimode vif, you can use the vif create command to group them by creating 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 an active/active configuration.

Note: You cannot use LACP vifs as second-level vifs.

Next topics

Guidelines for creating a second-level vif on page 126 *Creating a second-level vif* on page 126 *Enabling failover in a second-level vif* on page 127

Guidelines for creating a second-level vif

You can create a single-mode second-level vif over two multimode vifs. The ports of the underlying multimode vifs should be connected to the same switch. If you create a second-level vif over two multimode vifs that are connected to two different switches, you should connect the two switches with an inter-switch link (ISL).

For a single-mode vif, the switch ports must be in the same broadcast domain (for example, a LAN or a VLAN). Link-monitoring ARP packets are sent over the ports of a single-mode vif to detect whether the ports are in the same broadcast domain. If the ports are not in the same broadcast domain, the vif is configured to the down status.

When the ports of a single-mode vif are connected to different broadcast domains, it is called a *split-network condition*. Therefore, a second-level vif, created over two multimode vifs that are connected to two different switches without an ISL, is automatically configured to the down status.

Creating a second-level vif

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

Before you begin

You must 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 name for the second-level vif:
 - 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 to select as part of the vif.
- Configure all interfaces that are to be included in the vif to the down status, by using the ifconfig command.

Steps

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

```
vif create multi -b {rr|mac|ip|port} vif_name1 if1 if2
```

The *vif_name1* vif is composed of two physical interfaces, *if1* and *if2*.

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

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

mac-specifies the MAC address load-balancing option.

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

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

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

```
vif create multi -b {rr|mac|ip|port} vif_name2 if3 if4
```

The vif_name2 vif is composed of two physical interfaces, if3 and if4.

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

vif create single vif_name vif_name1 vif_name2

vif_name is the second-level vif created with two multimode vifs, vif_name1 and vif_name2.

Example

Use the following commands to create two vifs and a second-level vif. In this example, IP address 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

Related tasks

Changing the status of an interface on page 50

Enabling failover in a second-level vif

In a second-level single-mode vif over two or more multimode vifs, you can enable the vif.failover.link_degraded option for failing over to a multimode vif with higher aggregate

bandwidth. The failover happens regardless of whether the currently active vif is favored or not. By default, this option is off.

Step

1. To enable failover to a multimode vif with higher aggregate bandwidth when one or more of the links in the active multimode vif fail, enter the following command:

```
options vif.failover.link_degraded on
```

Second-level vifs in an active/active configuration

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

With 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. Therefore, by using the two configurations together, you can achieve a fully redundant storage system connectivity architecture.





When both storage systems are in operation, the following connections exist:

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

Firstlev2 and Firstlev3 are in standby mode.

If one of the switches fails, the following happens:

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

In the following figure, Switch 1 fails in an active/active configuration. Firstlev2 takes over the MAC address of Firstlev1 and maintains the connectivity through Switch 2.



Creating a second-level vif in an active/active configuration

You can create two second-level vifs in an active/active configuration so that you can access data from both storage systems even if one of the storage system in the configuration fails.

Before you begin

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

About this task

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

Assume *StorageSystem1* and *StorageSystem2* are the storage systems that are configured in an active/active configuration.

Steps

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

```
vif create multi -b {rr|mac|ip|port} vif_name1 if1 if2
vif create multi -b {rr|mac|ip|port} vif_name2 if3 if4
```

-b specifies the type of load-balancing method.

rr specifies the round-robin load-balancing option.

mac specifies the MAC address load-balancing option.

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

port specifies the port-based load-balancing option.

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

vif_name1 and vif_name2 are the names of the multimode vifs.

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

vif create single secondlev1 vif_name1 vif_name2

secondlev1 is the name of the second-level vif.

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

```
vif create multi -b {rr|mac|ip|port} vif_name3 if5 if6
```

vif create multi -b {rr|mac|ip|port} vif_name4 if7 if8

4. Enter the following command on *StorageSystem2* to create a second-level interface from the multimode vifs:

vif create single secondlev2 vif_name3 vif_name4

5. Enter the following command on *StorageSystem1* to configure the second-level vifs for takeover:

ifconfig secondlev1 partner secondlev2

6. Enter the following command on *StorageSystem2* to configure the second-level vifs for takeover:

ifconfig secondlev2 partner secondlev1

In steps 5 and 6, secondlev1 and secondlev2 (arguments to the partner option) must be interface names and not interface IP addresses. If secondlev1 is a vif, secondlev2 must also be a vif.

Example

Use the following commands to create a second-level vif in an active/active configuration. In this example, IP-based load balancing is used for the multimode vifs.

On StorageSystem1:

vif create multi Firstlev1 e1 e2

```
vif create multi Firstlev2 e3 e4
vif create single Secondlev1 Firstlev1 Firstlev2
On StorageSystem2:
vif create multi Firstlev3 e5 e6
vif create multi Firstlev4 e7 e8
vif create single Secondlev2 Firstlev3 Firstlev4
On StorageSystem1:
ifconfig Secondlev1 partner Secondlev2
On StorageSystem2:
ifconfig Secondlev2 partner Secondlev1
```

Related tasks

Changing the status of an interface on page 50

How CDP works with Data ONTAP

In a data center, you can use Cisco Discovery Protocol (CDP) to view network connectivity between a pair of physical or virtual systems and their network interfaces. CDP is also useful for verifying network connectivity before performing online migration of vFiler units.

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

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

Next topics

Data ONTAP support for CDP on page 133 Enabling or disabling CDP on your storage system on page 134 Configuring hold time for CDP messages on page 134 Setting the intervals for sending CDP advertisements on page 135 Viewing or clearing CDP statistics on page 135 Viewing neighbor information by using CDP on page 137

Data ONTAP support for CDP

By default, Cisco devices or CDP-compliant devices send CDPv2 advertisements. CDP-compliant devices send CDPv1 advertisements only when they receive CDPv1 advertisements. Data ONTAP supports only CDPv1. Therefore, when the storage controller sends CDPv1 advertisements, the immediately connected CDP-compliant devices send back CDPv1 advertisements.

You should consider the following information before enabling CDP on your storage system:

- CDP advertisements are sent only by the ports that are in the up state and configured with IP addresses.
- CDP must be enabled on both the transmitting and receiving devices for sending and receiving CDP advertisements.
- CDP advertisements are sent at regular intervals. You can configure the time interval.
- When IP addresses are changed at the storage system side, the storage system sends the updated information in the next CDP advertisement.

Note: Sometimes when IP addresses are changed at the storage system side, the CDP information is not updated at the receiving device side (for example, a switch). If you

encounter such problem, you should configure the network interface of the storage system to the down status and then to the up status.

- Only IPv4 addresses are advertised in CDP advertisements.
- For physical network ports with VLANs, all the IP addresses configured on the VLANs on that port are advertised.
- For physical ports that are part of a vif, all the IP addresses configured on that vif are advertised on each physical port.
- For a vif that hosts VLANs, all the IP addresses configured on the vif and the VLANs are advertised on each of the network ports.
- The number of IP addresses that can fit into a 1500 MTU sized packet is advertised for packets with MTU sizes equal and greater than 1500 bytes.

Enabling or disabling CDP on your storage system

To discover and send advertisements to CDP-compliant neighboring devices, CDP must be enabled on the storage system. You should use the cdpd.enable option to enable or disable CDP on your storage system.

About this task

When the cdpd.enable option is set to on, CDPv1 is enabled on all physical ports of the storage system.

Step

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

```
options cdpd.enable {on|off}
```

on-Enables CDP

```
off—Disables CDP
```

Configuring hold time for CDP messages

Hold time is the period of time for which all CDP advertisements are stored in a cache in the neighboring CDP-compliant devices. Hold time is advertised by the storage controller in each CDPv1 packet. You can use the cdpd.holdtime option to configure hold time.

About this task

The value of the cdpd.holdtime option applies to both members of an active/active configuration.

The default value of hold time is 180 seconds.

Step

1. To configure the hold time, enter the following command:

options cdpd.holdtime holdtime

holdtime is the time interval, in seconds, for which the CDP advertisements are cached in the neighboring CDP-compliant devices. You can enter values ranging from 10 seconds to 255 seconds.

Setting the intervals for sending CDP advertisements

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

About this task

The value of the cdpd.interval option applies to both the members of an active/active configuration.

Step

1. To configure the interval for sending CDP advertisements, enter the following command:

```
options cdpd.interval interval
```

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

Viewing or clearing CDP statistics

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

Step

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

If you want to	Enter the following command:
View the CDP statistics	cdpd show-stats
Clear the CDP statistics	cdpd zero-stats

Example of showing the statistics before and after clearing them

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

```
system1> cdpd show-stats
```

```
RECEIVE

Packets: 9116 | Csum Errors: 0 | Unsupported

Vers: 4561

Invalid length: 0 | Malformed: 0 | Mem alloc

fails: 0

Missing TLVs: 0 | Cache overflow: 0 | Other

errors: 0

TRANSMIT

Packets: 4557 | Xmit fails: 0 | No

hostname: 0

Packet truncated: 0 | Mem alloc fails: 0 | Other

errors: 0
```

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

Enter the following command to clear the statistics:

```
cdpd zero-stats
```

The following output shows the statistics after they are cleared:

```
system1> cdpd show-stats
```

```
RECEIVE

Packets: 0 | Csum Errors: 0 | Unsupported

Vers: 0

Invalid length: 0 | Malformed: 0 | Mem alloc

fails: 0

Missing TLVs: 0 | Cache overflow: 0 | Other

errors: 0

TRANSMIT

Packets: 0 | Xmit fails: 0 | No

hostname: 0

Packet truncated: 0 | Mem alloc fails: 0 | Other

errors: 0
```

```
OTHER
Init failures: 0
```

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

Viewing neighbor information by using CDP

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

Before you begin

CDP must be enabled on your storage system.

Step

1. To view information about all CDP-compliant devices connected to your storage system, enter the following command:

cdpd show-neighbors

Example

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

```
system1> cdpd show-neighbors
Local Remote Remote
                                                        Hold
                                         Remote
Remote
Port Device
                 Interface
                                         Platform
                                                        Time
Capability
 ----
e0a sw-215-cr(4C2) GigabitEthernet1/17 cisco WS-C4948
                                                         125
RSI
e0b sw-215-11(4C5) GigabitEthernet1/15 cisco WS-C4948
                                                         145
SI
      sw-215-11(4C5) GigabitEthernet1/16 cisco WS-C4948
                                                         145
e0c
SI
```

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

- R—Router
- T—Transparent bridge
- B—Source-route bridge
- S—Switch

138 | Data ONTAP 7.3 Network Management Guide

- H—Host
- I—IGMP
- r—Repeater
- P—Phone

How to monitor your storage system with SNMP

If you enable SNMP in Data ONTAP, the SNMP managers can query your storage system's SNMP agent for information. The SNMP agent gathers information and forwards it to the managers by using SNMP. The SNMP agent also generates trap notifications whenever specific events occur.

For diagnostic and other network management services, Data ONTAP provides an SNMP agent compatible with SNMP versions 1 and 3. SNMPv3 offers advanced security by using pass phrases and encryption. SNMPv3 supports the MIB-II specification and the MIBs of your storage system. The following MIB-II groups are supported:

- System
- Interfaces
- Address translation
- IP
- ICMP
- TCP
- UDP
- SNMP

Note: Transmission and EGP MIB-II groups are not supported.

Starting with Data ONTAP 7.3.1, IPv6 (RFC 2465), TCP (RFC 4022), UDP (RFC 4113), and ICMP (RFC 2466) MIBs, which show both IPv4 and IPv6 data, are supported.

Next topics

Types of SNMP traps in Data ONTAP on page 139 *What a MIB is* on page 140 *What the SNMP agent does* on page 140 *How to configure the SNMP agent* on page 140 *User-defined SNMP traps* on page 148

Types of SNMP traps in Data ONTAP

SNMP traps capture system monitoring information in Data ONTAP. There are two types of traps in Data ONTAP: built-in and user-defined.

- Built-in traps are predefined in Data ONTAP and are automatically sent to the network management stations on the traphost list if an event occurs. These traps are based on one of the following:
 - RFC 1213, which defines traps such as coldStart, linkDown, linkUp, and authenticationFailure.

- Specific traps defined in the custom MIB, such as diskFailedShutdown, cpuTooBusy, and volumeNearlyFull.
- User-defined traps are defined by snmp traps commands or the FilerView SNMP Traps windows. These traps are sent using proxy trap ID numbers 11 through 18, which correspond to a trap's MIB priority.

What a MIB is

A MIB file is a text file that describes SNMP objects and traps. MIBs are not configuration files. Data ONTAP does not read these files and their contents do not affect SNMP functionality.

Data ONTAP provides two MIB files:

- A custom MIB (/etc/mib/netapp.mib)
- An Internet SCSI (iSCSI) MIB (/etc/mib/iscsi.mib)

Data ONTAP also provides a short cross-reference between object identifiers (OIDs) and object short names in the /etc/mib/traps.dat file. This file is useful for creating user-defined traps.

Note: The latest versions of the Data ONTAP MIBs and traps.dat files are available online on the IBM NAS support site. However, the versions of these files on the Web site do not necessarily correspond to the SNMP capabilities of your Data ONTAP version. These files are provided to help you evaluate SNMP features in the latest Data ONTAP version.

What the SNMP agent does

The storage system includes an SNMP agent that responds to queries and sends traps to network management stations.

The SNMP agent on the storage system has only read privileges—that is, it cannot be used to take corrective action in response to a trap.

Note: Starting with Data ONTAP 7.3.1, the SNMP agent supports IPv6 transport.

How to configure the SNMP agent

You need to configure the SNMP agent on your storage system to set SNMP values and parameters. You can configure your SNMP agent through the command-line interface or with FilerView.

To configure the SNMP agent on your storage system, you need to perform the following tasks:

• Verify that SNMP is enabled.

Note: SNMP is enabled by default in Data ONTAP.

- If you are running SNMPv3, configure SNMPv3 for read-only access.
- Enable traps. Although SNMP is enabled by default, traps are disabled by default.
- Specify host names of one or more network management stations.

Traps can only be sent when at least one SNMP management station is specified as a traphost. Trap notifications can be sent to a maximum of eight network management stations.

Note: The SNMP agent can send traps over IPv6 transport to the traphosts whose IPv6 address is configured on the storage system. You can specify traphosts by their IPv6 addresses, but not by their host names.

You can perform the following tasks after configuring SNMP:

- Provide courtesy information about storage system location and contact personnel.
- Set SNMP access privileges. You can restrict SNMP access on a host or interface basis.
- Specify SNMP communities.

Community strings function as group names to establish trust between SNMP managers and clients. Data ONTAP imposes the following limitations on SNMP communities:

- No more than eight communities are allowed.
- Only read-only communities are supported.
- Enable query authentication.

You can enable authentication failure traps, which are generated when the agent receives queries with the wrong community string, for the SNMP agent. The traps are sent to all hosts specified as traphosts.

• Create and load user-defined traps.

Note: Storage systems in an active/active configuration can have different SNMP configurations. For more information, see the na_snmp(1) man page.

Next topics

Enabling or disabling SNMP using the command-line interface on page 142

Configuring SNMPv3 users on page 142

Setting SNMP access privileges on page 143

Viewing or modifying your SNMP configuration from the command-line interface on page 143 *Modifying your SNMP configuration from FilerView* on page 144

SNMP command syntax on page 144

SNMP security parameters on page 145

Example: SNMP commands on page 146

Related concepts

User-defined SNMP traps on page 148

Enabling or disabling SNMP using the command-line interface

You can enable or disable SNMP from the command-line interface by entering the options snmp.enable command.

Step

1. Enter the following command:

options snmp.enable {on|off}

on-Enables SNMP

off—Disables SNMP

Configuring SNMPv3 users

To access MIB objects by using SNMPv3, you should create users with login-snmp capability.

Steps

1. Enter the following command to create a role with login-snmp capability:

useradmin role add role_name -a login-snmp

role_name is the role name with login-snmp capability.

Example

useradmin role add myrole1 -a login-snmp

2. Enter the following command to create a group and add the created role to that group:

useradmin group add group_name -r role_name

group_name is the group name to which you want to add the created role, role_name.

Example

useradmin group add mygroup1 -r myrole1

3. Enter the following command to create a user and add the user to the created group:

useradmin user add user_name -g group_name

user_name is the user name belonging to the group, group_name.

Example

useradmin user add myuser1 -g mygroup1

4. Create a password for the new user.

Ensure that the password has a minimum of eight characters.

5. Enter the snmpwalk command through the system MIB:

snmpwalk -v 3 -u user_name -l authNoPriv -A password storage_system
system

password is the user's password that you entered in Step 3.

storage_system is the storage system that contains the MIBs.

Example

snmpwalk -v 3 -u myuser1 -l authNoPriv -A johndoel23 host1 system

Setting SNMP access privileges

You can set SNMP access privileges on a host or an interface by using the command-line interface. The snmp.access option defines a method to restrict access to the storage system on a protocol-by-protocol basis.

About this task

You cannot set access privileges from FilerView.

Step

1. Enter the following command:

options snmp.access access_spec

access_spec consists of keywords and their values. Access can be allowed or restricted by host name, IP address, and network interface name.

Example

To allow access to SNMP for network interfaces e0, e1, and e2, enter the following command:

options snmp.access if=e0,e1,e2

Viewing or modifying your SNMP configuration from the command-line interface

You can use the snmp command to view or modify your SNMP configuration values.

Step

1. Enter the following command:

snmp {options values}

options are the available options for the snmp command, such as authtrap, community, contact, init, location, traphost, and traps.

values are the values that you want to set for a particular option.

Related references

SNMP command syntax on page 144

Modifying your SNMP configuration from FilerView

You can use FilerView to modify your SNMP configuration.

Steps

1. From the list on the left pane, click **SNMP** > **Configure**.

The current SNMP configuration is displayed.

- **2.** To set or modify SNMP configuration values, enter configuration values in the drop-down lists or text fields.
- 3. Click Apply.

SNMP command syntax

If you specify one or more values for an option of the SNMP commands, the value of that option is set or changed. However, if no values are specified, the current value of that option is returned.

Command	Description
snmp	Displays the current values of all SNMP options, such as init, community, contact, and traphost.
snmp authtrap [0 1]	With a value: Enables (with value 1) or disables (with value 0) authentication failure traps on the SNMP agent.
	Without a value: Displays the current value of authtrap set in Data ONTAP.
snmp community	Displays the current list of communities.
snmp community add ro <i>community</i>	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 ro (read-only).
<pre>snmp community delete {all rocommunity }</pre>	Deletes one or all communities.

The following table describes the syntax and parameters of SNMP commands.
Command	Description
<pre>snmp contact [contact]</pre>	With a value: Sets the contact name for your storage system. You must enclose the contact string in single quotes (' ') if the string contains spaces.
	You can enter a maximum of 255 characters for the contact information.
	Without a value: Displays the current contact name set in Data ONTAP.
snmp init [0 1]	With a value: Enables (with value 1) or disables (with value 0) built-in traps and the traps defined using the snmp traps command.
	Without a value: Displays the current value of snmp init in Data ONTAP.
	Default value : By default, SNMP traps are disabled in Data ONTAP; the system uses the equivalent of snmp init 0.
<pre>snmp location [location]</pre>	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.
<pre>snmp traphost [{add delete} { hostname </pre>	With the option: Adds or deletes SNMP hosts that receive traps from Data ONTAP.
1padaress}]	When IPv6 is enabled on your storage system, IPv6 traphosts can be added and deleted. You can specify IPv6 addresses, and not host names, to identify IPv6 traphosts.
	Without the option: Displays the current traphosts set in Data ONTAP.
snmp traps[options]	Displays the list of user-defined traps set in Data ONTAP

SNMP security parameters

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

If the security level is set to authNoPriv, authentication is performed by using a user's authKey to sign the message being sent. The authProtocol parameter must be MD5. The authKey parameters are generated from a passphrase that must have a minimum of eight characters.

If the security level is set to authNoPriv, you must enter the following parameters:

Parameter	Command-line flag	Description	
securityName	-u Name	User name must not exceed 31 characters.	

Parameter	Command-line flag	Description		
authProtocol	-a (MD5)	Authentication type must be MD5.		
authKey	-A PASSPHRASE	Passphrase with a minimum of eight characters.		
securityLevel	-l (authNoPriv)	Security level: must be Authentication, No Privacy. Note: Data ONTAP does not support retrieving MIB values using the noAuthNoPriv security level.		
context	-n CONTEXTNAME	Sets the context name used for SNMPv3 messages.		

Example: SNMP commands

You can use the snmpget, snmpwalk, snmpbulkget, and snmpbulkwalk commands to retrieve information from network elements with SNMP agents.

snmpwalk

The following command retrieves all the variables under the system sys1:

```
snmpwalk -Os -c public -v 1 sys1 system
sysDescr.0 = STRING: Data ONTAP Release 7.3.1
sysObjectID.0 = OID: enterprises.789.2.3
sysUpTimeInstance = Timeticks: (121596665) 14 days, 1:46:06.65
sysContact.0 = STRING:
sysName.0 = STRING: sys1.lab.example.com
sysLocation.0 = STRING:
sysServices.0 = INTEGER: 72
```

The following command is an example of an SNMP request from an IPv6 client:

```
snmpwalk -v2c -c public udp6:[2001:0db8:85a3:0:0:8a2e:0370:99]:161 system
SNMPv2-MIB::sysDescr.0 = STRING: Data ONTAP Release 7.3.1
SNMPv2-MIB::sysObjectID.0 = OID:
SNMPv2-SMI::enterprises.789.2.3
DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (11415057) 1 day,7:42:30.57
SNMPv2-MIB::sysContact.0 = STRING:
SNMPv2-MIB::sysName.0 = STRING:
SNMPv2-MIB::sysLocation.0 = STRING:
SNMPv2-MIB::sysLocation.0 = STRING:
SNMPv2-MIB::sysServices.0 = INTEGER: 72
```

The following command is an example of an SNMPv3 request to retrieve all variables under the system sys1:

```
snmpwalk -v 3 -u joeblow -l authNoPriv -A joeblow12 sys1 system
SNMPv2-MIB::sysDescr.0 = STRING: Data ONTAP Release 7.3.1
SNMPv2-MIB::sysObjectID.0 = OID:
SNMPv2-SMI::enterprises.789.2.3
DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (121622059) 14
days, 1:50:20.59
SNMPv2-MIB::sysContact.0 = STRING:
```

SNMPv2-MIB::sysName.0 = STRING: sys1.lab.example.com SNMPv2-MIB::sysLocation.0 = STRING: SNMPv2-MIB::sysServices.0 = INTEGER: 72

Note: You need to enter authentication information for using SNMPv3.

snmpget

The following command retrieves the system.sysDescr.0 object from the host sys1 by using the public community string:

```
snmpget -c public sys1 system.sysDescr.0
system.sysDescr.0 = Data ONTAP Release 7.3.1 Mon Mar 16 16:56:43 IST 2009
```

The following command retrieves the value of an ICMP object (OID=56.1.1.1) from the host sys1:

```
snmpget -c public -v 2c sys1 .1.3.6.1.2.1.56.1.1.1.1
56.1.1.1.1.1 = Counter32: 0
```

snmpbulkget

The following command retrieves the system object sysDescr.0 and the first three objects in the ifTable:

```
snmpbulkget -v2c -Cn1 -Cr3 -Os -c public sys1 system ifTable
sysDescr.0 = STRING: Data ONTAP Release 7.3.1
ifIndex.1 = INTEGER: 1
ifIndex.2 = INTEGER: 2
ifDescr.1 = STRING: "lo0"
```

The following example shows a part of the output from retrieving all variables under the IPv6 object (OID=55.1):

```
snmpbulkget -c public -v 2c 192.0.2.19 .1.3.6.1.2.1.55.1
55.1.1.0 = 2
55.1.2.0 = 64
55.1.3.0 = Gauge32: 4
55.1.4.0 = Counter32: 0
55.1.5.1.1.1 = 1
55.1.5.1.2.1 = "ns0"
55.1.5.1.3.1 = OID: .ccitt.zeroDotZero
55.1.5.1.4.1 = 1500
55.1.5.1.6.1 = IpAddress: 00 00 00 00 00 00 00 02 05 00 FF FE 00 02 AB
55.1.5.1.8.1 = Hex: 00 05 00 00 02 AB
55.1.5.1.9.1 = 1
55.1.5.1.9.1 = 1
```

snmpbulkwalk

The following command retrieves all the variables under the system sys1:

```
snmpbulkwalk -v2c -Os -c public sys1 system
sysDescr.0 = STRING: Data ONTAP Release 7.3.1
sysObjectID.0 = OID: enterprises.789.2.3
sysUpTimeInstance = Timeticks: (121603434) 14 days, 1:47:14.34
sysContact.0 = STRING:
```

```
sysName.0 = STRING: sys1.lab.example.com
sysLocation.0 = STRING:
sysServices.0 = INTEGER: 72
```

The following example shows a part of the output from retrieving all the variables for the UDP object:

```
snmpbulkwalk -c public -v 2c 192.0.2.19 udp
udp.udpInDatagrams.0 = Counter32: 347
udp.udpNoPorts.0 = Counter32: 4
udp.udpInErrors.0 = Counter32: 0
udp.udpOutDatagrams.0 = Counter32: 138
udp.udpTable.udpEntry.udpLocalAddress.0.0.0.0.69 = IpAddress: 00 00 00 00
udp.udpTable.udpEntry.udpLocalAddress.0.0.0.0.111 = IpAddress: 00 00 00 00
```

User-defined SNMP traps

If the predefined built-in traps are not sufficient to create alerts for conditions you want to monitor, you can create user-defined traps in Data ONTAP.

Before you define a new trap, you should consult the Data ONTAP MIBs to see if any existing traps serve your purpose.

Next topics

How SNMP traps work on page 148 How to define or modify a trap on page 149 Viewing or modifying trap values from the command-line interface on page 149 Viewing or modifying trap values from FilerView on page 149 Defining traps in a configuration file on page 150 Example: Trap definitions on page 151 Command syntax for SNMP trap parameters on page 151 SNMP trap parameters on page 152

How SNMP traps work

You can set SNMP traps to inspect the value of MIB variables periodically. Whenever the value of a MIB variable meets the conditions you specify, a trap is sent to the network management stations on the traphost list. The traphost list specifies the network management stations that receive the trap information.

You can set traps on any numeric variable in the MIB. For example, you can set a trap to monitor the fans on your storage system 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.

Follow these guidelines when creating traps:

- Use the /etc/mib/traps.dat file to find Object Identifiers (OIDs) for objects in the MIB files of your storage system.
- Ensure that the trap can be generated in the storage system environment.
- Do not set traps on tabular data.

It is possible to set traps on row entries in a sequence—for example, an entry in a table. However, if the order in the table is changed by adding or removing rows, you will no longer be trapping the same numeric variables.

How to define or modify a trap

You can define traps or modify traps you have already defined by entering values from the command-line interface, in FilerView, or in a configuration file.

You must supply the following elements when you create or modify traps:

• Trap name

Trap name is the name of the user-defined trap you want to create or change. A trap name must not have any embedded periods.

- Trap parameters
- Parameter values

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

Viewing or modifying trap values from the command-line interface

You can view or modify your trap values by using the snmp traps command.

Step

1. Enter the following command:

```
snmp traps {options variables}
```

options are the options for SNMP traps such as walk, load, trapname, and so on.

variables are values for the specified option.

Viewing or modifying trap values from FilerView

You can use FilerView to view or modify a trap value.

Steps

- 1. From the list on the left pane, click **SNMP > Traps**.
- 2. Depending on whether you want to create, modify, or view a trap, perform the following step:

If you want to	Then			
Create a new trap	 a. Click Add. b. In the Add an SNMP Trap window, enter the requested information. c. Click Add again. 			
View or modify an existing trap	 a. Click Manage for the trap you want to view or modify. b. To modify the trap, click Modify in the Manage SNMP Traps window. 			

Defining traps in a configuration file

You can define SNMP traps in a configuration file and then load the file with the snmp traps load command. Data ONTAP automatically backs up your SNMP configuration as Snapshot copies, making it easy to transfer user-defined traps to other storage systems and to recover SNMP configurations in case of a disaster.

Steps

1. Create a traps configuration file on your storage system.

You can choose the name and location of the file.

Example

/etc/mib/mytraps

2. Enter the traps in the configuration file in the following format:

trapname.parmvalue

The parameters are the same as those used with the snmp traps command.

- **3.** Test each line of the file by entering the snmp traps command using the command-line interface or by specifying the trap using FilerView. Make corrections as needed.
- 4. Load the configuration file with the following command:

snmp traps load file_name

file_name is the name of the configuration file.

Example

snmp traps load /etc/mib/mytraps

Related references

SNMP trap parameters on page 152

Example: Trap definitions

You can define a group of traps by using the command-line interface or FilerView.

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

```
snmp traps cifstotalops.var snmp.1.3.6.1.4.1.789.1.7.3.1.1.1.0
[1]
snmp traps cifstotalops.trigger level-trigger
snmp traps cifstotalops.edge-1 1000000
[4]
snmp traps cifstotalops.interval 10
[2]
snmp traps cifstotalops.backoff-calculator step-backoff
[5]
snmp traps cifstotalops.backoff-step 3590
[5]
snmp traps cifstotalops.rate-interval 3600
[3]
snmp traps cifstotalops.priority alert
snmp traps cifstotalops.message snmp.1.3.6.1.4.1.789.1.7.3.1.1.1.0
```

A cifstotalops trap [1] is evaluated every 10 seconds [2]. The value received from the previous evaluation and the current value are used to calculate the number of CIFS operations per hour [3]. If the number exceeds one million [4], the trap triggers and continues to be triggered every hour [5] until the total number of CIFS operations drops below one million.

Command syntax for SNMP trap parameters

If you specify one or more values for an option of the SNMP commands, the value of that option is set or changed. However, if no values are specified, the current value of that option is returned.

Command	Description
snmp traps	Displays the list of user-defined traps set in Data ONTAP.
snmp traps [enable disable reset delete] <i>trapname</i>	Enables, disables, resets, or deletes the trap <i>trapname</i> . If you do not specify a trap name, all traps defined so far are acted on.
snmp traps walk prefix	Walks (traverses in order) the trap list by prefix; that is, lists all traps that have names beginning with <i>prefix</i> .

The following table describes the syntax and parameters for the snmp traps command.

Command	Description
snmp traps load <i>trap_list_filenam</i> e	Loads a set of traps from a configuration file. The file contains a list of traps and parameters without the snmp traps command preceding each trap. If the specified file name is defaults, traps are read from the /etc/ defaults/ traps file.
snmp traps trapname.parm value	Defines or changes a user-defined trap parameter.

SNMP trap parameters

You must specify certain parameters to create SNMP traps.

The following table lists SNMP trap parameters that you enter with the snmp traps command in the command-line interface and the equivalent parameters that you select in FilerView.

Parameter in command-line interface	Equivalent in FilerView		
var	OID		
trigger	Trigger		
edge-1	Edge 1		
edge-2	Edge 2		
edge-1-direction	Edge 1 Direction		
edge-2-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		

Next topics

The var parameter on page 153 *The trigger parameter* on page 153 *The edge-1 and edge-2 parameters* on page 154 The edge-1-direction and edge-2-direction parameters on page 154 The interval parameter on page 154 The interval-offset parameter on page 154 The rate-interval parameter on page 155 The backoff-calculator parameter on page 155 The backoff-step parameter on page 155 The backoff-multiplier parameter on page 156 The priority parameter on page 156 The message parameter on page 156

The var parameter

The var parameter associates a user-defined trap name (specified by the *trapname* variable in the snmp traps command or Trap Name in FilerView) with a specific MIB object. The MIB object is specified in the value field of the snmp traps command. It must be in the format snmp.oid, where oid is an Object Identifier (OID).

The traps.dat file, which is located in the /etc/mib directory on your storage system, can help you determine OIDs. This file maps MIB objects' short names in the Data ONTAP MIB files to their numeric OIDs. For more information about a particular OID, see the MIB.

In FilerView, it is necessary to enter only the numerical OID, and not the "snmp" prefix.

The trigger parameter

The trigger parameter specifies the type of triggers that you can set for a trap. If a trap is triggered, data about the event that caused the trigger is sent to the network management stations.

You can specify the following values for the trigger parameter:

single-edge- trigger	Triggers a trap and sends data when the value of the trap's MIB variable crosses an edge (a value that you specify) for the first time.
double-edge- trigger	Triggers a trap and sends data when either of two edges is crossed. A double-edge-trigger enables you to set two edges, each with its own direction.
level-trigger	Triggers a trap and sends data whenever the trap's value crosses a specified edge value.
change- trigger	Keeps track of the last value received from the trap. If the current value differs from the previously received value, the trap is triggered.
always- trigger	Enables a trap to always trigger at the specified evaluation interval (specified by the interval parameter). For example, a trap can trigger every 24 hours for the agent to send the total number of CIFS operations to an SNMP manager.

The edge-1 and edge-2 parameters

The edge-1 and edge-2 parameters of a trap specify the threshold values that are compared during trap evaluation to determine whether to fire a trap and send data.

The edge-1 parameter specifies the value for the edge in a single-edge-triggered trap or the first edge in a double-edge-triggered trap. The default value for the edge-1 parameter is MAXINT.

The edge-2 parameter specifies the value for the second edge in a double-edge-triggered trap. The default value for the edge-2 parameter is 0.

Note: The edge-2 parameter is not displayed in FilerView during trap creation unless double-edge-trigger is selected in the trigger parameter.

The edge-1-direction and edge-2-direction parameters

The edge-1-direction and edge-2-direction parameters enable you to set or change the direction that is used to evaluate a trap. The edge-triggered traps only send data when the edge is crossed in either the up or down direction.

The default value for the edge-1-direction parameter is up and for the edge-2-direction parameter is down.

Note: You enter the direction values on the same line as the edge value when you run the snmp traps command. The edge-2-direction parameter is not displayed in FilerView during trap creation unless double-edge-trigger is selected in the trigger parameter.

The interval parameter

The interval parameter is the time, in seconds, between evaluations of a trap.

A trap can only send data as often as it is evaluated, even if the edge values are exceeded sooner. The default value for the interval parameter is 3600.

Note: The maximum value that can be specified for the interval parameter in Data ONTAP is 2147482.

The interval-offset parameter

The interval-offset parameter is the amount of time, in seconds, until the first trap evaluation.

The default value for the interval-offset parameter is 0. You can set it to a nonzero value to prevent too many traps from being evaluated at once (for example, at system startup).

The rate-interval parameter

The rate-interval parameter specifies the time, in seconds, in which the change in value of a trap's variable (rate of change) is expressed.

If the rate-interval value is set for a trap, the samples of data obtained at the interval points (set using the interval parameter) for a trap variable are used to calculate the rate of change. If the calculated value exceeds the value set for the edge-1 or edge-2 parameter, the trap is fired.

For example, to obtain the number of CIFS operations per hour, you specify a rate-interval of 3600. If rate-interval is set to 0, no sampling at interval points occurs and trap evaluation proceeds as with any other kind of trap. The default value for the rate-interval parameter is 0.

The backoff-calculator parameter

The backoff-calculator parameter enables you to change the trap evaluation interval for a trap after a trap fires.

After a trap fires and sends data, you might not want it to be evaluated so often. For instance, you might want to know within a minute of when a file system is full, but only want to be notified every hour thereafter that it is still full.

The backoff-calculator parameter can take the following values in the value variable field:

- step-backoff
- exponential-backoff
- no-backoff

The default value for the backoff-calculator parameter is no-backoff.

The backoff-step parameter

The backoff-step parameter specifies the number of seconds by which the trap evaluation interval is increased.

If a trap interval is 10 and its backoff-step is 3590, the trap is evaluated every 10 seconds until it fires the first time and sends data, and once an hour thereafter. The default value for the backoff-step parameter is 0.

Note: The backoff step parameter is not displayed in FilerView during trap creation unless "step" is selected in the Backoff Style field.

The backoff-multiplier parameter

The backoff-multiplier parameter specifies the value by which to multiply a trap's evaluation interval each time it fires.

If you set backoff-calculator to exponential-backoff and backoff-multiplier to 2, the interval doubles each time the trap fires. The default value of the backoff-multiplier parameter is 1.

Note: The backoff multiplier parameter is not displayed in FilerView during trap creation unless "exponential" is selected in the Backoff Style field.

The priority parameter

The priority parameter sets the priority of a trap. If several traps are scheduled to be triggered at the same time, you can use the priority parameter to decide which trap is serviced first.

The possible values for the priority parameter, from highest to lowest, are as follows:

- emergency
- alert
- critical
- error
- warning
- notification
- informational
- debug

The default value for the priority parameter is notification.

The message parameter

The message parameter specifies a message that goes out with a trap.

The message can be a string of text or simply the SNMP OID, in the form snmp.oid. If you specify the OID as your message, Data ONTAP sends the information that was trapped concerning the OID. If you do not specify a message parameter for a trap, when the trap fires you see a string with the numerical OID value and its priority level.

For example, the following string is sent to the network management stations for the trap cpuUpTime if the message parameter is not set:

cpuUpTime == 10562288.priority == notification

Note: If the message is a string that includes spaces, you must enclose the string in quotation marks (" ").

You cannot set the message parameter in FilerView.

Internet Protocol Security

Internet Protocol Security (IPsec) is a security protocol suite that protects data from unauthorized disclosure. Using IPsec, you can add policies that configure encryption and authentication algorithms between the storage system and the client, and negotiate a security association (SA) between two end-stations that initiate and receive secure communications.

A security association is used for secure data exchanges between the storage system and the client systems.

Note: IPsec is not supported over IPv6.

Next topics

What security associations are on page 157 What security policies include on page 158 Key exchanges on page 158 IPsec implementation in Data ONTAP on page 159 IPsec in an active/active configuration on page 160 IPsec in a vFiler unit configuration on page 160 How to set up IPsec on page 161 Configuring certificate authentication on page 161 Kerberos support on page 169 Configuring preshared keys on page 169 Enabling or disabling IPsec on page 170 Security policies and IPsec on page 170 Viewing IPsec statistics on page 173 Viewing security associations on page 175

What security associations are

A security association (SA) is an authenticated simplex (uni-directional) data connection between two end-stations.

Security associations 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 one of the following:

• Authentication Header (AH)

• Encapsulating Security Payload (ESP)

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.

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.

What security policies include

IPSec security associations are created based on information collected in security policies, which determine how security is handled in a transfer of information. Security policies include specifications such as addresses of end-stations, authentication methods, and encryption mechanism.

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 endstation. Matching security policies must be configured on your storage system and clients.

Key exchanges

Key exchanges are a vital part of establishing security associations (SA). An IPsec SA is negotiated by means of the key management protocol, Internet Key Exchange (IKE). 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.

Three key exchange mechanisms using IKE are supported between storage systems and clients: certificate authentication, Kerberos, and preshared keys.

- Certificate authentication lets an end-station prove its identity by providing a certificate that has been digitally signed by a third-party certificate authority (CA), such as Verisign or Entrust. With certificate authentication, administrators need not configure keys between all IPsec peers. Instead, administrators request and install a certificate on each peer, enabling it to dynamically authenticate all other participating peers.
- Kerberos is a network authentication system in which end-stations prove their identities 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.

IPsec implementation in Data ONTAP

The IPsec implementation in Data ONTAP conforms to the Internet Engineering Task Force (IETF) Security Architecture for the Internet Protocol (RFC 2401) and related protocols.

The IPsec implementation in Data ONTAP has some restrictions that might affect its implementation on your storage system and its clients.

The following restrictions apply to the IPsec implementation in Data ONTAP:

- 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.
- IPsec is not supported over IPv6.

The following authentication methods are supported:

• For Solaris—preshared keys authentication and certificate authentication.

- For Windows—preshared keys authentication, certificate authentication, and Kerberos authentication; however, Kerberos authentication is available only for Windows Domains, not Windows Workgroups.
- Between storage systems—preshared keys authentication and certificate authentication.

The following restrictions apply to these authentication methods:

- 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.
- For certificate authentication, Data ONTAP supports v3 certificates in accordance with RFC 3280, but it does not support Certificate Revocation Lists (CRLs).
- 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 an active/active configuration

If you are considering implementing IPsec in an active/active configuration, you need to optimize IPsec to function in this environment.

The IPsec protocol, by its nature, does not work well in a failover environment (an environment in which one storage system in an active/active 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 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 re-sent.

For this reason, you should reduce the security association lifetime to a minimum value to optimize IPsec operation in an active/active 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 should 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 by using the vfiler run command.

Note: The IPsec policies and configurations must be set individually for each vFiler unit.

How to set up IPsec

You need to perform several steps to set up IPsec. These steps involve key exchanges, IPsec functionality, and security policies.

1. Select and configure one of the following key exchange mechanisms:

- Certificate authentication
- Kerberos
- · Preshared keys
- 2. Enable IPsec functionality on your storage system.
- 3. Create security policies.
- 4. View security associations.

Related concepts

Kerberos support on page 169 *Security policies and IPsec* on page 170

Related tasks

Configuring certificate authentication on page 161 *Configuring preshared keys* on page 169 *Enabling or disabling IPsec* on page 170

Configuring certificate authentication

To configure certificate authentication, you need to complete a number of steps on each storage system and Windows client that will be participating in IPsec communications.

Steps

1. Request a signed certificate from a certificate authority.

You can request a signed certificate from a Windows 2000 Server certificate authority or from a non-Windows 2000 certificate authority.

2. Install the signed certificate.

The proper installation method depends on whether the certificate was signed by a certificate authority and whether you are installing the certificate on a storage system or a Windows client.

3. Download and install one or more root certificates.

The storage system or Windows client can establish an IPsec connection with any other storage system or Windows client that uses a certificate signed by a certificate authority that you trust. To specify that you trust a specific certificate authority, you should install that certificate authority's root certificate. Then, you can optionally specify a subset of 1 to 15 certificates that Data ONTAP should use for certificate authentication.

4. Enable the IPsec certificate authentication mechanism.

Next topics

Requesting a signed certificate from a Windows 2000 certificate authority on page 162 Installing a certificate signed by a Windows 2000 certificate authority on a Windows client on page 163 Requesting a signed certificate from a non-Windows 2000 certificate authority on page 164 Installing a certificate signed by a non-Windows 2000 certificate authority on a Windows client on page 165 Installing a signed certificate on a storage system on page 166 Installing root certificates on a storage system on page 166 Specifying the subset of root certificates that Data ONTAP uses for certificate authentication on page 167 Viewing the subset of root certificates that Data ONTAP uses for certificate authentication on page 167 Installing root certificates on a Windows client on page 167 Enabling the IPsec certificate authentication mechanism on a storage system on page 168 Enabling the IPsec certificate authentication mechanism on a Windows client on page 168

Requesting a signed certificate from a Windows 2000 certificate authority

You can request a signed certificate from a Windows 2000 certificate authority.

Steps

1. Navigate to the Windows 2000 certificate authority in your Web browser.

The URL is: http://host/certsrv

host is the IP address or fully-qualified host name of the Windows 2000 Server hosting the certification authority.

- 2. Select Advanced request and click Next.
- 3. Select Submit a certificate request to this CA using a form and click Next.
- **4.** Under identifying information, type your name, e-mail address, company name, department name, state (as a two-letter abbreviation), and country (as a two-letter code).

Note: All symbols, such as ampersand (&) or at (@) symbols, should be spelled out in or omitted from the company and department names.

- 5. Under Intended Purpose, select Server Authentication Certificate.
- 6. In the Key size box, type 1024.

7. Select Mark keys as exportable.

Note: If you do not complete this step, you will not be able to export the certificate and private key into separate files, a step that is required during installation.

8. Click Submit.

After the certificate authority notifies you that your certificate has been issued, you can install the certificate.

Installing a certificate signed by a Windows 2000 certificate authority on a Windows client

If you have requested a certificate signed by a Windows 2000 certificate authority, you must then install it on your Windows client.

Before you begin

You must request for a signed certificate and receive a notification from the Windows 2000 certificate authority that your certificate has been issued.

Steps

1. Navigate to the Windows 2000 certificate authority in your Web browser.

The URL is: http://host/certsrv

host is the IP address or fully-qualified host name of the Windows 2000 Server hosting the certification authority.

- 2. Select Check on a pending certificate and click Next.
- 3. Select your certificate and click Next.
- 4. Click the link to install the certificate automatically.
- 5. Start the Microsoft Management Console (MMC). To do this, from the Start menu, select **Run**. Then enter mmc.
- **6.** If you have not done so already, add the Certificates Current User snap-in to the MMC by performing the following steps:
 - a. From the File menu, select Add/Remove Snap-in.
 - b. Click Add, select Certificates, and click Add.
 - c. Select My User Account.
 - d. Click Finish.

- **7.** Export the certificate from the Certificates Current User store by performing the following steps:
 - **a.** In the MMC, right-click the certificate, which is in the Personal or Certificates folder of the Certificates Current User store, and then select **Export** from the All Tasks menu.
 - **b.** Use the **Certificate Export** wizard to export the certificate, including its private key, to a file.
- **8.** Import the certificate into the Certificates (Local Computer) store by performing the following steps:
 - **a.** In the MMC, right-click the Certificates folder in the Certificates (Local Computer) store, and then select **Import** from the All Tasks menu.
 - b. Use the Certificate Export wizard to import the certificate.

Note: Although the MMC allows you to copy a certificate from one store to another, the installation will not succeed unless you export the certificate from the first store and import the certificate into the second store.

Requesting a signed certificate from a non-Windows 2000 certificate authority

You can request a signed certificate from a non-Windows 2000 certificate authority.

Before you begin

To request a signed certificate from a non-Windows 2000 certificate authority, you should follow the instructions on the certificate authority's Web site. Non-Windows 2000 certificate authorities typically require you to generate and submit a certificate signing request.

To generate a certificate signing request for a certificate that you install on a Windows client, you can use the openssl utility. For more information, search the Internet for "openssl."

Step

1. From the command-line interface, enter the following command:

```
keymgr generate cert cert_file_name KeyLen = key_length KeyFile =
key_file_name Common = storage_system_common_name Country =
two_character_country_code State = full_state_name Local =
organization_locality Organ = organization_name Unit = unit_name
```

cert_file_name is the name of the file into which to store the unsigned certificate. Data ONTAP stores this file in the /etc/keymgr/cert directory.

key_length is the length of the private key in bits. For example, 1024.

key_file_name is the name of the file in which to store the private key. Data ONTAP stores this file in the /etc/keymgr/key directory.

two_character_country_code is the two-character abbreviation (without punctuation) for the country where the storage system is located. For example, US or CA.

full_state_name is the full name of the state where the storage system is located. For example, California or Washington.

organization_name is the name of the organization or company running the storage system.

organization_locality is the city where the storage system is located. For example, Sunnyvale or Berkeley.

unit_name is name of the department or organization unit running the storage system.

Note: All symbols, such as ampersand (&) or at (@) symbols, must be spelled out in or omitted from the organization and unit names.

Installing a certificate signed by a non-Windows 2000 certificate authority on a Windows client

If you have requested a certificate signed by a non-Windows 2000 certificate authority, you must then install it on your Windows client.

Steps

1. Convert the signed certificate to the Windows PKCS12 (*.pfx) format.

For example, copy the certificate into a file and then use the openssl utility to convert it to the Windows PKCS12 (*.pfx) format. For more information, search the Internet for "openssl."

- 2. Start the Microsoft Management Console (MMC). To do this, from the Start menu, select **Run**. Then enter mmc.
- **3.** If you have not done so already, add the Certificates (Local Computer) snap-in to the MMC by performing the following steps:
 - a. From the File menu, select Add/Remove Snap-in.
 - b. Click Add, select Certificates, and click Add.
 - c. Select Computer Account and click Next.
 - d. Select Local Computer and click Finish.
- **4.** Import the certificate into the Certificates (Local Computer) store by performing the following steps:
 - **a.** In the MMC, right-click the Certificates folder in the Certificates (Local Computer) store, and then select **Import** from the All Tasks menu.
 - **b.** Use the **Certificate Import** wizard to import the certificate.

Installing a signed certificate on a storage system

You need to install a signed certificate on a storage system if you are going to use the certificate authentication method for the IPsec protocol.

Before you begin

If the certificate was signed by a Windows 2000 certificate authority, you should install the certificate on a Windows client and export the certificate, including its private key, to a file.

Steps

1. Copy the signed certificate onto the root volume of the storage system.

Example

Mount the storage system's root volume on an NFS client, such as your administration console, and then copy the file containing the signed certificate onto the storage system's root volume.

2. If the signed certificate is in the Windows PKCS12 (*.pfx) format, convert it to the X.509 (*.pem) format.

You can use the openssl utility to convert the certificate to the X.509 (*.pem) format. For more information, search the Internet for "openssl."

3. Install the signed certificate by entering the following command:

keymgr install cert signed_certificate_file_name

signed_certificate_file_name is the full path to the file containing the signed certificate.

Installing root certificates on a storage system

You must install one or more root certificates in each of the storage systems that will be part of a security association among clients and storage systems.

Steps

- **1.** Download the root certificate (in PEM format, if possible) from the certificate authority's Web site.
- 2. Copy the root certificate onto the root volume of the storage system.

Example

Mount the storage system's root volume on an NFS client, such as your administration console, and then copy the file containing the root certificate onto the storage system's root volume.

3. If the root certificate is not in PEM format, convert it to PEM format.

You can convert the certificate using the openssl utility. For more information, search the Internet for "openssl."

4. Install the root certificate. From the storage system command line, enter the following command:

keymgr install root path

path is the full path and file name of the root certificate.

Specifying the subset of root certificates that Data ONTAP uses for certificate authentication

By default, Data ONTAP uses all of your storage system's root certificates for certificate authentication. You can specify that Data ONTAP should use only a subset of these root certificates for certificate authentication.

Step

1. From the storage system command line, enter the following command:

ipsec cert set -r file_names

file_names is a space-delimited list of 1 to 15 names of files containing root certificates that you downloaded and installed previously. Data ONTAP uses this subset of root certificates for certificate authentication, ignoring all other root certificates.

Note: To remove root certificates from this subset, repeat this step, specifying a new subset.

Viewing the subset of root certificates that Data ONTAP uses for certificate authentication

You can use the ipsec cert show command to view the subset of root certificates that Data ONTAP is currently using for certificate authentication.

Step

1. From the command-line interface, enter the following command:

ipsec cert show

Installing root certificates on a Windows client

After you have installed root certificates on your storage system, you must also install them on your Windows clients.

Steps

- 1. Download the root certificate (in CER format, if possible) from the certificate authority's Web site.
- 2. If the root certificate is not in CER format it, convert it to CER format.

You can convert the certificate using the openssl utility. For more information, search the Internet for "openssl."

3. Start the Microsoft Management Console (MMC).

From the Start menu, select **Run**. Then enter mmc.

- 4. Import the root certificate by performing the following steps:
 - **a.** Right-click the Trusted Root Certification Authorities folder in the Certificates (Local Computer) store, and then select **Import** from the All Tasks menu.
 - **b.** Use the **Certificate Import** wizard to import the root certificate.

Enabling the IPsec certificate authentication mechanism on a storage system

After your certificates are installed on a storage system, you must enable the IPsec security authentication mechanism.

Step

1. From the command line, enter the following command:

```
ipsec cert set -c signed_certificate_file -k private_key_file
```

signed_certificate_file is the full path to the file containing the signed certificate. *private_key_file* is the full path to the file containing the private key for the signed certificate.

Enabling the IPsec certificate authentication mechanism on a Windows client

You must enable the IPsec certificate authentication mechanism on a Windows client before you can use IPsec.

Steps

- 1. Start the Microsoft Management Console (MMC). From the Start menu, select **Run**. Then enter mmc.
- 2. If you have not done so already, add the IP Security Policies on Local Computer snap-in to the MMC by performing the following steps:
 - a. From the File menu, select Add/Remove Snap-in.
 - b. Click Add, then select IP Security Policy Management, and click Add.
 - c. Select Local computer and click Finish.

- 3. Right-click IP Security Policies on Local Computer, and then select Create IP Security Policy.
- 4. Use the IP Security Policy wizard to create an IPsec policy.
- **5.** In the MMC console, right-click your new IPsec policy, which is in the IP Security Policies on Local Computer store, and then select **Properties**.
- 6. Select Add.
- 7. Use the Security Rule wizard to create a security rule.
- **8.** For the authentication method, select **Use a certificate from this certificate authority (CA)**, click **Browse**, and then select the certificate that you installed previously.

Kerberos support

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 later 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

Note: To authenticate a client by using the Kerberos key-exchange mechanism, the storage system should have enough space in its root volume to store the security credentials of the client. If Kerberos support is enabled, the system administrator must ensure that the storage system has at least four kilobytes of free space in its root volume at all times.

Configuring preshared keys

You can configure preshared keys if you want to use a simple encryption system on a pair of endstations. Preshared key configuration requires manual management on each end-station. 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.

Steps

- 1. Create a file named psk.txt in the /etc directory.
- 2. Decide an ASCII text key that you use for authenticating the client and storage system.
- 3. In the psk.txt file, enter a line using the following format:

ip_address key

ip_address is the IP address of the client.

key is the preshared key you decided upon.

Example

192.0.2.1 ag8key

4. Copy this file to both the storage system and the client.

Enabling or disabling IPsec

You can use the options ip.ipsec.enable command to enable or disable IPsec on your storage system.

Step

1. From the command line, enter the following command:

```
options ip.ipsec.enable {on off}
```

on-Enables IPsec.

off—Disables IPsec.

Security policies and IPsec

You can use the ipsec command to add, modify, display, delete, and monitor security policies in your Security Policy Database (SPD) and on your storage system.

Next topics

Creating a security policy on page 170 Security policy options on page 171 Displaying existing security policies on page 172 Deleting a security policy on page 172

Creating a security policy

You need to create a security policy for your storage system and its clients to implement IPsec.

Step

1. Enter the following command:

```
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 ] [-1 {restrict|permit}]
```

Example

```
ipsec policy add -s 192.0.2.5 -t 192.0.1.12/24[139] -p esp -e des -a ah
-d in -l restrict
```

Security policy options

You must select a number of security policy options when you create a security policy on your storage system and its Windows clients.

When you create security policies, you must select from the following required and optional parameters on your storage system. You must also select corresponding values 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)		
encryption	-e	Optional. If the ESP protocol is selected, DES, triple DES, or no encryption can be specified. If this option is not specified, the best algorithm is selected based on the peer capabilities.		
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.		
protocol	-f	Optional. Specifies an upper-layer protocol by number.		
permission level	-1	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 ipsec policy show command.		

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 parameters.

About this task

You can display the contents of the Security Policies Database (SPD) by a combination of these parameters:

- Source and destination addresses
- Security protocol (AH or ESP)
- Direction (relative to your storage system)
- Specifications of upper-layer protocols

Step

1. From the command line, enter the following command:

```
ipsec policy show [-s src_ip] [-t dst_ip] [-f ip_protocol] [-d {in|out}]
[-p {esp|ah}]
```

src_ip is the source IP address.

dst_ip is the destination IP address.

ip_protocol is an upper-layer protocol expressed as a numeric protocol number. For example, the protocol number is 6 for TCP and 17 for UDP.

Example

The following example displays security policy information for the device that has a source IP address (-s) of 192.0.2.17:

```
ipsec policy show -s 192.0.2.17
```

Index IPAddress /prefix/port/protocol Dir/Policy Alg/SecLevel
-----1 192.0.2.17 / 0/ [any]/any in /IPSEC esp/Default

Deleting a security policy

You can delete an obsolete security policy and replace it with an up-to-date one.

About this task

You can remove entries from the security policy database 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

Step

1. From the command line, enter the following command:

```
ipsec policy delete {all|-i index} [[-s src_ip|-t dst_ip] -d {in|out} [-
m]]
```

index is the SPD index number of the policy that you want to delete.

src_ip is the source IP address.

dst_ip is the destination IP address.

Note: You must delete the same policies from corresponding clients.

Viewing IPsec statistics

You can use the ipsec stats command to view the cumulative IPsec statistics. You can use these statistics to verify IPsec configuration and monitor protocol processing, and to view IPsec violations.

About this task

The ipsec stats 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. It keeps track of the 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

Step

1. Enter the following command, depending on whether you want to view or clear the IPsec statistics:

If you want to	Enter the following command		
View IPsec statistics	ipsec stats		
Reset the IPsec statistics counter	ipsec stats -z		

Example

The following output shows the statistics provided by the ipsec stats command.

```
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
Ouick : 7310
IKE output packets
Identity Protection : 108
Informational : 10
Ouick : 3663
```

Viewing security associations

You can view the currently active security associations on your storage system.

About this task

You can view any of the following security associations:

- The entire contents of the Security Associations Database (SAD)
- An individual entry in the SAD identified by the Security Parameter Index (SPI)
- 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

Note: To learn the SPI for a database entry, you must first display the entire contents of the SAD.

Step

1. From the command line, enter the following command:

ipsec sa show [spi options]

spi is the Security Parameter Index number that identifies an individual entry in the Security Associations Database.

options include the source and destination IP addresses, and the encryption protocol, either esp (ESP based on RFC 2405) or ah (AH based on RFC 2402).

Example

The following example displays security association information for the device that has a source IP address of 192.0.2.17:

```
ipsec sa show 1 -s 192.0.2.17 -p esp
Alg/State/Spi Current Bytes/CreatedTime SrcIPAddr->DstIPAddr
esp/M/0001388 0/20 Aug 2002 17:28:19 192.0.2.17->192.0.2.20
```

The possible values for state are:

- M—Mature and active
- D—Dead
- d—Dying

• L—Larval (uninitiated)

How to diagnose network problems

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

cdpd	The cdpd command displays information of the devices that advertise themselves by using the CDPv1 protocol. You can use this command to view information about the CDP neighbors of the storage system and therefore, detect network connectivity.
netdiag	The netdiag command continuously gathers and analyzes statistics, and performs diagnostic tests. These diagnostic tests identify and report problems with your physical network or transport layers and suggest remedial action.
	For a full description of the netdiag command along with all available options, see the na_netdiag(1) man page.
ping	You can use the ping command to test whether your storage system can reach other hosts on your network.
	For a full description of the ping command, see the na_ping(1) man page.
pktt	You can use the pktt command to trace the packets sent and received in the storage system's network.
	For a full description of the pktt command, see the na_pktt(1) man page.
ping6	To reach IPv6 hosts, you can use the ping6 command. Starting with Data ONTAP 7.3.3, you can use the ping command for reaching IPv6 hosts. You can use the ping6 and ping commands with all types of IPv6 addresses.
	The -d option in the ping6 command specifies the number of data bytes to be sent. In addition, you might need to specify the -b option to extend the socket buffer size.
	You must use the -b option with the ping6 command when pinging hosts with jumbo frames. For the pinging to succeed with jumbo frames, the buffer must be large enough to reassemble IP fragments.
	For example, when pinging an IPv6 address with an 8900 byte payload and specifying a 9000 byte buffer, you should use the following command:
	ping6 -d 8900 -b 9000 2001:0db8::99
	In the previous example, setting the buffer size to 8901 or 8902 bytes might not be

In the previous example, setting the buffer size to 8901 or 8902 bytes might not be adequate and might cause the ping6 command to fail. Increasing the buffer size to 10000 allows the ping to succeed in both directions.

ndp You can use the ndp command to control the address mapping table used by Neighbor Discovery Protocol (NDP).

For a full description of the ndp command, see the na_ndp(1) man page.

traceroute6 You can use the traceroute6 command to trace the route that the IPv6 packets take to a network node.

For a full description of the traceroute6 command, see the na_traceroute6(1) man page.

Next topics

Diagnosing transport layer problems on page 178 Viewing diagnostic results on page 179 How to diagnose ping problems on page 180 Protecting your storage system from forged ICMP redirect attacks on page 181

Related references

Error codes for the netdiag command on page 213

Diagnosing transport layer problems

You can use the netdiag -t command to diagnose problems with the transport layer in your storage system.

Step

1. Enter the following command:

netdiag -t

Example

A storage system whose TCP window size is smaller than the recommended value displays the following output:

Performing transport layer diagnostics.... The TCP receive window advertised by CIFS client 192.0.2.13 is 8760. This is less than the recommended value of 32768 bytes. You should increase the TCP receive buffer size for CIFS on the client. Press enter to continue.

Viewing diagnostic results

You can use the netdiag -s command to view a summary of the various diagnostic checks and tests performed on your storage system.

About this task

If you enable the IPv6 option, you can view the IPv4 and IPv6 statistics in the network layer diagnostic summary.

Step

1. Enter the following command:

netdiag -s

Example

The following output shows some issues in IPv6 configuration of the network layer.

netdiag -s

Physical Layer Diagnostics Summary:

Interface	H/W Status	Link	Configured UP	Speed Mismatch	Duplex Mismatch	AutoNeg Mismatch
e0a	OK	-	N	-	-	-
e0b	OK	-	Ν	-	-	-
e0c	OK	Y	Y	N	-	-
e0d	OK	-	Ν	-	-	-
Network La	aver Dia	aqnost	ics Summary:			
Protocol		Statu	3			
IP		OK				
IPv6		Prob				
Transport	Layer 1	Diagno	stics Summar	y:		
Protocol		Statu	3			
TCP		OK				
UDP		OK				
Use netdia	ag with	out the	e -s option	for detail	S	

How to diagnose ping problems

You can use the Data ONTAP ping throttling mechanism and its ip.ping_throttle.drop_level option to help avoid denial-of-service attacks that can occur when using ICMP.

The ping throttling mechanism is active in intervals of 1 second. If the number of ICMP echo and reply packets that the storage system receives in a one-second interval exceeds the ping throttling threshold, the storage system drops all subsequent packets that are received within that one-second interval.

Note: Regardless of whether the ping throttling threshold has been reached, clients that send more than 16 packets per second to a storage system might experience packet loss. To allow clients to send more than 16 packets per second, you must disable ping throttling.

If your storage system supports a very large number of CIFS clients that use ICMP pings to determine CIFS shares accessibility, you can increase the ping throttling threshold value in the <code>ip.ping_throttle.drop_level</code> option.

If a large number of CIFS clients are experiencing temporary or persistent unavailability of the storage system, you should check to see if the ping throttling threshold has been exceeded for the storage system. If so, you can increase the ping throttling threshold value.

Next topics

Increasing the ping throttling threshold value on page 180 Checking the ping throttling threshold status on page 181 Disabling ping throttling on page 181

Increasing the ping throttling threshold value

If your storage system supports a very large number of CIFS clients that use ICMP pings to determine CIFS shares accessibility, you might need to increase the ping throttling threshold value.

Step

1. Enter the following command:

options ip.ping_throttle.drop_level packets_per_second

packets_per_second specifies the maximum number of ICMP echo or echo reply packets (ping packets) that the storage system accepts per second. Any further packets within 1 second are dropped. The default value is 150.
Checking the ping throttling threshold status

If a large number of CIFS clients are experiencing temporary or persistent unavailability of the storage system, you should check if the ping throttling threshold has been exceeded for the storage system.

Step

1. Enter the following command:

```
netstat -p icmp
```

Result

The resulting report lists the number of pings and ping replies that have been dropped, if any.

If the number of pings dropped, the number of ping replies dropped, or the number of both pings and ping replies dropped is greater than zero, you should increase the <code>ip.ping_throttle.drop_value</code> option to a number that is higher than the current value.

Disabling ping throttling

To allow clients to send more than 16 packets per second, you need to disable ping throttling.

Step

1. Enter the following command:

```
options ip.ping_throttle.drop_level 0
```

Protecting your storage system from forged ICMP redirect attacks

You can disable ICMP redirect messages to protect your storage system against forged ICMP redirect attacks.

About this task

To efficiently route a series of datagrams to the same destination, your storage system maintains a route cache of mappings to next-hop gateways. If a gateway is not the best next-hop for a datagram with a specific destination, the gateway forwards the datagram to the best next-hop gateway and sends an ICMP redirect message to the storage system. By forging ICMP redirect messages, an attacker can modify the route cache on your storage system, causing it to send all of its communications through the attacker. The attacker can then hijack a session at the network level, easily monitoring, modifying, and injecting data into the session.

Step

1. Enter the following command:

options ip.icmp_ignore_redirect.enable on

Your storage system now ignores ICMP redirect messages.

For more information about the ip.icmp_ignore_redirect.enable option, see the na_options(1) man page.

Note: By default, the ip.icmp_ignore_redirect.enable option is off.

Network interface statistics

You can use the ifstat command to view statistics for the network interfaces supported by Data ONTAP. To determine the Ethernet controllers in your system, you can use the sysconfig command.

Next topics

Statistics for Gigabit Ethernet controller IV - VI and G20 interfaces on page 183 Statistics for Gigabit and 10 Gigabit Ethernet controllers T204, T210, and T320 interfaces on page 187 Statistics for the N3700 network interfaces on page 190 Statistics for the BGE 10/100/1000 Ethernet interface on page 193

Statistics for Gigabit Ethernet controller IV - VI and G20 interfaces

The ifstat command output displays several statistics when you use the command for the Gigabit Ethernet controllers and G20 interfaces.

The statistics in this section are for the following controllers:

- 10/100 Ethernet controller IV
- 10/100/1000 Ethernet controller IV through VII
- Gigabit Ethernet controller IV through VI
- 10/100/1000 Ethernet controller G20
- Gigabit Ethernet controller G20

RECEIVE section statistics

The following table describes the statistics in the RECEIVE section of the ifstat command output.

Statistic	Definition
Frames/second	Rate of received frames per second.
Bytes/second	Rate of received bytes per second.
Errors/minute	Rate of errors (which led to frames being lost) per minute.
Discards/minute	Rate per minute of packets discarded due to unavailable resources.
Total frames	Total frames that are received on the interface.
Total bytes	Total bytes that are received on the interface.

Statistic	Definition
Total errors	Total errors that occur on the interface.
Total discards	Total number of packets that were discarded even though no errors were detected. This number is a sum of the "No buffers", "Bus overruns", and "Queue overflows" statistics.
Multi/broadcast	Total number of multicast or broadcast packets received.
Alignment errors	Number of frames that are both misaligned and contain CRC errors.
Non-primary u/c	Number of Ethernet frames received for the partner's MAC address after a failover in an active/active 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.
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.

Statistic	Definition
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.
Jumbo	Number of good packets received that were larger than the standard Ethernet packet size when jumbo frames are enabled.
Reset	Number of times the driver reset the NIC because the NIC was in a bad state.
Reset1	Number of times the driver reset the NIC because the NIC was in a bad state.
Reset2	Number of times the driver reset the NIC because the NIC was in a bad state.

TRANSMIT section statistics

The following table describes the statistics in the TRANSMIT section of the ifstat command output.

Statistic	Meaning
Frames/second	Rate of transmitted frames per second.
Bytes/second	Rate of transmitted bytes per second.
Errors/minute	Rate of errors (which led to frames being lost) per minute.
Discards/minute	Rate per minute of packets discarded due to unavailable resources.
Total frames	Total frames that are transmitted on the interface.
Total bytes	Total bytes that are transmitted on the interface.
Total errors	Total errors that occur on the interface.
Total discards	Total number of packets that were discarded even though no errors were detected. This number is a sum of the "No buffers" and "Queue overflows" statistics.
Multi/broadcast	Total number of multicast or broadcast packets transmitted.

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 frame size (1,518 bytes).

LINK INFO section statistics

The following table describes the statistics in the LINK INFO section of the *ifstat* command output.

Statistic	Meaning
Current state	Current state of the interface:
	• up or down—The state of the link.
	• cfg_down—The interface is configured to the down status.
	• enabling—The interface is moving to the up status.
Up to downs	Number of times the link switched between the up status and the down status.

Statistic	Meaning
Auto	 Operational state of autonegotiation: 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 might also indicate that the interface is in the down status.
Speed	Speed of link negotiated or set.
Duplex	Duplex of the link negotiated or set.
Flow control	The operational flow control setting.

Related tasks

Viewing or clearing network interface statistics on page 54

Statistics for Gigabit and 10 Gigabit Ethernet controllers T204, T210, and T320 interfaces

The ifstat command output displays several statistics when you use the command for the 10/100/1000 Ethernet controllers T204V and T204E, and the 10 Gigabit Ethernet controllers T210 and T320.

RECEIVE section statistics

The following table describes the statistics in the RECEIVE section of the ifstat command output when you use the command on the 10/100/1000 Ethernet controllers T204V and T204E, and the 10 Gigabit Ethernet controllers T210 and T320.

Statistic	Meaning
Frames/second	Rate of received frames per second.
Bytes/second	Rate of received bytes per second.
Errors/minute	Rate of errors (which led to frames being lost) per minute.
Discards/minute	Rate per minute of packets discarded due to unavailable resources.
Total frames	Total frames that are received on the interface.
Total bytes	Total bytes that are received on the interface.
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", "Bus overruns", and "Queue overflows" statistics.
Multi/broadcast	Total number of multicast or broadcast packets received.
Alignment errors	Number of frames that are both misaligned and contain CRC errors.
Non-primary u/c	Number of Ethernet frames received for the partner's MAC address after a failover in an active/active 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.
Vlan untag drop	Number of untagged frames dropped on an interface that is configured to be part of a VLAN.
CRC errors	Number of packets received with bad CRC.
Runt frames	Number of received frames that were less than the minimum size (64 bytes) and had a valid CRC.
Long frames	Number of received frames that were greater than the maximum size and had a valid CRC.
Jabber	Number of received frames that were greater than the maximum size and had a bad CRC.
No buffer	Number of times the driver could not allocate a buffer and a packet was dropped. This might happen when your storage system is very busy. If the count increases continually, it might indicate that a software component is not returning buffers.
Jumbo	Number of good packets received that were larger than the standard Ethernet packet size when jumbo frames are enabled.
Pause Frames	Number of MAC Control PAUSE frames sent to the link partner to inhibit transmission of data frames for a specified period of time. This can help the partner from overrunning the controller's receive buffers.

TRANSMIT section statistics

The following table describes the statistics in the TRANSMIT section of the ifstat command output when you use the command on the 10/100/1000 Ethernet controllers T204V and T204E, and the 10 Gigabit Ethernet controllers T210 and T320.

Statistic	Meaning
Frames/second	Rate of transmitted frames per second.
Bytes/second	Rate of transmitted bytes per second.
Errors/minute	Rate of errors (which led to frames being lost) per minute.
Discards/minute	Rate per minute of packets discarded due to unavailable resources.
Total frames	Total frames that are transmitted on the interface.
Total bytes	Total bytes that are transmitted on the interface.
Total errors	Total errors that occur on the interface.
Total discards	Total number of packets that were discarded even though no errors were detected. This number is a sum of the "No buffers" and "Queue overflows" statistics.
Multi/broadcast	Total number of multicast or broadcast packets transmitted.
No buffers	Number of times the driver failed to allocate a buffer for the transmit packet.
Queue overflow	Number of outgoing packets dropped because the driver's queue was full. It might indicate a system problem.
Bus Underruns	FIFO goes empty before an internal End-Of-Packet indicator is read.
Pause Frames	Number of MAC Control PAUSE frames sent to the link partner to inhibit transmission of data frames for a specified period of time. This can help the partner from overrunning the controller's receive buffers.

LINK INFO section statistics

The following table describes the statistics in the LINK INFO section of the ifstat command output when you use the command on the 10/100/1000 Ethernet controllers T204V and T204E, and the 10 Gigabit Ethernet controllers T210 and T320.

Statistic	Meaning
Current state	 Current state of the interface: up or down—The state of the link. cfg_down—The interface is configured to the down status. enabling—The interface is coming to the up status.
Up to downs	Number of times the link switched between the up status and the down status.

Statistic	Meaning
Speed	Speed of link negotiated or set.
Duplex	Duplex of the link negotiated or set.
Flow control	The operational flow control setting.

Statistics for the N3700 network interfaces

The ifstat command output displays several statistics when you use the command on the N3700 network interfaces of the storage system.

RECEIVE section statistics

The following table describes the statistics in the RECEIVE section of the ifstat command output when you use the command on the N3700 network interfaces.

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 an active/active 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 address 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 address that had a code error signaled by the Physical (PHY) layer.
Dribble errors	The number of frames received by the MAC address with an alignment error. This is not used for 1000 Mb/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 the N3700 network interfaces.

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.	
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 failed to allocate a buffer for the transmit packet.	

Statistic	Meaning
CRC errors	Number of packets transmitted by the MAC address with CRC errors. This can happen only when the MAC address 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 address.
Long frames	Number of packets larger than the maximum frame size transmitted by the MAC address.
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 address due to excessive collisions during transmission.
	If 16 consecutive collisions occur during the transmission of a packet, the transmission is deferred and the MAC address 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 the N3700 network interfaces.

Statistic	Meaning
Current state	 Current state of the interface: up or down—The state of the link. cfg_down—The interface is configured to the down status. enabling—The interface is coming to the up status.
Up to downs	Number of times the link switched between the up status and the down status.
Speed	Speed of the link negotiated or set.
Duplex	Duplex of the link negotiated or set.
Flow Control	The operational flow control setting.

Statistics for the BGE 10/100/1000 Ethernet interface

The ifstat command output displays several statistics when you use the command on the BGE 10/100/1000 Ethernet interface.

RECEIVE section statistics

The following table describes the statistics in the RECEIVE section of the ifstat command output when you use the command on the BGE 10/100/1000 Ethernet interface.

Statistic	Meaning	
Frames/second	Rate of received frames per second.	
Bytes/second	Rate of received bytes per second.	
Errors/minute	Rate of errors (which led to frames being lost) per minute.	
Discards/minute	Rate per minute of packets discarded due to unavailable resources.	
Total frames	Total frames that are received on the interface.	
Total bytes	Total bytes that are received on the interface.	
Total errors	Total errors that occur on the interface.	
Total discards	Total number of packets that were discarded even though no errors were detected. This number is a sum of the "No buffers", "Bus overruns", and "Queue overflows" statistics.	
Multi/broadcast	Total number of multicast or broadcast packets received.	
Alignment errors	Number of frames that are both misaligned and contain CRC errors.	
Non-primary u/c	Number of Ethernet frames received for the partner's MAC address after a failover in an active/active 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.	
Vlan untag drop	Number of untagged frames dropped on an interface that is configured to be part of a VLAN.	
CRC errors	Number of packets received with bad CRC.	
Runt frames	Number of received frames that were less than the minimum size (64 bytes) and had a valid CRC.	

Statistic	Meaning
Fragment	Number of received frames that were less than the minimum size and had a bad CRC.
Long frames	Number of received frames that were greater than the maximum size and had a valid CRC.
Jabber	Number of received frames that were greater than the maximum size and had a bad CRC.
No buffer	Number of times the driver could not allocate a buffer and a packet was dropped. This might happen when your storage system is very busy. If the count increases continually, it might indicate that a software component is not returning buffers.
Xon	Number of XON frames received when receive or full flow control is enabled.
Xoff	Number of XOFF frames received when receive or full flow control is enabled.
Jumbo	Number of good packets received that were larger than the standard Ethernet packet size when jumbo frames are enabled.
Ring full	Not used. Ignore.
Jumbo error	Error detected while processing a jumbo packet. Packet is discarded.

TRANSMIT section statistics

The following table describes the statistics in the TRANSMIT section of the ifstat command output when you use the command on the BGE 10/100/1000 Ethernet interface.

Statistic	Meaning
Frames/second	Rate of transmitted frames per second.
Bytes/second	Rate of transmitted bytes per second.
Errors/minute	Rate of errors (which led to frames being lost) per minute.
Discards/minute	Rate per minute of packets discarded due to unavailable resources.
Total frames	Total frames that are transmitted on the interface.
Total bytes	Total bytes that are transmitted on the interface.
Total errors	Total errors that occur on the interface.

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.	
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.	
Jumbo	Number of packets transmitted that were larger than the standard Ethernet packet size when jumbo frames are enabled.	
Deferred	Number of frames for which the first transmission was delayed because the medium was busy.	
MAC Internal	Number of frames not transmitted due to an internal MAC sublayer error.	

LINK INFO section statistics

The following table describes the statistics in the LINK INFO section of the ifstat command output when you use the command on the BGE 10/100/1000 Ethernet interface.

Statistic	Meaning
Current state	 Current state of the interface: up or down—The state of the link. cfg_down—The interface is configured to the down status. enabling—The interface is coming to the up status.
Up to downs	Number of times the link switched between the up status and the down status.
Speed	Speed of link negotiated or set.
Duplex	Duplex of the link negotiated or set.
Flow control	The operational flow control setting.

Related tasks

Viewing or clearing network interface statistics on page 54

Ways to improve your storage system's performance

You can improve your storage system's performance by performing certain configuration procedures, such as using vifs, correcting duplex mismatches, and upgrading to Ethernet interfaces.

The following configuration procedures might improve the performance of your storage system:

- Using static or dynamic multimode vifs to aggregate the bandwidth of multiple interfaces
- Using jumbo frames with your network interfaces to reduce CPU processing overhead
- Upgrading to a faster network interface You can increase the storage system's performance by upgrading to a faster network interface (10 Gigabit Ethernet interfaces).
- Correcting duplex mismatches on 10Base-T or 100Base-T Ethernet networks On 10Base-T or 100Base-T Ethernet networks, the speed and duplex settings for the interfaces at both ends of a link must match exactly. You can use the ifconfig *interface* command to check the duplex setting of your storage system's interface.

If the setting is to autonegotiate, the ifconfig command displays a setting that begins with auto (for example, auto-100tx-fd-up). Otherwise, the ifconfig command displays the configured media type setting (for example, 100tx-fd-up).

Note: If one end of the link is set to autonegotiate, the other end must also be set to autonegotiate; otherwise, a mismatch might occur. You can determine the negotiated setting with the ifstat command.

- Using iSCSI multiconnection sessions to balance the load across interfaces For each iSCSI session, multiple connections are created. The number of allowed connections is negotiated during login and session creation. While it is possible to create multiple connections over a single physical interface, it is best to use multiple physical interfaces for bandwidth enhancement.
- Enabling fast path on your storage system Fast path provides load balancing by sending responses on the same network interface that receives the incoming requests and improved performance by skipping routing table lookups.

Related concepts

Static multimode vif on page 109 Dynamic multimode vif on page 110 What jumbo frames are on page 29

Related tasks

Specifying a media type for a network interface on page 43 *Enabling or disabling fast path* on page 64

IP port usage on a storage system

The Data ONTAP services file is available in the /etc directory. The /etc/services file is in the same format as its corresponding UNIX system's /etc/services file.

Next topics

Host identification on page 199 /etc/services NNTP and TTCP ports on page 202 NFS-enabled ports on page 202 Ports not listed in /etc/services on page 203 FTP on page 204 **SSH** on page 204 *Telnet* on page 205 *SMTP* on page 205 *Time service* on page 206 **DNS** on page 206 **DHCP** on page 207 *TFTP* on page 207 *HTTP* on page 207 Kerberos on page 208 NFS on page 208 CIFS on page 209 SSL on page 209 **SNMP** on page 210 **RSH** on page 211 Syslog on page 211 The routed daemon on page 211 *NDMP* on page 212 SnapMirror and Snap Vault on page 212

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.

The following table gives a sample content of the /etc/services file.

Service	Port/ Protocol	Description
ftp-data	20/tcp	# File transfer protocol
ftp	21/tcp	# File transfer protocol
ssh	22/tcp	# SecureAdmin rsh replacement
telnet	23/tcp	# Remote login (insecure)
smtp	25/tcp	# outbound connections for autosupport
time	37/tcp	# Time Service
time	37/udp	# Time Service
domain	53/udp	# DNS - outbound only
domain	53/tcp	# DNS zone transfers - unused
dhcps	67/udp	# DHCP server - outbound only
dhcp	68/udp	# DHCP client - only first-time setup
tftp	69/udp	# Trivial FTP - for netboot support
http	80/tcp	# HTTP license, FilerView, SecureAdmin
kerberos	88/udp	# Kerberos 5 - outbound only
kerberos	88/tcp	# Kerberos 5 - outbound only
portmap	111/udp	# aka rpcbind, used for NFS
portmap	111/tcp	# aka rpcbind, used for NFS
nntp	119/tcp	<pre># unused, shouldn't be listed here.</pre>

Service	Port/ Protocol	Description
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	<pre># rsh, insecure remote command execution.</pre>
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	<pre># primary NFS service</pre>
nfsd	2049/tcp	<pre># primary NFS service</pre>
ttcp	5001/udp	# unused, shouldn't be listed here.
ttcp	5001/tcp	<pre># unused, shouldn't be listed here.</pre>
ndmp	10000/tcp	# for network backups
snapmirror	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 not used by your storage system and should never be detected by a port scanner.

NFS-enabled ports

Some ports (port numbers in the 600 range) on the storage system are NFS-enabled.

UDP	602	NFS mount daemon (mountd)
ТСР	603	NFS mount daemon (mountd)
UDP	604	NFS status daemon (statd, statmon)
ТСР	605	NFS status daemon (statd, statmon)
UDP	606	NFS lock manager (lockd, nlockmgr)
ТСР	607	NFS lock manager (lockd, nlockmgr)
UDP	608	NFS quota daemon (quotad, rquotad)

The following ports are found on the storage system with NFS-enabled:

On other systems, the ports appear as follows:

UDP	611	NFS mount daemon (mountd)
ТСР	612	NFS mount daemon (mountd)
UDP	613	NFS status daemon (statd, statmon)
ТСР	614	NFS status daemon (statd, statmon)
UDP	615	NFS lock manager (lockd, nlockmgr)
ТСР	616	NFS lock manager (lockd, nlockmgr)
UDP	617	NFS quota daemon (quotad, rquotad)

The following command on UNIX systems obtains the correct information by querying the port mapper on port 111:

```
toaster# rpcinfo -p sys1
```

program	vers	proto	port	service
100011	1	udp	608	rquotad
100021	4	tcp	607	nlockmgr
100021	3	tcp	607	nlockmgr
100021	1	tcp	607	nlockmgr
100021	4	udp	606	nlockmgr
100021	3	udp	606	nlockmgr
100021	1	udp	606	nlockmgr
100024	1	tcp	605	status
100024	1	udp	604	status
100005	3	tcp	603	mountd
100005	2	tcp	603	mountd
100005	1	tcp	603	mountd
100005	3	udp	602	mountd
100005	2	udp	602	mountd
100005	1	udp	602	mountd
100003	3	udp	2049	nfs
100003	2	udp	2049	nfs
100000	2	tcp	111	rpcbind
100000	2	udp	111	rpcbind

Note: The port numbers listed for mountd, statd, lockd, and quotad are not committed port numbers. These services can be running on other ports of the storage systems. Because the system selects these port numbers at random when it boots, they are not listed in the /etc/services file.

Ports not listed in /etc/services

Some ports appear in a port scan but are not listed in the /etc/services file, for example, TCP ports 22 and 443.

Protocol	Port	Service
ТСР	22	SSH (SecureAdmin)
ТСР	443	SSL (SecureAdmin)
ТСР	3260	iSCSI-Target
UDP	XXXX	Legato ClientPack for your storage system runs on random UDP ports and is now deprecated. It is best to use NDMP to back up your storage system using Legato Networker.

The following ports appear in a port scan but are not listed in the /etc/services file.

Note: Disable open ports that you do not need.

FTP

File Transfer Protocol (FTP) uses TCP ports 20 and 21.

If you use FTP to transfer files to and from your storage system, the FTP port is required; otherwise, use FilerView or the following CLI command to disable the FTP port:

options ftpd.enable off

FTP is not a secure protocol for two reasons:

• When users log in to the system, user names and passwords are transmitted over the network in clear text format that can easily be read by a packet sniffer program.

These user names and passwords can then be used to access data and other network resources. You should establish and enforce policies that prevent the use of the same passwords to access storage systems and other network resources.

• FTP server software used on platforms other than storage systems contains serious securityrelated flaws that allow unauthorized users to gain administrative (root) access and control over the host.

Starting with Data ONTAP 7.3.1, FTP over IPv6 is supported.

For a detailed description of the FTP support for your storage system, see the *Data ONTAP File Access and Protocols Management Guide*.

SSH

Secure Shell (SSH) protocol is a secure replacement for RSH and runs on TCP port 22. This port appears in a port scan only if the SecureAdmin software is installed on your storage system.

There are three commonly deployed versions of the SSH protocol:

- SSH version 1—is secure than RSH or Telnet, but is vulnerable to TCP session attacks. This vulnerability to attack lies in the SSH protocol version 1 itself and not in the associated storage system products.
- SSH version 2—has a number of feature improvements over SSH version 1 and is less vulnerable to attacks.
- SSH version 1.5—is used to identify clients or servers that support both SSH versions 1 and 2.

To disable SSH support or to close TCP port 22, you must use the following CLI command:

secureadmin disable ssh

Telnet

Telnet is used for administrative control of your storage system and uses TCP connections on port 23. Telnet is more secure than RSH, as secure as FTP, and less secure than SSH or Secure Socket Layer (SSL).

Telnet is less secure than SSH and SSL because:

• When users log in to a system, such as your storage system, user names and passwords are transmitted over the network in clear text format.

Clear text format can be read by an attacker by using a packet sniffer program. The attacker can use these user names and passwords to log in to your storage system and execute unauthorized administrative functions, including destruction of data on the system. If administrators use the same passwords on your storage system as they do on other network devices, the attacker can use these passwords to access the resources of the storage system as well.

Note: To reduce the potential for attack, you must establish and enforce policies preventing administrators from using the same passwords on your storage system that they use to access other network resources.

• Telnet server software used on other platforms (typically in UNIX environments) have serious security-related flaws that allow unauthorized users to gain administrative (root) control over the host.

Telnet is also vulnerable to the same type of TCP session attacks as SSH protocol version 1. However, TCP session attacks are less common because a packet sniffing attack is easier.

To disable Telnet, you must set options telnet.enable to off.

Starting with Data ONTAP 7.3.1, Telnet supports IPv6. However, if you have enabled the IPv6 option when the storage system is in operation (not during setup), you must restart the Telnet service. That is, you need to turn off and then turn on the Telnet service for connecting over IPv6.

SMTP

Simple Mail Transport Protocol (SMTP) uses TCP port 25. Your storage system does not listen on this port but makes outgoing connections to mail servers using this protocol when sending AutoSupport e-mail.

Time service

Your storage system supports two different time service protocols, TIME protocol and Simple Network Time Protocol (SNTP).

The following are the two different time service protocols:

- TIME protocol (also known as rdate)—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.
- SNTP—specified in the RFC 2030 standard and is provided only on UDP port 123.

When your storage system has the timed.enable option set to on and a remote protocol (rdate or SNTP) 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 SNTP. You can use the rdate command to use the rdate TIME protocol.

You should set the timed.enable option to on in an active/active configuration.

DNS

The Domain Name System (DNS) uses UDP port 53 and TCP port 53. Your storage system does not typically listen on these ports because it does not run a domain name server. However, if DNS is enabled on your storage system, it makes outgoing connections using UDP port 53 for host name and IP address lookups.

The storage system never uses TCP port 53 because this port is used explicitly for communication between DNS servers. Outgoing DNS queries by your storage system are disabled by turning off DNS support. Turning off DNS support protects against receiving bad information from another DNS server.

Because your storage system does not run a domain name server, the name service must be provided by one of the following:

- Network information service (NIS)
- An /etc/hosts file
- Replacement of host names in the configuration files (such as /etc/exports, /etc/ usermap.cfg, and so on) with IP addresses

DNS must be enabled for participation in an Active Directory domain.

DHCP

Clients broadcast messages to the entire network on UDP port 67 and receive responses from the Dynamic Host Configuration Protocol (DHCP) server on UDP port 68. The same ports are used for the BOOTP protocol.

DHCP is used only for the first-time setup of your storage system. Detection of DHCP activity on your storage system by a port scan other than the activity during the first-time setup indicates a serious configuration or software error.

TFTP

Trivial File Transfer Protocol (TFTP) uses TCP port 69. It is used mostly for booting UNIX or UNIX-like systems that do not have a local disk (this process is also known as netbooting) and for storing and retrieving configuration files for devices such as Cisco routers and switches.

Transfers are not secure on TFTP because it does not require authentication for clients to connect and transfer files.

Your storage system's TFTP server is not enabled by default. When TFTP is enabled, the administrator must specify a directory to be used by TFTP clients, and these clients cannot access other directories. Even within the TFTP directory, access is read-only. TFTP should be enabled only if necessary. You can disable TFTP using the following option:

options tftpd.enable off

You can configure the maximum number of simultaneous connections handled by the TFTP server by using the tftpd.max_connections option. The default number of TFTP connections is 8. The maximum number of connections supported is 32.

HTTP

Hypertext Transport Protocol (HTTP) runs on TCP port 80 and is the protocol used by Web browsers to access Web pages.

Your storage system uses HTTP to access the following:

- Files when HTTP is enabled
- FilerView for graphical user interface (GUI) administration
- Secure FilerView when SecureAdmin is installed

Starting with Data ONTAP 7.3.1, HTTP over IPv6 is supported and can be used for file access. Starting with Data ONTAP 7.3.3, HTTP and HTTPS over IPv6 can also be used to access FilerView.

The SecureAdmin SSL interface accepts connections on TCP port 443. SecureAdmin manages the details of the SSL network protocol, encrypts the connection, and then passes this traffic through to

the normal HTTP FilerView interface through a loopback connection. This loopback connection does not use a physical network interface. HTTP communication takes place inside your storage system, and no clear text packets are transmitted.

HTTP is not vulnerable to security attacks because it provides read-only access to documents by unauthenticated clients. Although authentication is not typically used for file access, it is frequently used for access to restricted documents or for administration purposes, such as FilerView administration. The authentication methods defined by HTTP send credentials, such as user names and passwords, over the network without encryption. The SecureAdmin product is provided with SSL support to overcome this shortcoming.

Note: You can stop your storage system from listening for new connections by setting the options httpd.enable and httpd.admin.enable to off. If either of the options is set to on, your storage system will listen for new connections.

Kerberos

There are four Kerberos ports in the /etc/services file: TCP port 88, UDP port 88, TCP port 750, and UDP port 750. These ports are used only for outbound connections from your storage system. Your storage system does not run Kerberos servers or services and does not listen on these ports.

Kerberos is used by your storage system to communicate with the Microsoft Active Directory servers for both CIFS authentication and, if configured, NFS authentication.

NFS

Network File System (NFS) is used by UNIX clients for file access. NFS uses port 2049.

NFSv3 and NFSv2 use the portmapper service on TCP or UDP port 111. The portmapper service is consulted to get the port numbers for services used with NFSv3 or NFSv2 protocols such as mountd, statd, and nlm. NFSv4 does not require the portmapper service.

NFSv4 provides the delegation feature that enables your storage system to grant local file access to clients. To delegate, your storage system sets up a separate connection to the client and sends callbacks on it. To communicate with the client, your storage system uses one of the reserved ports (port numbers less than 1024). To initiate the connection, the client registers the callback program on a random port and informs the server about it.

With delegations enabled, NFSv4 is not firewall-friendly because several other ports need to be opened up as well.

Starting with Data ONTAP 7.3.1, IPv6 over NFS is supported.

You can disable the TCP and UDP ports by setting the nfs.tcp.enable and nfs.udp.enable options to off.

To disable NFS, you should use the nfs off command.

CIFS

Common Internet File Service (CIFS) is the successor to the server message block (SMB) protocol. CIFS is the primary protocol used by Windows systems for file sharing.

CIFS uses UDP ports 137 and 138, and TCP ports 139 and 445. Your storage system sends and receives data on these ports while providing CIFS service. If it is a member of an Active Directory domain, your storage system must also make outbound connections destined for DNS and Kerberos.

Starting with Data ONTAP 7.3.1, CIFS over IPv6 is supported. CIFS over IPv6 uses only port 445. Ports 137, 138, and 139 are used by NetBIOS, which does not support IPv6.

CIFS is required for Windows file service. You can disable CIFS using FilerView or by issuing the cifs terminate command on your storage system console.

Note: If you disable CIFS, be aware that your storage system's /etc/rc file can be set up to automatically enable CIFS again after a reboot.

SSL

The Secure Sockets Layer (SSL) protocol provides encryption and authentication of TCP connections. Data ONTAP supports SSLv2, SSLv3, and Transport Layer Security (TLS) version 1.0. You should use TLSv1.0 or SSLv3 because it offers better security than previous SSL versions.

When SecureAdmin is installed and configured on your storage system, it listens for SSL connections on TCP port 443. It receives secure Web browser connections on this port and uses unencrypted HTTP, running on TCP port 80, through a loopback connection to pass the traffic to FilerView. This loopback connection is contained within your storage system and no unencrypted data is transmitted over the network.

You can enable or disable SSL by using FilerView or with the following command:

```
secureadmin {enable|disable} ssl
```

For TLS to be used for communication, both the client requesting the connection and the storage system must support TLS.

TLS is disabled by default, and setting up SSL does not automatically enable TLS. Before enabling TLS, ensure that SSL has been set up and enabled. To enable or disable TLS, enter the following command:

```
options tls.enable {on|off}
```

SNMP

Simple Network Management Protocol (SNMP) is an industry-standard protocol used for remote monitoring and management of network devices over UDP port 161.

SNMP is not secure because of the following reasons:

• Instead of using encryption keys or a user name and password pair, SNMP uses a community string for authentication. The community string is transmitted in clear text format over the network, making it easy to capture with a packet sniffer.

Within the industry, devices are typically configured at the factory to use public as the default community string. The public password allows users to make queries and read values but does not allow users to invoke commands or change values. Some devices are configured at the factory to use private as the default community string, allowing users full read-write access.

• Even if you change the read and write community string on a device to something other than private, an attacker can easily learn the new string by using the read-only public community string and asking the router for the read-write string.

There are three versions of SNMP:

- SNMPv1 is the original protocol and is not commonly used.
- SNMPv2 is identical to SNMPv1 from a network protocol standpoint and is vulnerable to the same security problems. The only differences between the two versions are in the messages sent, messages received, and types of information. These differences are not important from a security perspective.
- SNMPv3 is the latest protocol version and includes security improvements but is difficult to implement and many vendors do not yet support it. SNMPv3 supports several different types of network encryption and authentication schemes. It allows for multiple users, each with different permissions, and solves SNMPv1 security problems while maintaining an important level of compatibility with SNMPv2.

SNMP is required if you want to monitor a storage system through an SNMP monitoring tool, such as DataFabric Manager. The SNMP implementation in the storage system allows read-only access. Regardless of the community string used, the user cannot issue commands or change variables using SNMP on your storage system.

You should use the snmp.access option to restrict SNMP access to a named set of trusted hosts.

You can disable SNMP entirely by setting the snmp.enable option to off to disable SNMP entirely.

The snmp community delete and snmp community add commands are used to change the community string to something other than the default value.

RSH

Remote Shell (RSH) protocol is used for remote command execution. It is less secure than TFTP and uses TCP port 514.

RSH is not secure because passwords are not required for login and commands are easy to misconfigure. Therefore, you should disable RSH by setting the rsh.enable option to off.

You should use the SSH supplied with SecureAdmin for remote command execution and login. If this is not possible, Telnet is preferred to RSH.

If RSH is the only alternative, follow these guidelines when using RSH:

- Specify only secure, trusted hosts in the /etc/hosts.equiv file.
- Always use IP addresses rather than host names in the /etc/hosts.equiv file.
- Always specify a single IP address with a single user name on each line in /etc/hosts.equiv file.
- Use the rsh.access option instead of the trusted.hosts option for access control.
- Make sure the ip.match_any_ifaddr option is set to off.

Syslog

Your storage system sends messages to hosts specified by the user in the /etc/syslog.conf file by using the syslog protocol on UDP port 514. It does not listen on this port, nor does it act as a syslog server.

The routed daemon

The routed daemon, routed, listens on UDP port 520. It receives broadcast messages from routers or other hosts using Routing Information Protocol (RIP). These messages are used by your storage system to update its internal routing tables to determine which network interfaces are optimal for each destination.

Your storage system never broadcasts RIP messages containing routes because Data ONTAP is not capable of acting as a router.

RIP is not secure because an attacker can easily send artificial RIP messages and cause hosts running the routed daemon (such as your storage system) to redirect network traffic to the attacker. The attacker can then receive and shift this traffic for passwords and other information and send it on to the actual destination, where the intrusion is undetected. This method can also be used as a starting point for TCP session attacks.

Because of these security issues, use static routes (those set up using the route command on your storage system) instead of using the routed daemon.

NDMP

Network Data Management Protocol (NDMP) runs on TCP port 10000 and is used primarily for backup of network-attached storage (NAS) devices, such as storage systems.

The protocol defines three authentication methods:

- NONE—allows authentication without restriction
- TEXT-sends a clear text password over the network, similar to Telnet or FTP
- MD5—uses the MD5 message digest algorithm along with a challenge-response message exchange to implement a secure login mechanism

Your storage system supports both the TEXT and MD5 authentication methods. Most NDMPenabled backup software uses MD5 by default.

To entirely disable the TEXT authentication method, you should set the ndmpd.authtype option to challenge.

To restrict NDMP commands to certain authorized backup hosts, you should use the ndmp.access option.

Regardless of the authentication method used, NDMP sends backup data in decrypted format over the network, as does most other backup software. A separate network optimized for backup is a common means to increase performance while retaining data security.

To disable NDMP, you should set the ndmp.enable option to off.

SnapMirror and SnapVault

SnapMirror and SnapVault use TCP port 10566 for data transfer. Network connections are always initiated by the destination system; that is, SnapMirror and SnapVault *pull* data rather than *push* data.

Authentication is minimal with both SnapMirror and SnapVault. To restrict inbound TCP connections on port 10566 to a list of authorized hosts or IP addresses, you should configure the snapmirror.access or snapvault.access option. When a connection is established, the destination storage system communicates its host name to the source storage system, which then uses this host name to determine if a transfer is allowed. You should confirm a match between the host name and its IP address. To confirm that the host name and the IP address match, you should set the snapmirror.checkip.enable option to on.

To disable SnapMirror, you should set the snapmirror.enable option to off. To disable SnapVault, you should set the snapvault.enable option to off.

Error codes for the netdiag command

Network error codes are generated by the netdiag command. They describe the network problems and suggest the actions that you can take.

The following table lists some network error codes, describes the problems that the error codes point to, and suggests actions that you can take to fix the problems.

Note: Only a small fraction of the possible network error messages are presented in the following table. If you receive any problem code not listed in this table, contact your technical support.

Error code	Description	Recommended actions
201	Link not detected.	Complete the following steps until you detect a link:
		1. Ensure that the cable is connected between the switch port and your storage system interface, and that both ends are securely attached.
		2. Ensure that the switch port and interface are both configured to the up status, and one of the following is true:
		 Autonegotiation is enabled on both sides Autonegotiation is disabled on both sides, and the duplex and speed settings match
		3. Because the switch port, cable, or NIC might be faulty, replace them, one by one, to locate the fault.
		4. If the problem persists, contact your technical support.
203	No link is detected because of a speed mismatch.	Change the interface configuration or peer switch port configuration to match the speed.
204	The interface is not configured to the up status.	Configure the interface state to the up status.
205	Duplex mismatch.	Change the interface or peer switch port duplex setting so that they match.
206	Link capacity problem.	Upgrade to a faster interface.

Error code	Description	Recommended actions
207	207 The interface is not	Complete the following steps:
receiving.	receiving.	1. Pull the network cable out from the network interface card.
		2. Reinsert the cable.
		3. Use ifstat to display statistics.
		• 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. Reset the interface card.
		2. Check the cables.
		3. If the problem persists, contact your technical support.
209	Excessive unsupported protocol packets are being sent to your	The problem is not with your storage system. Contact your network administrator to resolve the problem.
201	The Declaration of the second se	
301	The IP address and the netmask are inconsistent with the assigned broadcast address.	Change the configuration by using the if config 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.

Error code	Description	Recommended actions
402	There is excessive packet loss on the sending side.	The problem is not with your storage system. Examine the network and the client for congestion.
403	There is excessive packet loss on the receiving side.	The problem is not with your storage system. Examine the network and the client for congestion.
404	The average TCP packet size is poor on the receiving side because the network, client, or both are not enabled to support jumbo frames.	The problem is not with your storage system. Enable support for jumbo frames in network devices and the client.
405	The average TCP packet size is poor on the receiving side because of a problem with the network, client, or both.	The problem is not with your storage system. Examine the network and client for configured MTUs.
406	The average TCP packet size is poor on the receiving side because of a client application problem.	The problem is not with your storage system. Examine the client application data transmission strategy.
407	Excessive TCP listen socket drops because the system is overloaded or under security attack.	Contact your network administrator to resolve the problem.
408	There are excessive filtered TCP port drops because the system is under security attack.	Check your network. Contact your network administrator to resolve the problem.

Error code	Description	Recommended actions
409	There are excessive embryonic TCP connection drops because the system is under security attack or because a client has a bug.	A packet trace might assist in locating the problem. Contact your network administrator to resolve the problem.
410	Excessive TCP checksum errors. These errors can be caused by bad hardware on the client, in the network infrastructure (for example, blade in switch or router), or on the NIC. These errors can also be caused by a bug in the client.	 Check your client system for bugs. Replace hardware components until the problem is resolved. Contact your network administrator to resolve the problem.
411	There are packets because of a client. Your system might be under a security attack.	The problem is not with your storage system.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

/etc/gateways file 61 /etc/hosts file about 71 adding, host name 72 changing host name 73 creating, from NIS master 83 editing, with FilerView 73 hard limits 73 host-name resolution 71 /etc/nsswitch.conf file 71, 90, 91 /etc/resolv.conf file 75, 77, 79 /etc/resolv.conf file, hard limits 76 /etc/services file 199, 203 /etc/syslog.conf file 211

10 Gigabit Ethernet interface statistics 187–189

A

A record 76 AAAA record 76 address autoconfiguration 35 address resolution 35 alias address creating 49 deleting 49 anycast address 31

B

Baseboard Management Controller (BMC) 27 blocking protocols 51 BMC how to configure 28 managing with Data ONTAP commands 28

С

CDP configuring hold time 134 configuring periodicity 135 Data ONTAP support 133 disabling 134

enabling 134 online migration 133 viewing neighbor information 137 viewing statistics 135 CDP (Cisco Discovery Protocol) 133 certificate authentication about 158 adding a signed certificate 166 configuring 161 enabling on a storage system 168 enabling, on a Windows client 168 installing 165 installing, root certificates 167 root certificate 161, 167 viewing, subset of root certificates 167 certificate authority non-Windows 164 Windows 2000 162, 163 CIFS (Common Internet File Service) 209 Cisco Discovery Protocol (CDP) 133 commands dns flush 77 dns info 77 ifconfig 39, 40, 44, 49, 50, 63, 100 ifconfig -a 52 ifstat 44, 52, 54, 187-189 ipsec 170 ipsec cert set 167, 168 ipsec cert show 167 ipsec policy add 170 ipsec policy delete 172 ipsec policy show 172 ipsec sa show 175 ipsec stats 173 keymgr install cert 166 ndp 177 netdiag 177, 178, 213 netdiag -s 179 netstat 52, 53 netstat -p icmp 181 netstat -rn 66, 67 nis info 88 ping 177 ping6 177 pktt 177 route 61, 62, 68

route -s 66 route add 63 routed 64 routed status 67 snmp 143, 144 snmp authtrap 144 snmp community 144 snmp contact 144 snmp init 144 snmp location 144 snmp traphost 144 snmp traps 139, 144, 149, 151, 152 snmp traps load 150 snmpbulkget 146, 147 snmpbulkwalk 146, 147 snmpwalk 142, 146, 147 sysconfig 25 traceroute6 177 useradmin group add 142 useradmin role add 142 useradmin user add 142 vfiler run 160 vif add 121 vif create 113 vif create lacp 119 vif create multi 118 vif create single 114 vif delete 121 vif destroy 125 vif favor 116 vif nofavor 117 vif stat 124 vif status 122, 123 vlan add 98, 101 vlan create 98 vlan delete 98, 102 vlan modify 98, 103 vlan stat 98, 104 Common Internet File Service (CIFS) 209

D

DAD (Duplicate Address Detection) 36 default route 63, 65, 67 default router list 62 DHCP 207 diagnose network problems 177 diagnostic tests 177, 179 DNS about 74

configuration information 77 configuring, from the command-line interface 75 configuring, with FilerView 89 disabling 75 disabling, dynamic updates 80 dynamic updates 78, 79 enabling 75 enabling, dynamic updates 80 fully qualified domain names (FODN) 77 host-name resolution 71, 74, 76 lookup 76 modifying dns.update.ttl 81 name cache 77 port used 206 time-to-live (TTL) 79 Domain Name System (DNS) 74 duplex settings, correcting mismatches 197 Duplicate Address Detection (DAD) 35, 36 dynamic DNS about 78, 79 disabling 80 disabling, for an IP address 80 enabling 80 in Data ONTAP 79 Dynamic Host Configuration Protocol (DHCP) 74, 207

Е

e0M 26 error messages error code, netdiag 213 networking 213 Ethernet frame jumbo frame 30

F

fast path about 59, 60 disabling 64 enabling 64 IPv4 60, 61 IPv6 60, 61 with asymmetric routing 59, 60 with NFS-over-UDP 59, 60 with ping utility 59, 60 with TCP 59, 60 with Telnet 59, 60 fats path differences between IPv4 and IPv6 60, 61

similarities between IPv4 and IPv6 60, 61 File Transfer Protocol (FTP) 204 FilerView changing host-name search order 91 configuring DNS 89 configuring NIS 89 configuring SNMP 144 network interface settings 50 network report 57 routed daemon 65 SNMP traps 149 viewing routing table 68 viewing, network interface statistics 57 flow control about 30, 44 options 44 frame about 29 characteristics 29 Ethernet 29 flow 30 frame size 29 jumbo 29 jumbo frame 29 MTU size 29 Pause Off 30 Pause On 30 **FTP 204** fully qualified domain names (FQDN) 77

G

GARP (Generic Attribute Registration Protocol) 95 GARP VLAN Registration Protocol (GVRP) 95 Generic Attribute Registration Protocol (GARP) 95 Gigabit Ethernet controller LINK INFO statistics 183, 185, 186 RECEIVE statistics 183, 185, 186 statistics 183, 185, 186 TRANSMIT statistics 183, 185, 186 Gigabit Ethernet interface interface statistics 187–189 statistics 193–195 GVRP 97 GVRP (GARP VLAN Registration Protocol) 95

H

host identification 199

naming 23, 24 host name about 23, 24 adding, in /etc/hosts file 72 changing 73 changing search order 91 resolution 90 resolution, with /etc/hosts file 71 resolution, with DNS 74 resolution, with NIS 81, 85 search order 90 host-name resolution about 71.91 FilerView 91 using /etc/hosts file 71 using DNS 74, 76 using NIS 81, 85 **HTTP 207** Hypertext Transport Protocol (HTTP) 207

I

ICMP 180, 181 ICMP redirect messages 181 ICMP Router Discovery Protocol (IRDP) 61 IEEE 802.1Q standards 97 inter-switch link (ISL) 126 interfaces statistics for N3700 190-192 Internet Key Exchange (IKE) 158 IP address alias 49 broadcast 42 configuration 39 configuring 40 flow control 44 media type 43 MTU size 43 partner interface 46 partner IP 45 prefix length 42 removing 47 subnet mask 41 IP ports 199 **IPsec** Windows authentication 169 about 157 active/active 160 adding a signed certificate 166 anti-replay service 173

authentication methods 159 certificate authentication 161, 168 certificate for Windows client 165 configuring preshared keys 169 disabling 170 enabling 170 implementing 159 installing a Windows 2000 certificate 163 installing root certificates 166 installing, root certificates 167 Kerboros support 169 key exchange 161 key exchanges 158 requesting a signed certificate 162, 164 restrictions 159 security associations 175 security policies 161, 170 security policy 158 Security Policy Database (SPD) 158 set up 161 specifying, root certificates 167 subset of root certificates 167 verifying configuration 173 vFiler configuration 160 viewing statistics 173 IPv6 address autoconfiguration 34 address scopes 32 address states 32 address types 31 configure addresses 31 disabling 33 dual stack mechanism 33 dynamic routing 62 enabling 33 Router Advertisement 63 stateless address autoconfiguration 34 support in Data ONTAP 31

J

jumbo frames advantages 29, 30 configuring 30 network requirements 30 size 29

K

Kerberos 158, 169, 208 Key Distribution Center (KDC) 158, 169

L

LACP (Link Aggregation Control Protocol) 110 LACP log file 119 Link Aggregation Control Protocol (LACP) 110 LINK INFO statistics 10 Gigabit Ethernet interface 187–189 Gigabit Ethernet interface 187–189, 193–195 load balancing IP address based 112 MAC address based 112 multimode vifs 112 port-based 112 round-robin 112 localhost 71

Μ

MIB /etc/mib/iscsi.mib 140 /etc/mib/netapp.mib 140 custom mib 140 iSCSI MIB 140 Microsoft Management Console (MMC) 168 multicast address 31 multimode vifs load balancing, IP address based 112 load balancing, MAC address based 112 load balancing, port-based 112 load balancing, round-robin 112

N

N3700 interfaces, statistics 190–192 NDMP (Network Data Management Protocol) 212 negotiated failover 46 Neighbor Discovery 35, 36 Neighbor Solicitation 48 neighbor unreachability detection 35 network connectivity discovering 133 Network Data Management Protocol (NDMP) 212 Network File System (NFS) 208 Network Information Service (NIS) 81 network interface 10 Gigabit Ethernet 23 10/100/1000 Ethernet 23 automatic takeover 46

blocking protocols 51 changing status 50 clearing statistics 54 configuration 39 configuring 39 dad attempts 48 down, status 50 flow control 44 Gigabit Ethernet 23 maximum number 25 modifying with FilerView 50 naming 23, 24 nfo 46 partner 46 statistics 183 statistics for T204E 187-189 statistics for T204V 187-189 statistics for T210 187-189 statistics for T320 187-189 trusted 44 types 23 unblocking protocols 51 untrusted 44 up, status 50 viewing context statistics 53 viewing settings 49 viewing statistics 54 viewing statistics, FilerView 57 network interfaces viewing statistics 52 next-hop determination 35 NFS port used 208 NIS about 81 administrative commands vpcat 84 ypgroup 84 vpmatch 84 yppush 82 ypwhich 84 configure 85 configuring, with FilerView 89 creating /etc/hosts file 83 disabling 85 enabling 85 enabling slave 87 host-name resolution 71, 81, 85 hosts map 81, 83 ipnodes map 81, 83

IPv6 support 81 master 83 netgroup cache 88 selecting the master server 83 slave 82.83 specifying domain name 86 specifying servers 86 statistics 88 viewing information 88 viewing, performance statistics 88 NIS (Network Information Service) 81 NIS slave about 82 enabling 87 guidelines 83 improve performance 82 **NNTP 202**

0

OID 140 options dns.cache.enable 77 dns.update.enable 80 dns.update.ttl 81 ip.fastpath.enable 64 ip.icmp_ignore_redirect.enable 63, 181 ip.ipsec.enable 170 ip.ping_throttle.drop_level 180, 181 ip.v6.enable 33 ip.v6.ra_enable 35 nis.domainname 86 nis.enable 85 nis.server 83 nis.servers 83, 86 nis.slave.enable 87 snmp.access 143 snmp.enable 142

P

parameter discovery 35 pause frame 30 performance, storage system 197 ping command 177 diagnose problems 180 throttling 180 throttling, disabling 181 throttling, threshold status 181 throttling, threshold value 180 port for SnapMirror 212 for SnapVault 212 NDMP 212 NFS 208 port usage 199 ports TCP 202 UDP 202 ports, IP 199 ports, NFS-enabled 202 prefix discovery 35 prefix list 62 preshared keys 158, 169

R

RECEIVE statistics 10 Gigabit Ethernet interface 187-189 Gigabit Ethernet interface 187-189, 193-195 on N3700 interfaces 190-192 redirect by routers 35 Remote LAN Module (RLM) 27 Remote Shell (RSH) 211 reverse lookup 76 RLM how to configure 27 managing with Data ONTAP commands 27 root certificate installing 166, 167 viewing, subset 167 route default 63 route metric 67 routed daemon about 61 disable 64 enable 64 port usage 211 turning off 62, 64 turning off, with FilerView 65 turning on 64 turning on, with FilerView 65 Router Advertisement 63 router advertisement (RA) 35 router discovery 35 router-advertised messages disabling 35 enabling 35

routing about 59 default route 63, 65, 67 fast path 59, 60, 64 FilerView 65 managing routing table 61 methods 59 modifying routing table 68 routed daemon 61, 62, 64 routing table 63, 65, 66 vFiler units 62 viewing with FilerView 68 routing information 67, 68 Routing Information Protocol (RIP) 61, 211 routing protocols 67, 68 routing table commands to manage 61 flags 67 IPv6 62 modify, circumstances 63 modifying 68 vFiler units 62 viewing 65, 66 viewing with FilerView 68 **RSH 211**

S

second-level vif guidelines for creating 126 Secure Shell (SSH) 204 Secure Sockets Layer (SSL) 209 security association (SA) 157 security associations about 157 key exchange 158 lifetime 160 security policy 158 viewing 175 security associations (SA) 157 security policy creating 170 deleting 172 viewing 172 Security Policy Database (SPD) 158 security policy options 171 services file 199 Simple Mail Transport Protocol (SMTP) 205 Simple Network Management Protocol (SNMP) 139, 210Simple Network Time Protocol (SNTP) 206

SMTP 205 SNMP about 139 access privileges, setting 143 agent 139, 140 agent, configure 140 authKey security 145 authNoPriv security 145 authProtocol security 145 commands 144, 146, 147 configuring group, v3 142 configuring role, v3 142 configuring users, v3 142 configuring, with FilerView 144 disabling 142 enabling 142 examples 146, 147 IPv6 support 139 login-snmp capability, v3 142 MIBs 139, 140 modifying configuration 143 modifying, with FilerView 144 noAuthNoPriv security 145 port usage 210 restricting access 143 security parameters 145 traps 140 traps, configuration file 150 traps, define 149 traps, examples 151 traps, guidelines for creating 148 traps, modify 149 traps, modifying 149 traps, modifying with FilerView 149 traps, parameter 155 traps, parameters 151–156 traps, types 139 traps, user-defined 148 traps, viewing 149 traps, viewing with FilerView 149 version 3 (SNMPv3) 139 viewing configuration 143 SNMP (Simple Network Management Protocol) 139 **SNMP** traps backoff-calculator parameter 155 backoff-multiplier parameter 156 backoff-step parameter 155 built-in 139 commands 151 configuring, in a file 150

creating 149 defining 150 edge-1 parameter 154 edge-1-direction parameter 154 edge-2 parameter 154 edge-2-direction parameter 154 example 151 guidelines 148 interval parameter 154 interval-offset parameter 154 loading 150 message parameter 156 modifying 149 modifying, with FilerView 149 parameters 151, 152 priority parameter 156 rate-interval parameter 155 trigger parameter 153 user-defined 139, 148 var parameter 153 viewing 149 viewing, with FilerView 149 SNMPv3 about 139 configuring group 142 configuring role 142 configuring users 142 example 146, 147 login-snmp capability 142 **SNTP 206** split-network condition 126 **SSH 204** SSL 209 statistics Gigabit Ethernet interface 193-195 on N3700 interfaces 190-192 syslog 211

Т

Telnet 205 TFTP 207 time service 206 time-to-live (TTL) 79, 81 TLS 209 TRANSMIT statistics 10 Gigabit Ethernet interface 187–189 Gigabit Ethernet interface 187–189, 193–195 LINK INFO statistics on N3700 interfaces 190–192

224 | Data ONTAP 7.3 Network Management Guide

on N3700 interfaces 190–192 Transport Layer Security version (TLS) 209 transport layer, diagnosing 178 Trivial File Transfer Protocol (TFTP) 207 TTCP 202

U

unblocking protocols 51 unicast address 31

V

vif

about 107 adding interfaces 121 creating single-mode 114 deleting interfaces 121 destroying 125 dynamic multimode 108, 110, 119 dynamic multimode, LACP log 119 failover, second-level 127 in an active/active configuration 128 LACP 119 load balancing 112 load balancing, IP address based 112 load balancing, MAC address based 112 manage 113 naming 23, 24

second-level 126, 128, 129 selecting preferred interface 116 single-mode 108, 109 specifying nonfavored interface 117 static multimode 108, 109, 118 status 123 types 108 viewing statistics 124 viewing status 122 vifs single-mode, failure scenarios 117 VLAN about 93 adding an interface 101 advantages 96 commands 98 configuring 97, 100 configuring GVRP 95 creating 98 deleting 102 GVRP 95 link-local address 101 membership 93, 94 modifying 103 naming 23, 24 prerequisites 97 tags 95 viewing statistics 104

IBM.®

NA 210-04777_A0, Printed in USA

GC52-1280-05

