Data ONTAP™ Software 5.3
System Administrator's Guide

Network Appliance, Inc.
2770 San Tomas Expressway
Santa Clara, CA 95051, USA
Tel: +1 408 367-3000
Fax: +1 408 367-3151
Support tel: (888) 4-NETAPP
Support: http://www.netapp.com/support/form.html
Information email: info@netapp.com
Web: http://now.netapp.com

Part number 210-02334
March 1999

Copyright and trademark information

Copyright information

Copyright © 1998 Network Appliance, Inc. All rights reserved. Printed in the U.S.A.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Trademark information

Network Appliance, NetCache, Data ONTAP, SmartSAN, SnapCopy, SnapMirror, SnapRestore, Snapshot, WAFL, Web Filer, FilerView, BareMetal, and SecureShare are trademarks of Network Appliance, Inc. FAServer, NetApp, and the Network Appliance logo are registered trademarks of Network Appliance, Inc.

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

Communications regulation information

Compliance

This equipment has been tested and found to comply with the limits for a Class A digital device pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference in which case the users will be required to correct the interference at their own expense.

Cables

To maintain compliance with applicable emission standards, use cables shielded with metallic RFI/EMI connector hoods on this device.

Canada notice

This digital apparatus does not exceed the Class A limits for radio noise emissions from digital apparatus set out in the radio interference regulations of the Canadian Department of Communications.

Avis Canadien

Le présent appareil numérique n'émet pas de bruits radioélectriques dépassant les limites applicables aux appareils numériques de la classe A prescrites dans le Règlement sur le brouillage radioélectrique édicté par le ministère des Communications du Canada.

BZT Hinweis

Hiermit wird bescheinigt, daß dieses Gerät in Übereinstimmung mit den Bestimmungen der BAPT-AmtsblVfg 243/1991 funkentstört ist. Der vorschriftsmäßige Betrieb mancher Geräte (z.B. Meßsender) kann allerdings gewissen Einschränkungen unterliegen. Beachten Sie deshalb die Hinweise in der Bedienungsanleitung.

Dem Bundesamt für Zulassungen in der Telekommunikation wurde das Inverkehrbringen dieses Gerätes angezeigt und die Berechtigung zur Überprüfung der Serie auf die Einhaltung der Bestimmungen eingeräumt.

Safety warnings

Failure to follow the directions contained in the following statements could result in bodily harm or loss of life.

- To reduce the risk of personal injury from hot surfaces, allow the internal components time to cool before touching or installing options.
- To reduce the risk of personal injury or damage to the equipment, ensure that the equipment is properly supported or braced when installing options.
- This equipment is designed for connection to a grounded outlet. The grounding type plug is an important safety feature. To avoid the risk of electrical shock or damage to the equipment, do not disable this feature.
- This equipment is furnished with one or more lithium batteries. There is a danger of explosion and risk of personal injury if the battery is incorrectly replaced or mistreated. Do not attempt to recharge the battery, disassemble it, or dispose of it in a fire.

Table of Contents

Preface xvii

Chapter 1 **Introducing Network Appliance Filers** 1

About filers 1

How you administer a filer 3

About filer main unit components 5

About disk shelves 7

About Data ONTAP 8

Filer administration with Data ONTAP 12

Chapter 2 **Filer Administration Basics**. 15

Overview 15

Using the administration host. 16

The root volume 18

Editing configuration files 19

Obtaining access to the filer shell 22

About multiple administrative users 24

Creating administrative users. 25

Deleting administrative users. 27

Listing administrative users 28

Changing an administrative user password. 29

Halting and rebooting the filer 31

Understanding the filer default configuration 33

Default exported and shared directories 34

The /etc/rc file 38

Naming conventions for network interfaces 43

Recovering from configuration errors 46

Core files 47

Message logging 49
Configuring filer options 52
The options command 53
The vol options command 54
Sending automatic email to Network Appliance 55
Filer system time synchronization 60
Time synchronization with the rdate command 61
Time synchronization with SNTP 63
Synchronizing filer system time 64
Using options command options to maintain filer security. 66
Software licenses 67

Chapter 3 **Disk and File System Management**. 71
Disk concepts 72
Volume concepts 80
Procedures for managing disks and volumes 84
Disk management tasks. 85
Volume management tasks 89
File statistics for volumes. 93
Getting a file statistics summary 96
Filestats command options 97
Volume reversion using SnapRestore 101
Reverting a volume to a selected snapshot 105

Chapter 4 **Network Administration** 109
About this chapter 109
Working with large files 110
Using SNMP 111
Host name resolution 117
Routing 123

Using ifconfig to configure an interface 127
EtherChannel trunking 132
Kinds of trunks 134
Hardware requirements for trunks 136
Virtual interfaces 137
Trunking virtual interfaces 138
Second-level virtual interfaces on a single filer 139
Second-level virtual interfaces in a cluster 140
Virtual interface management 142
Creating a single-mode trunk 143
Specifying a preferred link in a single-mode trunk. 144
Removing a link from preferred status in a single-mode trunk. 145
Creating a multiple-mode trunk 146
Creating a second-level virtual interface on a single filer 147
Creating second-level virtual interfaces in a cluster 148
Adding physical interfaces to a trunk. 150
Displaying the status of a trunk. 151
Displaying trunk statistics 152
Destroying a trunk 153
Database file protection. 154
Enabling and disabling database file protection with nvfail 156
Using the nvfail_rename file for additional database protection 157

Chapter 5 **ATM LANE and SPANS Administration** 159
Overview 159
Overview of ATM 160
ATM LANE 161
Configuring the ATM adapter for LANE 165
Getting started with LANE configuration 167
Configuring the LECS address 171

Configuring the ATM adapter for an ELAN 174
Checking and completing the ELAN configuration 177
FORE/IP over SPANS 183

Chapter 6 **File Sharing between NFS and CIFS Users** 189
About this chapter 189
File-locking interactions 191
Managing symbolic links for CIFS access 192
NFS and CIFS use of the read-only bit 195
Naming files used by both NFS and CIFS 197
Languages and character sets 199
Displaying a list of supported languages 205
Setting the console encoding 206
Setting the language of a volume 207
Creating a volume that uses a specified language 208
Displaying which volume uses which language 209
Directory formats 210
Directory conversion 211
Forcing Unicode default directory creation 212
Returning to pre-Data ONTAP 4.0 default directory creation 213
Converting directories to Unicode format upon any access 214
Returning to default Unicode format directory conversion 215
CIFS file name case. 216
Forcing CIFS file names to lowercase 217
Preserving the case of CIFS file names. 218
Directory conversion time 219
Speeding up conversion time by renaming NFS directories 220
How to manage UNIX access to NTFS files 221
How to manage the WAFL credential cache 222
Setting how long each WAFL credential cache entry is valid 225

Adding an entry to the WAFL credential cache 226
Deleting entries from the WAFL credential cache 227
Displaying WAFL credential cache statistics 228
Displaying a mapping result for a UNIX name. 230
Displaying a mapping result for a Windows name 231
Toggling CIFS login tracing 232

Chapter 7 **NFS Administration** 235
NFS administration covers four areas 235
Introducing the /etc/exports file 237
Rules for exporting volumes and directories 239
Default /etc/exports entries 242
Restricting access to volumes and directories 243
The exportfs command 245
The /etc/netgroup file 247
Exporting to subnets 249
About configuring a filer for (PC)NFS 251
Introducing the (PC)NFS daemon 252
Enabling the daemon 253
Setting up (PC)NFS user entries 254
Setting the default umask 256
About configuring a filer for WebNFS 258
Setting up WebNFS. 259
Managing WebNFS. 260
About displaying NFS statistics 262

Chapter 8 **CIFS Administration** 267
CIFS limitations 270
Limits on CIFS open files, sessions, and shares 271
Changing or viewing the filer’s description 273

Adding CIFS users to the filer 275
Adding local groups to the filer. 278
Using CIFS commands with a remote shell program 281
Enabling guest and generic access 282
Displaying a filer's shares 284
Creating and changing a share 287
Displaying information about shares 294
Deleting a share. 297
Creating a home share for each user 299
Assigning and changing access rights 302
Displaying access rights to an NTFS file. 307
Changing UNIX permissions and DOS attributes from Windows 308
Sending a message to all users on a filer 311
Event auditing 313
Event log detail displays 314
Event auditing overview 317
Enabling CIFS access logging 318
Disabling CIFS access logging 319
Specifying the active event log 320
Setting a system ACL on a file 321
Viewing events in a security log 325
Using oplocks 328
Displaying CIFS statistics 330
Displaying CIFS session information 331
Stopping and restarting CIFS sessions 333
Reconfiguring the filer for CIFS 337

Chapter 9 **HTTP Administration**. 339
Chapter contents 339
Starting HTTP service 340

Protecting Web pages with passwords 342
Using the HTTP virtual firewall 345
Using virtual hosting 346
Specifying MIME Content-Type values 348
Translating URLs 350
Displaying HTTP connection information 353
Displaying HTTP statistics 354

Chapter 10 **Cluster Administration** 355
Introducing clusters 355
Understanding cluster concepts 357
Configuring interfaces for takeover 361
Configuring dedicated and hot standby interfaces 365
Configuring interfaces with mirrored interfaces 367
Understanding takeover 369
Understanding giveback 372
Managing a cluster 374
Managing the cluster in normal mode 375
Managing takeover 377
Managing filers in takeover mode 379
Managing a virtual filer 381
Performing a giveback 386

Chapter 11 **Snapshots** 389
Understanding snapshots 390
Snapshot commands and options 394
Understanding snapshot disk consumption 399
Managing snapshot disk consumption 403
Accessing snapshots from clients 408

Chapter 12 **Qtree administration** 411
Using qtrees. 413
Qtree security styles 415
Qtree file access models 418
Creating a qtree 420
Modifying the security style of a qtree 421
Modifying qtree oplocks settings 422
Displaying qtree information 424

Chapter 13 **Quotas and Maximum Number of Files** 425
Restricting or tracking disk usage by using disk quotas 426
Increasing the maximum number of files. 437
The df command 438

Chapter 14 **Data Backup** 441
Introduction to data backup. 441
How the dump command works 443
Format of the backup data 447
How the dump command writes and stores data on tape 449
Determining the amount of backup data 451
Determining the number of tapes for the backup. 453
Prerequisites for the dump command. 455
Recommendations for performing a backup 457
The dump command syntax 460
Using the dump command to back up data to tape 463

Chapter 15 **Data Recovery** 471
Introduction to data recovery 471
The restore command syntax 473
Using the restore command. 475

Performing a full restore of a volume containing qtrees 478
Examples of the restore command 479
Restarting the restore command 484
How to use a filer tape drive to restore files to another system 486

Chapter 16 **Tape Device Management** 487
Introduction to tape device management 487
How the filer displays information about various tape drives 488
Displaying tape device information 489
Using the mt command to control tape devices 492

Chapter 17 **Network Data Management Protocol (NDMP)** 495
Introduction to NDMP 495
Tape devices that can be used by the filer's NDMP server 496
The ndmpd command. 498
Log messages generated by NDMP 500
Considerations when using NDMP on the filer 501

Chapter 18 **Volume Copy Using the vol copy Command Set** 503
About this chapter 503
Introduction to the filer's commands for copying volumes 504
Requirements and recommendation for copying a volume. 506
Details about copying one volume to another 509
Management of a volume copy operation when it is in progress. 514

Chapter 19 **Data Replication Using SnapMirror** 517
About this chapter 517
Purposes of SnapMirror 518
How SnapMirror works. 520
Snapshots created during data replication 522

How SnapMirror works with quotas 525
How SnapMirror works with the dump command 526
The /etc/snapmirror.allow file 527
The /etc/snapmirror.conf file 528
Replicating a volume 530
Disabling data replication for the entire filer 534
Resuming data replication for the entire filer. 535
Disabling data replication for one volume 536
Checking data replication status 538
Converting a mirror to a regular volume 540
Differences between the vol copy command and SnapMirror 542

Chapter 20 **System Information and Performance** 543
Displaying the Data ONTAP version. 544
Displaying filer configuration information 545
Displaying volume information 547
Displaying filer statistics 548
Displaying network statistics 550
Displaying interface statistics. 551
Improving filer performance 561

Chapter 21 **Troubleshooting** 563
Getting technical assistance 566
Booting from system boot diskette 567
Restarting a shut down filer 570
NVRAM problems 571
Volume problems 573
Failed mounts and stale file handles 574
Volume name problems 575
Disk problems 576

Disk failure without a hot spare disk 577
Disk failure with a hot spare disk. 578
Disk errors 579
Inconsistent file system. 581
Disk operations in Maintenance mode 582
Configuration problems. 583
How to reset the filer password. 585
How to initialize all disks and create a new file system 586
Network problems 587
NFS problems. 589
Windows access problems 591
UNIX cpio problems 595
UNIX df problems 596
Cluster messages 597
Serious error messages 603

Chapter 22 **SecureAdmin management** 605
SecureAdmin overview. 605
SecureAdmin concepts 606
Setting up and starting SecureAdmin. 609
Reinitializing SecureAdmin 610
Disabling SecureAdmin 612
Restarting SecureAdmin 613
Determining whether SecureAdmin is running. 614
Displaying the version of SecureAdmin 615

Chapter 23 **Detailed Options Information** 617
About options. 617
Autosupport options 618
CIFS options 620

DNS options 624
HTTP options 625
NFS and (PC)NFS options 627
NIS options 629
RAID options 630
Timed options. 631
Volume options 633
Miscellaneous options 635

Glossary 641

Index 649

Preface

Introduction

This guide describes how to configure, operate, and manage Network Appliance™ filers that run Data ONTAP™ 5.3 software. It covers all filer models.

Audience

This guide is for system administrators who are familiar with operating systems that run on the filer's clients, such as UNIX, Windows 95, and Windows NT. It also assumes that you are familiar with how the NFS, CIFS, and HTTP protocols are used for file sharing or transfers. This guide doesn't cover basic system or network administration topics, such as IP addressing, routing, and network topology; it emphasizes the characteristics of the Network Appliance filer.

Terminology and conventions

When the term "filer" is used, the information applies to all Network Appliance filer models.

You can enter filer commands on either the system console or from any client computer that can access the filer through `telnet`.

In examples that illustrate commands executed on a UNIX workstation, this guide uses the command syntax of SunOS 4.1.x. The command syntax and output might differ, depending on your version of UNIX.

This guide uses the term "type" to mean pressing one or more keys on the keyboard. It uses the term "enter" to mean pressing one or more keys and then pressing the Enter key.

Keyboard conventions

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

Typographic conventions

The following table describes typographic conventions used in this guide.

Convention	Type of information
Italic type	Words or characters that require special attention. File names. Placeholders for information 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. Man page names. Book titles in cross-references.
`Monospaced` font	Command and daemon names. Information displayed on the system console or other computer monitors. The contents of files.
`Bold monospaced` font	Words or characters you type. What you type is always shown in lowercase letters, unless you must type it in uppercase letters.

Special messages

This guide contains special messages that are described as follows:

Note

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

Caution

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

WARNING

A warning contains instructions that you must follow to avoid personal injury.

Introducing Network Appliance Filers 1

About filers

What a filer is

A filer is a hardware and software system. It acts on network requests from clients and processes them by writing data to or retrieving data from disks in disk shelves that are connected to it. The software that enables the filer to perform these tasks is the Data ONTAP™ operating system.

Components of a filer

A filer consists of the following major components listed in the following table.

Component	Function	To learn more, read...
Filer main unit	The piece of hardware that receives and sends data.	"About filer main unit components" on page 5
Disk shelves	Hardware that holds disks and is connected to the filer.	"About disk shelves" on page 7
Data ONTAP	Software that is the operating system for the filer.	"About Data ONTAP" on page 8

Filer and filer main unit

Often, the filer main unit is referred to simply as a filer.

What a filer does

A filer provides the following features.

Feature	Description
Network file service	The filer enables users on client workstations to create, delete, modify, and access files stored on it. Client workstations are connected to the filer through network connections.

Feature	Description
Multiprotocol file sharing	Clients can use the following protocols to access data on the filer: ◆ CIFS (Common Internet File System)—used by Windows clients. ◆ HTTP (HyperText Transmission Protocol)—used by the World Wide Web. ◆ NFS (Network File System)—used by UNIX systems. ◆ (PC)NFS (Personal Computer NFS)—used by PCs to access NFS. Files written using one protocol are accessible to clients of any protocol, provided that filer licenses and permissions allow it. For example, an NFS client can access a file created by a CIFS client, and a CIFS client can access a file created by an NFS client.
Data protection	The filer protects disk data in the following ways: ◆ Network transaction logging—the filer records network transactions in case of failures and reconstructs transactions on recovery. ◆ Disk redundancy—the filer reconstructs data disks in case of disk failure.
Autosupport	Mail notifications about filer problems is automatically sent to Network Appliance Technical Support.

How you administer a filer

You administer the filer using the Data ONTAP operating system. You can use the following methods of administering a filer:

- Command execution through the filer's command line
- Command execution through Windows NT
- Configuration file editing
- Command execution through FilerView

Command execution through the filer's command line

You use the filer's command line to execute all Data ONTAP administrative commands from the command line, with the exception of some Windows NT administrative commands.

You can access the filer's command line from

- A serial terminal connected to the Console port of the filer
- A `telnet` session to the filer

You can also access some commands through a remote shell program, such as the UNIX `rsh` utility.

Command execution through Windows NT

You use Windows NT commands to perform filer administrative tasks related to Windows NT.

You can execute Windows NT commands that affect the filer using native Windows NT administration tools such as Server Manager and User Manager.

Configuration file editing

You edit configuration files to supply information that Data ONTAP needs to perform certain tasks.

You can access configuration files by mounting the root directory of the filer on a UNIX client or by mapping the administrative share (C$) to a drive on a Windows client, then editing the file from the client.

Command execution through FilerView

You use FilerView to perform most administrative tasks from a Web-based interface.

You can use FilerView even if you did not purchase a license for the HTTP protocol.

About filer main unit components

Two kinds of components

The filer main unit has two kinds of components: internal components that enable it to function, and slots and ports that connect it to networks and disk shelves.

Internal filer components

The following internal components enable the filer to function.

Component	Description
System board	The system board is the main board of the filer. It has upgradable firmware. All components are connected to the system board.
System memory	System memory stores information temporarily.
NVRAM (Nonvolatile Random Access Memory)	Data ONTAP uses NVRAM to log network transactions as a data integrity measure. In case of a system or power failure, Data ONTAP uses the contents of NVRAM to restore network data to disk.
Diskette drive	In an emergency and for major upgrades, Data ONTAP uses the diskette drive to boot from diskettes.
LCD and LEDs	The filer displays status information on the LCD and LEDs.

Component	Description
Environmental adapter	The environmental adapter performs the following functions: ◆ Monitors the filer's temperature and fans ◆ Sends critical information to the filer's LCD ◆ Logs information ◆ Shuts down the filer if its temperature is beyond a critical range or the fans cease operating

Slots and ports

The filer has slots for external connections and ports for a console and diagnostic hardware.

Component	Description
Slots	The filer contains expansion slots for network interface cards, tape drive adapters, and disk shelf adapters.
Serial ports	The two serial ports are as follows: ◆ The console port connects to the filer a serial terminal that you can use as a console. ◆ The diagnostics port connects diagnostic equipment such as the environmental monitor unit (EMU) of a StorageShelf 2 storage shelf.

About disk shelves

Disk shelves contain disks

A disk shelf contains the disks that store the data that the filer serves.

Disk shelf environmental information

Disk shelves collect information about the presence of disks, fan status, power supply status, and temperature. Disk shelves send messages to the console if parameters exceed permissible operating conditions.

About Data ONTAP

Data ONTAP overview

You administer the filer using Data ONTAP commands. Data ONTAP uses only commands and files that are relevant to file service administration. Data ONTAP contains no extraneous utilities such as text editors.

Data ONTAP manages data in the following three ways

- Data storage management
- Data organization management
- Data access management

Data storage management

Data ONTAP stores data on disks in disk shelves. Disks are organized into RAID groups, and RAID groups are organized into volumes. These items are explained in the following table, along with what aspects of them you can administer.

Data Storage Concept	Explanation	Storage administrative Actions
RAID (Redundant Array of Independent Disks)	A feature that enables file access even if one disk in a RAID group is damaged.	
RAID group	A RAID group consists of a parity disk, up to 52 data disks before Data ONTAP 5.0 and 28 data disks in Data ONTAP 5.0 and later, and optional spare disks.	You can control the size of a RAID group. This enables you to customize backups and disk failure recovery.

Data Storage Concept	Explanation	Storage administrative Actions
Data disks	Hold the data that clients access.	You can control the number of data and spare disks that a filer can use. This enables you to manage recovery if a disk fails.
Parity disk	Contains information that Data ONTAP uses to reconstruct data if a data disk fails. Each RAID group has one parity disk.	
Spare disks	Replace failed data disks automatically with reconstructed data.	
Volumes	Pools of storage composed of multiple disks that store client data. A filer can support up to 23 volumes or can have only one volume.	You can specify the size of a volume, add disks to it, copy volumes, and control the number of RAID groups in a volume. This gives flexibility in managing data storage tasks such as backups and restores, and enables you to customize volumes.

Data organization management

Data is organized in files, which are the smallest unit of data management. Users organize files into directories, and you organize directories into file systems, which are known as volumes. You can also organize directories into special directories called qtrees. Major data organization concepts that are special to filers are explained in the following table.

Data Organization concept	Explanation	Data Organization Administrative Actions
volume	An independent file system.	You can specify the following features for a volume: ◆ A security style, which determines whether a volume can contain files that use UNIX security, NTFS security, or both type of files ◆ Whether it uses CIFS oplocks ◆ Disk space and file limits This enables you to customize volumes for the needs of your users.
qtree	A special subdirectory of the root directory of a volume. It has the following special attributes: ◆ A security style ◆ A CIFS oplocks setting ◆ Disk space and file limits	You can specify the following features for a qtree: ◆ A security style like that of volumes ◆ Whether it uses CIFS oplocks ◆ Disk space and file limits This enables you customize areas for projects and to keep users and projects from monopolizing resources.

Data access management

Data ONTAP manages access to data by performing the following operations:

- Checks file access permissions against file access requests.
- Checks write operations against file and disk usage quotas that you set.
- Takes snapshots and makes them available so that users can access deleted or overwritten files. Snapshots are read-only copies of the entire file system.

Data ONTAP enables you to perform the following actions:

- Administer network connections.
- Administer protocols.
- Configure pairs of filers into clusters for failover.
- Dump data to tape and restore it to the filer.
- Copy volumes
- Mirror volumes

Filer administration with Data ONTAP

Filer administration activities

Administering a filer involves the following activities:

- Configuring the filer
- Monitoring and maintaining client access to data
- Monitoring and maintaining network access
- Monitoring and maintaining filer hardware
- Performing periodic administration tasks

Configuring the filer

You configure the filer by running either the Setup Wizard or the `setup` program or both if necessary, then editing configuration files if necessary.

Monitoring and maintaining client access

You monitor client access to data on the filer by the following methods:

- Gathering network statistics so that you can verify and improve the performance of the filer
- Gathering file statistics so that you can schedule snapshots
- Monitoring accesses to a file by CIFS clients so that you can monitor potential security problems

You maintain client access by the following methods:

- Configuring volumes and qtrees to accommodate the needs of Windows and UNIX clients
- Configuring volumes so that there is enough disk space for data
- Configuring RAID groups to maintain maximum data availability
- Configuring snapshots so that data is recoverable in case of accidental deletions
- Setting quotas to make sure that users have enough resources for their work

Monitoring and maintaining network access

You monitor network access by gathering network statistics so that you can verify and improve the performance of the filer and its network interfaces.

You maintain network access by following general filer network administration procedures and procedures specific to your protocol. These procedures are described in later chapters in this guide.

Monitoring and maintaining filer hardware

You monitor hardware by gathering statistics about your hardware and analyzing them for performance.

You maintain the hardware as described in the appropriate hardware guides.

Periodic administration tasks

The following table lists essential tasks that you should consider performing periodically, suggests how often to perform a task, and lists the command or file to use for the task.

How Often	**Task**	Command or File
Daily	Review filer performance and CPU utilization.	`sysstat` command
	Review filer disk usage on a per-volume and per-snapshot level.	`df` command
	Check for drive failures, reboots, and other logged events	*/etc/messages* file
	Review CIFS utilization, if licensed for CIFS.	`cifs stat` command
	Review HTTP utilization, if licensed for HTTP.	`httpstat` command
	Review NFS utilization, if licensed for NFS.	`nfsstat` command

How Often	Task	Command or File
Weekly	Review filer disk usage on a per-user level.	`quota report` command
	Test connectivity to NT Domain Controller.	`cifs testdc` command
	Review current shared CIFS directories.	`cifs shares` command
	Review NFS exports	*/etc/exports* file
	Review filer network traffic statistics.	`netstat` command
	Review network interface performance statistics on a per-interface level.	`ifstat` command
Monthly	Make sure that the disk shelves are connected properly.	`shelfchk` command
	Review filer uptime.	`uptime` command
	Review filer OS revision.	`version` command
	Review current NIS server.	`ypwhich` command

Filer Administration Basics 2

Overview

About this chapter

This chapter describes routine filer administration procedures that you need regardless of the file-sharing protocols licensed for your filer.

This chapter emphasizes the filer characteristics that distinguish the filer from a general-purpose server.

Other protocol-independent administration procedures

Use the following table to locate procedures for managing other filer features that are independent of file-sharing protocols.

Read...	If you want to manage...
Chapter 3, "Disk and File System Management"	Disks, file systems (volumes), or RAID groups.
Chapter 11, "Snapshots"	Snapshots.
Chapter 13, "Quotas and Maximum Number of Files	Quotas or the maximum number of files the filer stores.
Chapter 14, "Data Backup"	Backups.

Using the administration host

About the administration host

The filer recognizes a single client computer as the administration host. The administrator who set up the filer specified the name of the administration host using the `setup` program.

Administration host privileges

The filer granted root permissions to the administration host after the `setup` procedure was completed. The following table describes the administration host's privileges.

If the administration host is...	You can...
an NFS client	◆ Mount the filer root directory and edit configuration files from the administration host. ◆ Enter filer commands by using a remote shell program such as `rsh`.
a CIFS client	Edit configuration files from any CIFS client as long as you connect to the filer as root or "Administrator."

Administration host entry in the /etc/hosts.equiv file

The `setup` procedure placed the administration host name in the */etc/hosts.equiv* file automatically.

Administration host as the mail host

The administration host is the mail host for the filer. The filer configures the administration host to send periodic email messages to Network Appliance Technical Support about the status of your filer automatically.

For the administration host to send email, it must provide a server for the SMTP protocol, such as the `sendmail` program, or the Microsoft Exchange server.

Designating a different mail host

You can designate another host at your site to be the mail host at any time. Refer to "Sending automatic email to Network Appliance" on page 55 for more information about how the filer sends automatic email messages. Refer to "Use the options command to configure autosupport" on page 57 for information about how to specify a different mail host.

Requirements for using an NFS client as the administration host

If you plan to use an NFS client to manage the filer, the client must

- support a text editor that can display and edit text files containing lines ending with the newline character
- support the `telnet` and `rsh` commands
- be able to mount directories using the NFS protocol

Requirements for using a CIFS client as the administration host

If you plan to use a CIFS client to manage the filer, the client must support the `telnet` and `rsh` commands.

The root volume

About the root volume

Every filer has a root volume. It is the volume from which the filer reads configuration files. During `setup`, the filer creates a default volume, named *vol0*, and designates it as the root volume.

Designating the root volume

If you add volumes to your filer, you choose a volume to be the root volume during the multivolume configuration process.

For more information about designating the multiple volume configuration process, read Chapter 3, “Disk and File System Management,” on page 71.

About the volume name prefix

Volume names begin with the following prefix:

/vol/

Syntax to refer to the root volume from NFS clients

To refer to the root volume when mounting the root volume to an NFS client, use the following syntax:

/vol/vol0

Editing configuration files

What editor to use

The filer does not include a local editor. You must use an editor from a client to change filer configuration files.

Where configuration files reside

Configuration files reside in the */etc* directory in the filer's root volume.

Choosing an NFS or a CIFS client

The following table describes which procedure to use to edit configuration files from an NFS client or a CIFS client.

If the client is...	Read...
an NFS client	"Editing files from an NFS client" on page 19
a CIFS client	"Editing files from a CIFS client" on page 20

The procedure for modifying configuration files is different depending on whether you edit the files from an NFS client or a CIFS client. If you use an NFS client, edit the files from the administration host. If you use a CIFS client, connect to the filer as Administrator.

Editing files from an NFS client

The following table describes how to edit configuration files from an NFS client.

Step	Action	
1	**If the NFS client is...**	**Then...**
	the administration host	Mount the filer root volume to the host.
	not the administration host	1. Mount the filer root volume to the administration host. 2. From the administration host, edit the */etc/exports* file on the root volume to grant root permission to the client. 3. Use the filer console, a `telnet` client, or the `rsh` command to issue the following command to the filer: **`exportfs`** 4. Mount the filer root volume to the client.
2	From the client, use a text editor to edit the files in the */etc* directory.	

Editing files from a CIFS client

The following table describes how to edit configuration files from a CIFS client.

Step	Action
1	Connect from a CIFS client to the filer as Administrator. After `setup` finished, the default */etc/passwd* and */etc/group* files on the root volume were set up to enable you to share files on the filer as Administrator.

Step	Action
2	Display the contents of the filer's *C$* share and select a file to edit. After `setup` finished, the filer root directory was shared automatically as *C$*. The Administrator has read, write, and execute rights to the share. The *C$* share is a "hidden" share; you can get to it only by specifying the path manually (for example, as *\\filer\C$*), rather than accessing it through the Network Neighborhood icon.

Obtaining access to the filer shell

Ways to access the command line

The filer supports a command-line interface. You can access the command line

- directly, from the system console
- remotely, using `telnet`
- remotely, using a remote shell such as `rsh`

Sharing a single telnet and console session

The console and `telnet` share a single session. Everything entered through `telnet` is echoed at the console; everything entered at the console is echoed to the `telnet` session

telnet session restriction

Only one `telnet` session can be open at a time.

Closing a telnet session

To close a `telnet` session, press Ctrl-] to log out of the filer, then press Ctrl-D to log out of `telnet`.

telnet and console password requirement

Although `telnet` and the console share the same shell session, `telnet` and the console each prompt you for a password. Both the console and `telnet` connections use the same password.

rsh support

The filer supports `rsh` with trusted remote hosts—those remote hosts listed in */etc/hosts.equiv* on the root volume.

In addition to entering `rsh` commands manually, you can use a shell script or `crontab` file to enter some commands.

Note

You can use `rsh` only to enter filer commands. You cannot use `rsh` to remotely log in to the filer. To log in to a filer remotely from a host, use `telnet`.

Commands accepted from rsh

The following table lists the filer commands that you can execute using `rsh`.

`cf`	`halt`	`quota`	`sysconfig`
`cifs`	`httpstat`	`raid`	`sysstat`
`date`	`mt`	`rdate`	`timezone`
`df`	`netstat`	`reboot`	`uptime`
`disk`	`nfsstat`	`restore`	`version`
`download`	`options`	`route`	`vif`
`dump`	`partner`	`snap`	`vol`
`exportfs`	`qtree`	`snmp`	`ypwhich`

Use Ctrl-C to abort the command that is running

Ctrl-C aborts whatever command is being run from the console or a `telnet` session.

Because the console and `telnet` share a single session, a command entered at the console or through `telnet` can be terminated inadvertently from either location.

To ensure that a command is not terminated by Ctrl-C, start the command from a trusted host through `rsh`.

Changing the system password

A system password is required to establish a console or `telnet` connection with the filer. The password was specified during `setup`.

You can change the system password at any time with the `passwd` command. When you enter the `passwd` command, the filer prompts you to enter the old password, if any, and then requests the new password twice.

Where to go to learn more about security

For information about ways to increase filer security in addition to password protection, use the options described in "Using options command options to maintain filer security" on page 66.

About multiple administrative users

What is an administrative user?

An administrative user is a named account that exists on a filer. Administrative users have the same privileges as root, but can have a different password than root.

Multiple administrative users increase filer security

Adding multiple administrative users to a filer means you no longer need to share the root password. Instead, administrative users access the filer, either locally or remotely, with a unique login name and password. The filer records each user name at login with a syslog message in */etc/messages*, to enable auditing.

Command to use to create administrative users

Use the `useradd` option of the `useradmin` command to create multiple administrative users on a filer.

Ways to access the filer using an administrative login name

Using an administrative login name, you can access a filer through any of the following tools:

- Filer console
- `telnet`
- `rsh`

Creating administrative users

Description

The `useradmin` command enables you to increase filer security by creating multiple administrative users, rather than sharing the root name and password among multiple administrators.

Prerequisites

You must be logged in as root or as an existing administrative user to use the command.

Restrictions

- The user name is case-sensitive. It cannot contain any of the following 15 characters:

Character	Character
* (asterisk)	< (less than sign)
\ (back slash)	| (pipe)
: (colon)	+ (plus sign)
, (comma)	? (question mark)
= (equal sign)	] (right bracket)
/ (forward slash)	; (semicolon)
> (greater than sign)	space
[(left bracket)	

- The password should contain at least six characters, including at least two alpha characters and one numeric or special character.

Steps to create a new administrative user using a console or telnet

To create a new administrative user using the filer console or a `telnet` session, complete the following steps.

Step	Action
1	Enter the following command: **`useradmin useradd`** ***`user_name`*** *user_name* is the new administrative user login name.
2	Enter a password for the new user when prompted.

Steps to create a new administrative user using rsh

To create a new administrative user using `rsh`, complete the following steps.

Step	Action
1	Enter the following command: **`useradmin useradd`** ***`user_name password`*** *user_name* is the new administrative user login name. *password* is the new password associated with the new user name.

Deleting administrative users

Description

Use the `userdel` option to delete administrative users created with the `useradd` option of the `useradmin` command. You can use the `userdel` option at any time to maintain filer security by keeping the administrative user list current.

Caution

You are not prompted for a password when deleting a user.

Step

To delete an administrative user, complete the following step.

Step	Action
1	Enter the following command: **`useradmin userdel`** ***`user_name`*** *user_name* is the login name you want to delete.

Listing administrative users

Description

Use the `userlist` option to get a list of administrative users on a filer. You can use the `userlist` option at any time to determine whether you need to add or delete administrative users.

Step

To list one or more administrative users using a console, `telnet` or `rsh`, complete the following step.

Step	Action
1	Enter the following command: **`useradmin userlist [username_list]`** *username_list* is a space-separated list of login names you want to display. **Note** If you do not specify any user names, all administrative users are listed.

Changing an administrative user password

Description

Use the `passwd` command from a console, `telnet`, or `rsh` to change your administrative password on a filer at any time. Changing your password at regular intervals can provide you with better filer security.

Restrictions

When you use `telnet` or a console, the `passwd` command prompts you for the user name if administrative users, other than root, exist.

The `passwd` command only accepts arguments when you use `rsh`.

Steps to change an administrative password using a console or telnet

To change an administrative user password using a console or `telnet`, complete the following steps.

Step	Action
1	Enter the `passwd` command.
2	When prompted, enter the administrative user name of the user whose password you want to change.
3	Enter the old password.
4	Enter the new password.

Step to change an administrative password using rsh

To change an administrative user password using `rsh`, complete the following step.

Step	Action
1	Enter the following command: **`passwd`** ***`old_password new_password username`*** **Note** If you do not use the third argument, the `passwd` command works as before, and expects you to enter the old and new password for root.

Halting and rebooting the filer

Data storage in NVRAM

The filer stores requests it receives in nonvolatile random-access memory (NVRAM). The use of NVRAM

- improves system performance
- prevents loss of data in case of a system or power failure

NVRAM event during orderly shutdown

The `halt` and `reboot` commands perform an orderly shutdown. During an orderly shutdown, the contents of NVRAM are flushed to disk.

Procedure to halt the filer

To halt the filer, enter the following command:

```
halt
```

The filer displays the following prompt:

```
ok
```

Procedure to boot the filer

Complete the following step to boot the filer.

Step	Action
1	Ensure that the **ok** prompt is displayed on the console.
2	Enter the following command: **boot**

Procedure to reboot the filer

You can halt and reboot the filer in a single operation by entering the following command:

```
reboot
```

Where the filer boots from

When the filer boots, it uses the boot diskette in its diskette drive, if there is one. Otherwise, the filer boots from its hard disk.

Use the halt command to avoid data loss

You should always execute the `halt` command before turning the filer Off for the following reasons:

- The `halt` command flushes all data from NVRAM to disk, eliminating a potential point of failure.
- The `halt` command avoids potential data loss on CIFS clients.

  If a CIFS client is disconnected from the filer, the users' applications are terminated and changes made to open files since the last save are lost.

Caution

Never interrupt CIFS service by halting the filer without giving advance warning to CIFS users.

Before turning the filer Off, use the `halt` or `cifs terminate` command to send a warning message to CIFS users. This method gives users an opportunity to save files and exit applications within the time period that you specified before the actual shutdown.

Note

Clients using Windows 95 or Windows for Workgroups can display the CIFS shutdown messages only when the clients' WinPopup program is configured to receive messages. The capability to display messages from the filer is built into Windows NT.

For more information

For more information about `cifs terminate`, refer to "Stopping and restarting CIFS sessions" on page 333.

Understanding the filer default configuration

About the default configuration

The default configuration of a filer depends on whether the filer is running NFS, CIFS, or both; and NIS.

Although the default configuration is usable for small sites, it is probably not secure enough for large sites or for sites connected to the Internet.

Subsequent sections in this chapter describe in greater detail how you might want to modify the default configuration to suit your needs.

Default exported and shared directories

Default directories created

When `setup` finishes, two default directories are made available for access by clients. The default directories are

- the root directory
- the */home* directory

Permissions for the default directories

The following table shows the permissions that are assigned to the default directories when `setup` finishes.

This directory...	From this client...	Has these permissions...
The root directory	NFS	◆ full permissions for the root user on the administration host ◆ no permissions for any other user or host ◆ no permissions
	CIFS	◆ read and write permissions to all files for the Administrator user when logged in to the filer using the root password ◆ no permissions for other users
The */home* directory	NFS	permissions associated with individual users and with groups through UNIX security database
	CIFS	permissions assigned by the filer administrator for the HOME$ share

Accessing the directories

The following table shows how to access the default directories. Replace *filer* with the host name of your filer.

To access...	From this client...	Do this...
The root directory	NFS	Mount */filer/vol/vol0*.
	CIFS	Map a drive to *\\filer\C$*.
The */home* directory	NFS	Mount */filer/vol/vol0/home*.
	CIFS	Map a drive to *\\filer\HOME$*. Or Use the Network Neighborhood icon to locate the filer and locate the *HOME$* share.

Contents of the etc directory

The root directory contains an *etc* directory in which the filer configuration files are stored. You can modify the configuration files from the administration host.

The following table describes the files in the *etc* directory. Note that some of the configuration files might not exist on your filer if you use the filer for CIFS or NFS only.

File name	Contents	File-sharing protocol
.cifs.cat	Domain information (only if the filer is a member of a domain)	CIFS only
cifsconfig.txt	CIFS commands that the filer used for configuration The filer maintains this file automatically whenever you enter a `cifs` command, which can cause changes you make manually to be lost. To edit this file, terminate CIFS service by using the `cifs terminate` command, edit and save the file, then reboot the filer.	CIFS only

File name	Contents	File-sharing protocol
exports	NFS export points	NFS only
hosts	Known hosts and their IP addresses	All
hosts.equiv	Trusted hosts and users for `rsh`	All
group	CIFS group names, GIDs (group identification numbers), and members' names Not used if you use NIS to authenticate groups.	CIFS only
netgroup	Network groups	NFS only
nsswitch.conf	The order in which the filer contacts name services	All
passwd	Users' names, UIDs (user identification numbers), and primary GIDs Not used if you use NIS to authenticate users.	CIFS and (PC)NFS
rc	Script of commands to be executed when the filer is initializing	All
serialnum	Filer serial number and license codes, which help Network Appliance Technical Support process automatic email sent by the `autosupport` daemon	All
shadow	Encrypted password strings and password aging information Not used if you use NIS to authenticate users.	CIFS and (PC)NFS

How the /home directory is used

How the */home* directory is used depends on the file-sharing protocol used by the client. The following table describes how NFS and CIFS clients use the directory.

For NFS clients	For CIFS clients
The */home* directory is exported with read and write permissions to all NFS clients, and with root access to the administration host. Clients can mount directories after the filer host name is added to the filer */etc/hosts* file or to the client's name server. The svtx bit (also known as the sticky bit) is set on the */home* directory to prevent users from deleting each other's files. This is a safe configuration for networks that use UNIX permissions to provide security for network-accessible files. You can clear the sticky bit with the `chmod` command.	The */home* directory can be shared by CIFS users as the home share immediately after `setup` finishes. Before users can read and write files in the home share, you must follow the procedures in "Creating a home share for each user" on page 299. By default, CIFS users can write and delete their own files in */home*. For more information about access rights, refer to "Assigning and changing access rights" on page 302.

The /etc/rc file

How the filer uses the /etc/rc file

The filer executes the commands in the */etc/rc* file on the root volume at boot time to configure the filer.

If your filer is licensed to run the CIFS protocol, the */etc/rc* file must be present at boot time for the CIFS protocol to be enabled. No CIFS-specific information is entered in */etc/rc* as a result of the `setup` procedure.

All the commands in the */etc/rc* file are executable from the command line—there are no commands that are restricted to being executed from within the */etc/rc* file.

Procedure for editing the /etc/rc file

The following table describes how to make changes in the filer configuration that take effect every time the filer is booted.

Step	Action
1	Make a backup copy of the */etc/rc* file.
2	Edit the */etc/rc* file. **Note** Do not add CIFS commands in */etc/rc*. See "Editing configuration files" on page 19 for instructions about editing files from NFS and CIFS clients.
3	Save the edited file.
4	Reboot the filer to test the new configuration. If the new configuration does not work as desired, repeat Steps 2 through 4.

Default /etc/rc file contents

The best way to understand the commands used in */etc/rc* on the root volume is to examine the following sample */etc/rc* file:

```
#Auto-generated /etc/rc Fri May 30 14:51:36 PST 1997
hostname toaster
ifconfig e0 ‘hostname‘-0
ifconfig e1 ‘hostname‘-1
ifconfig fi0 ‘hostname‘-fi0
ifconfig a5 ‘hostname‘-a5
route add default MyRouterBox
routed on
options dns.domainname netapp.com
options dns.enable on
options nis.domainname netapp.com
options nis.enable on
timezone US/Pacific
savecore
exportfs -a
nfs on
```

Explanation of default /etc/rc contents

The following table explains the commands in the sample */etc/rc* file.

Command	Explanation
`hostname toaster`	Sets the filer’s host name.

Command	Explanation

`ifconfig e0 ‘hostname‘-0`
`ifconfig e1 ‘hostname‘-1`
`ifconfig f0 ‘hostname‘-f0`
`ifconfig a5 ‘hostname‘-a5`
	Sets the IP address for the filer’s Ethernet, Gigabit Ethernet, FDDI (Fiber Distributed Data Interface), and ATM (Asynchronous Transport Mode) interfaces with a default network mask. The arguments in single backquotes expand to “toaster” if you specify “toaster” as the host name during `setup`. The actual IP addresses are obtained from the */etc/hosts* file on the root volume; you might prefer to enter IP addresses directly in */etc/rc* on the root volume. If the specified network interface is not present, `ifconfig` issues an error message and has no other effect.
	
	The actual interface names and numbers depend on the filer model. Refer to “Naming conventions for network interfaces” on page 43 for more information about interface names.
	
	If you change the filer’s host name, you must modify the */etc/hosts* file on the root volume to substitute the new host name. If you don’t, `ifconfig` fails.
	
	To override the default network mask, explicitly specify the network mask in the `ifconfig` command after the host name; for example:
	`ifconfig e1 ‘hostname‘-1 netmask 255.255.0.0`
	`route add default LocalRouter 1`
	
	The preceding command specifies the default router. You can add route commands to */etc/rc* on the root volume to set static routes for the filer. The network address for LocalRouter must be in */etc/hosts* on the root volume.

`routed on`	Starts the routing daemon. See “Routing” on page 123 for more information about routing.
`options dns.domainname netapp.com` `options dns.enable on`	These options set the DNS domain.

Command	Explanation
`options nis.domainname netapp.com` `options nis.enable on`	These options set the NIS domain name and enable NIS.
`timezone US/Pacific`	Sets the time zone. The argument to the `timezone` command specifies which file in the */etc/zoneinfo* directory on the root volume describes the time zone you want.
`savecore`	Saves the core file from a system panic, if any, in the */etc/crash* directory on the root volume. Core files are created only during the first boot after a system panic.
`exportfs -a`	Exports all directories specified in the */etc/exports* file on the root volume. This command is included only if the filer runs NFS.

`nfs on`	Turns on NFS file service. This command is included even if NFS is not licensed for your filer. When the filer runs this command from */etc/rc* on the root volume on a filer without an NFS license, the command fails and the following messages appear:
	
	`NFS service is not licensed.`
	`(Use the "license" command to license it.)`
	`NFS server is NOT running.`

Changing SNMP commands in /etc/rc

Regardless of whether SNMP is enabled at your site, add the following commands to */etc/rc* on the root volume:

`snmp contact "`*`your email and telephone number`*`"`

`snmp location "`*`location of your filer`*`"`

`snmp init 1`

For example:

```
snmp contact "jdoe@abc.com 415-555-1212"
snmp location "ABC corporation, engineering lab"
snmp init 1
```

The information provided by these commands enables Network Appliance Technical Support to contact you after receiving automatic email messages about your filer from your site. For information about automatic email messages, refer to “Sending automatic email to Network Appliance” on page 55.

Naming conventions for network interfaces

Interface types the filer supports

The filer supports the following interface types:

- Ethernet
- Gigabit Ethernet
- FDDI
- ATM
- Virtual

How interfaces are numbered

Interface numbers are assigned based on the slot in which the interface card is installed. For more information about how network interfaces are numbered, refer to the hardware guide for your filer model.

How multiple ports are identified

Some Ethernet interface cards support four ports. These cards are referred to as quad-port interfaces. The filer uses a letter to refer to each port on a quad-port interface. The following table shows the relationship of port numbers to letters.

Port number	Letter
1	a
2	b
3	c
4	d

How interfaces are named

The following table shows how interfaces other than virtual interfaces are named, and how interface names are combined with card slot numbers and port letters to make their names unique.

Interface type	Letter used in name	Examples of names
Ethernet (single)	e	e0 e1
Ethernet (quad-port)	e	e0a e0b e0c e0d e1a e1b
FDDI	f	f0 f1
ATM	a	a0 a1

Virtual interface names

The name of a virtual interface is a string that is no longer than 15 characters that meets the following criteria:

- It must begin with a letter.
- It must not contain a space.
- It must not already be in use for a virtual interface.

Virtual interface names are case-sensitive.

About using interface names in scripts

When you write or modify shell scripts that involve interface names, remember the interface naming conventions that the filer uses. Because the slot in which an interface card is installed is a part of the interface name, different filers might have different interface names.

Filer host names

The first time the setup program runs, the filer creates a host name for each installed interface by appending the interface name to the host name.

Host name examples

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

Interface	Host name
Single-port Ethernet card in slot 0	toaster-0
Quad-port Ethernet card in slot 1	toaster-1a toaster-1b toaster-1c toaster-1d

Example 2: A filer named toaster with an FDDI interface in slot 1 would have the host name `toaster-f1`.

Recovering from configuration errors

Reasons to follow a special recovery procedure

Certain configuration errors can require you to follow a special recovery procedure because

- The filer does not have a local text editor.
- Problems with interface configuration can make the filer inaccessible to clients from which the */etc/rc* file can be edited.

Procedure when the filer does not boot

If configuration errors prevent the filer from booting from the hard disk, try booting from the diskette. For more information, see "Booting from system boot diskette" on page 567.

Procedure when administration host cannot access the filer

If your filer becomes inaccessible from the administration host after you change the */etc/rc* file, use the procedure in the following table to recover from the error.

Step	Action
1	Enter commands on the console to configure the interface with the correct address.
2	Enter the `exportfs` command to export the filer root directory to the administration host.
3	Edit the filer */etc/rc* file from the administration host.
4	Reboot the filer. If the changes do not correct the problem, repeat Steps 1 through 4.

Core files

About core files

When a hardware or software failure causes the filer to crash, the filer creates a core file that Network Appliance Technical Support can use to troubleshoot the problem.

Core file storage in /etc/crash

On the first boot after a system crash, the filer stores the core file in the */etc/crash* directory on the root volume.

What the savecore command does

The `savecore` command, which is included in the default */etc/rc* file on the root volume

- Produces a *core.n.nz* file. The *n* in the file name is a number. The string nz indicates that the file is compressed.
- Displays a message on the system console.
- Logs a message in */etc/messages* on the root volume.

Core dump space needed

A core dump file contains the contents of memory and NVRAM. Core dumps are written to a set of areas at the beginning of all the disks. The core dump area on each disk has a fixed size of approximately 20,447,232 bytes. Therefore, a filer or appliance with a large amount of memory can have an insufficient amount of core dump disk space to store a full core dump.

The following table shows how many disks are needed to store a full core dump for the amount of memory you might have in a specific filer or appliance.

Memory (in MB)	Disks Needed
32	2
64	4
96	5
128	7
160	9

Memory (in MB)	Disks Needed
192	10
224	12
256 or more	14

Automatic technical support notification

Your filer sends email automatically to Network Appliance Technical Support upon each system reboot.

Manual technical support notification

If you disabled automatic email, you should contact Network Appliance Technical Support when your system creates a core file.

Message logging

About message logging

The filer maintains messages in the */etc/messages* file on the root volume.

The level of information that the filer records in the */etc/messages* file is configurable.

About the syslogd daemon and the /etc/syslog.conf file

The message logging daemon, `syslogd`, uses the */etc/syslog.conf* configuration file on the filer root volume to determine how to log system messages.

You can configure `syslogd` to direct system messages to the console, to a file, or to a remote system based on their severity and origin.

By default, all system messages (except for those of severity level debug) are written to the console and to the */etc/messages* file on the root volume.

The /etc/syslog.conf file format

The */etc/syslog.conf* file consists of lines with two tab-separated (not space-separated) fields of the following form:

```
facility.level action
```

The facility parameter

The `facility` parameter specifies the subsystem from which the message originated. The following table describes the `facility` parameter keywords.

Keyword	Description
`auth`	messages from the authentication system, such as login
`cron`	messages from the internal `cron` facility
`daemon`	messages from filer daemons, such as `rshd`
`kern`	messages from the filer kernel
`*`	messages from all facilities

The level parameter

The `level` parameter describes the severity of the message. The following table describes the `level` parameter keywords arranged in order from highest to lowest severity.

Level	Description
`emerg`	panic condition that causes a disruption of normal service
`alert`	condition that you should correct immediately, such as a failed disk
`crit`	critical conditions, such as disk errors
`err`	errors, such as those caused by a bad configuration file
`warning`	conditions that might become errors if not corrected
`notice`	conditions that are not errors, but might require special handling
`info`	information, such as the hourly uptime message
`debug`	used for diagnostic purposes (suppressed by default)
`*`	specifies all levels

The action parameter

The `action` parameter specifies where messages should be sent. Messages for the specified level or higher are sent to the message destination. The following table describes the possible actions and gives examples of each action.

Action	Example
Send messages to a file specified by a path.	*/etc/messages*
Send messages to a host name preceded by an @ sign.	*@adminhost*
Send messages to the console.	*/dev/console* or *

Example line from /etc/syslog.conf

The following example causes all kernel messages of levels `emerg`, `alert`, `crit`, and `err` to be sent to the */etc/messages* file:

```
err.kern /etc/messages
```

The /etc/messages file restart schedule

Every Sunday at midnight, the */etc/messages* file is copied to */etc/messages.0*, the */etc/messages.0* file is copied to */etc/messages.1*, and so on. The system saves messages for up to six weeks.

Checking the /etc/messages file daily

Check the */etc/messages* file once a day for important messages. You can automate checking this file by creating a script on the administration host that periodically searches */etc/messages* and then alerts you.

Sample /etc/syslog.conf file

The following example shows a customized */etc/syslog.conf* file:

```
# Log anything of level info or higher to /etc/messages.
*.info                        /etc/messages

# Log all kernel messages, and anything of level err or
# higher to the console.
*.err;kern.*                  /dev/console

# Log all kernel messages and anything of level err or
# higher to a remote loghost system called adminhost.
*.err;kern.*                  @adminhost

# Log messages from the authentication system of level notice
# or higher to the /etc/secure.message file. This file has
# restricted access.
auth.notice                   /etc/secure.message
```

For more information

For more information about the `syslog.conf` file, see the *na_syslog.conf*(5) man page.

Configuring filer options

Commands to use to set options

The filer recognizes two commands, `options` and `vol options`, to set options.

"The options command" on page 53 describes the syntax for the options command.

"The vol options command" on page 54 describes the syntax for the vol options command.

The options command

What the options command does

When used interactively, the `options` command displays option values or changes the filer's behavior temporarily; the system returns to the state specified in the */etc/rc* file when rebooted. To make changes permanent, you must include the `options` commands in the */etc/rc* file.

Syntax of the options command

The syntax of the `options` command is as follows:

`options [ option [value] ] ...`

If you omit a value for an option, the command displays the current value of the option.

The following table describes the variables.

Variable	Description
option	the name of the option
value	the value of the option

Example of the options command

The following command enables the `pcnfsd` daemon:

```
options pcnfsd.enable on
```

The vol options command

vol options command configures volume-level behavior

You use the `vol options` command to configure volume-level behavior. You can use this command only with volume options, which are listed in "Volume options" on page 633.

Changes made with the `vol options` command are persistent between reboots—you do not need to add them to the /etc/rc file.

Syntax of the vol options command

The syntax of the `vol options` command is as follows:

`vol options` *`volname option`* [*`value`*]

If you omit a value for an option, the command displays the current value of the option.

The following table describes the variables.

Variable	Description
`volname`	the name of the volume that the option applies to
`option`	the name of the option
`value`	the value of the option

Example of the vol options command

The following command sets the maximum size of a RAID group in the volume named myvol to 12:

```
vol options myvol raidsize 12
```

Sending automatic email to Network Appliance

How automatic email messages are controlled

The filer uses the `autosupport` daemon to control how automatic email messages are sent from your site to Network Appliance Technical Support.

How the autosupport daemon works

The `autosupport` daemon is enabled by default on the filer. The daemon triggers automatic email messages to members of Network Appliance Technical Support, alerting them to potential filer problems.

Technical support response

If necessary, Network Appliance contacts you at the email address that you specify to help resolve a potential system problem.

Mail host requirement for autosupport

Because the filer doesn't function as a mail host, it relies on another host at your site that listens on the SMTP port (25) to send mail. Therefore, `autosupport` requires at least one host reachable by the filer that runs an SMTP server or a mail forwarder, such as the `sendmail` program or Microsoft Exchange server. By default, the administration host defined during `setup` is used as a mail host. You can specify more mail hosts.

About configuring autosupport

You can specify up to five addresses of email recipients. The default recipient is Network Appliance Technical Support. For receiving automatic messages, the email address of Network Appliance Technical Support is autosupport@netapp.com.

Refer to "Use the options command to configure autosupport" on page 57 for more information about specifying the email address and other options.

Events that trigger autosupport email

The mail host sends email to Network Appliance Technical Support about your filer after any of the events listed in the following table.

Event	Subject line of the email message
Low NVRAM lithium battery	`BATTERY_LOW!!!`
Disk failure	`DISK_FAIL!!!`
Disk scrub occurred	`DISK_SCRUB!!!`
Fan failure	`FAN_FAIL!!!`
Shutdown because of overheating	`OVER_TEMPERATURE_SHUTDOWN!!!`
Partial RPS failure[a]	`POWER_SUPPLY_DEGRADED!!!`
System reboot	`REBOOT`
Disk shelf error	`SHELF_FAULT!!!`
Spare disk failure	`SPARE FAIL!!!`
Weekly backup of */etc/messages*	`WEEKLY_LOG`
`option autosupport.doit` command	`THE STRING SPECIFIED IN OPTION autosupport.doit`
Successful cluster takeover of partner	`CLUSTER TAKEOVER COMPLETE`
Unsuccessful cluster takeover of partner	`CLUSTER TAKEOVER FAILED`
Cluster takeover of virtual filer	`REBOOT (CLUSTER TAKEOVER)`
Cluster giveback	`CLUSTER GIVEBACK COMPLETE`

a. This email message is sent only by a filer that uses a redundant power supply. It means that one of the two power supplies has ceased functioning.

Contents of automatic email messages

Each email message generated by `autosupport` contains the following types of information:

- date and time stamp of the message
- NetApp software version
- system ID of the filer

- host name of the filer
- software licenses enabled for the filer
- SNMP contact name and location (if specified in */etc/rc*)
- output of the following commands (some are applicable only to the licensed protocols):

  `sysconfig -v`

  `options`

  `ifconfig -a`

  `nfsstat -c`

  `cifs stat`

  `cifs sessions`

  `cifs shares`

  `httpstat`

  `df`

  `df -i`

  `snap sched`

  `sysconfig -r`

  `cf monitor`
- the following information, if clustering is licensed:
  - system ID of the cluster partner
  - host name of the cluster partner
  - cluster node status
- contents of */etc/messages*
- contents of */etc/serialnum*

Use the options command to configure autosupport

To change the default behavior of the `autosupport` daemon, use the `options` command. As with other filer commands, you can add `options` commands to the */etc/rc* file if you want to execute them automatically when the filer reboots.

Disabling or enabling the autosupport daemon

The `autosupport` daemon is enabled by default. The syntax of the command to disable or enable the daemon is as follows:

```
options autosupport.enable on|off
```

Example: `autosupport.enable off`

Specifying mail hosts: The command to specify hosts that send autosupport email messages to Network Appliance Technical Support is as follows:

`options autosupport.mailhost` *`hostname`*`,...`

You can specify up to five mail host names. Separate names by commas and do not include spaces in the list. The default is the administration host.

Example: `options autosupport.mailhost host1,host2,host3`

Specifying addresses for autosupport mail

The command for specifying the recipients of automatic email messages sent by the `autosupport` daemon is as follows:

`options autosupport.to` *`addresses`*`,...`

You can specify up to five email addresses. Separate email addresses by commas and do not include spaces in the list. The default is autosupport@netapp.com.

Example: `options autosupport.to autosupport@netapp.com,d@c.com`

Be sure to enter the command on a single line.

Specifying the filer administrator's address

Sometimes Network Appliance Technical Support might need to contact the administrator of the filer after receiving automatic email. The `options` command for specifying the filer administrator is as follows:

`options autosupport.from` *`address`*

Example: `options autosupport.from jdoe@abc.com`

If Network Appliance Technical Support needs to determine whether further action is necessary for solving or preventing filer problems, it sends email to jdoe@abc.com.

Sending an immediate message

Immediate messages contain the same filer data as automatic messages.

How to send immediate messages: The command to send an automatic email message immediately is as follows:

`options autosupport.doit string`

The string is used in the subject line of the email message to explain why the email was sent.

Example: **`options autosupport.doit TESTING`**

Sending a short message

The `options autosupport.noteto` command specifies the recipients of short email messages sent by `autosupport`. The short email messages are for urgent events, such as disk failures or filer reboots. The following example shows a short message:

```
Return-Path: <autosupport>
Received: from toaster by netapp.com (4.1/SMI-4.1)
id AA14370; Thu, 26 Sep 96 07:51:31 PDT
Message-Id: <9609261451.AA14370@netapp.com>
From: autosupport
To: jdoe
Date: Thu, 26 Sep 1996 07:51:27 -0700
Resent-Date: Thu, 26 Sep 1996 7:58:42 PDT
Subject: System Alert from toaster

REBOOT on Thu Sep 26 07:51:27 PDT 1996
```

Short messages are useful if the person who should be notified of urgent events reads email on a small screen, such as the screen on an alphanumeric pager.

Filer system time synchronization

Commands for synchronizing time

The filer uses two commands and a set of options to control and synchronize the system time.

- The `date` command sets the time locally.

  For example, the following command sets the system date and time to 9:25 a.m. on May 22, 1999:

  `date 199905220925`

- The `rdate` command synchronizes the filer's time with the time on another host with an accuracy of one second.
- SNTP (Simple Network Time Protocol) synchronizes the filer's time with the time on a time server with a theoretical maximum accuracy of one nanosecond; actual accuracy depends on the accuracy of the time server and network delays.

Time synchronization with the rdate command

The `rdate` command synchronizes the filer's time with a target computer's time using the UDP time service (typically, UDP port 37). This service is available on most UNIX computers. Check */etc/services* and */etc/inetd.conf* on the target computer to see whether this service is supported. For detailed information about the protocol that the `rdate` command uses, see RFC 868.

Note

There is currently no Windows 95 or Windows NT counterpart to the `rdate` command. If you want to use the `rdate` command, you must have a UNIX workstation that supports `rdate` on your filer's network.

When to use the rdate command

You use the `rdate` command when you need time accuracy only to the second, or when you want to keep your current method of synchronizing time.

Filer clock accuracy

To keep the filer's clock accurate, regularly run the `rdate` command on the filer with the target machine, a computer that maintains accurate time and that supports the port 37 UDP time service. For example, if the name of the target computer is time_node, enter

```
rdate time_node
```

Use of cron jobs to run rdate

A typical scheme is to have a UNIX computer run a periodic `cron` job that executes the appropriate `rdate` command on your filer through `rsh`. The computer and the user (if not root) running the `cron` job must be in the filer's */etc/hosts.equiv* file on the root volume.

cron job example

For example, on the UNIX computer named adminhost, the user named adminuser sets up a `cron` job to run every day at 3 a.m. It directs the filer named toaster to request a time update from the computer named time_node, which maintains an accurate time and supports the UDP time service.

crontab entry: The `crontab` entry on the UNIX system is as follows:

```
0 3 * * * rsh toaster -l root rdate time_node
```

Entry in /hosts/equiv file: The */etc/hosts.equiv* file on toaster must contain the following line:

```
adminhost adminuser
```

The adminhost and time_node host names must be known to the filer.

Time synchronization with SNTP

You control how time is synchronized using SNTP and whether time changes are logged by using the `timed` options. For detailed information about SNTP, see RFC 2030.

When to use SNTP

You use SNTP when you need a high degree of accuracy or when you are using clusters.

List of timed options

Before using SNTP, you set the `timed` options as described in "Synchronizing filer system time" on page 64 or accept the default settings. The `timed` options are listed in the following table. For details about the `timed` options, see "Timed options" on page 631.

Timed option	Function	Default
`timed.enable`	Determines whether the timed daemon runs.	Off
`timed.log`	Determines whether to log time changes initiated by the `timed` daemon to the console.	Off
`timed.max_skew`	Sets the maximum allowable discrepancy between filer time and server time.	30m
`timed.proto`	Selects whether to use the protocol used by the `rdate` command or SNTP.	ntp
`timed.sched`	Schedules when to synchronize the time with a time server.	hourly
`timed.servers`	Specifies up to five time servers in order of contact priority.	None

Synchronizing filer system time

Description

This procedure describes how to set the `timed` options to synchronize filer system time. You can do this any time, but if you have clusters, do it as soon as possible. After you complete this procedure, filer time is continually updated automatically.

Prerequisites

You must have available the name of at least one time server. You can get a list of NTP (Network Time Protocol) time servers, which SNTP can use, from *http://www.eecis.udel.edu/~mills/ntp/servers.htm*.

Steps

To set the `timed` options so that you can synchronize filer time, complete these steps.

Step	Action
1	Turn on the `timed` daemon by entering the following command: **`options timed.enable on`**
2	If desired, turn on console logging of time changes initiated by the `timed` daemon by entering the following command: **`options timed.log on`**
3	Specify the maximum discrepancy allowed before time is not synchronized by entering the following command: **`options timed.max_skew`** ***`nu`*** ***n*** is the number of units (specified by *u*). ***u*** is one of the following ◆ s for seconds ◆ m for minutes ◆ h for hours The default is 30m, which stands for 30 minutes.

Step	Action
4	Specify a protocol by entering the following command: **`options timed.proto`** ***`protocol`*** ***`protocol`*** is one of the following: ◆ rdate for the protocol used by the `rdate` command ◆ ntp for SNTP The default is ntp.
5	Specify a synchronization schedule by entering the following command: **`options timed.sched`** ***`schedule`*** ***`schedule`*** is one of the following: ◆ hourly to synchronize hourly ◆ multihourly to synchronize every six hours ◆ daily to synchronize every day at midnight ◆ a number followed by m to specify an interval of minutes or h to specify an interval of hours The default is hourly.
6	Specify up to five time servers by entering the following command: **`options timed.servers`** ***`servers`*** ***`servers`*** is a list of the host names or IP addresses of up to five time servers, separated by commas.

Using options command options to maintain filer security

What the options to the options command do

The following options to the `options` command help maintain filer security:

- The `telnet.hosts` option restricts `telnet` access to a limited number of hosts. Use this option to specify a comma-separated list of up to five hosts that can log in to the filer using `telnet`. Alternatively, you can disable `telnet` for all hosts by specifying a hyphen (-). By default, the option argument consists of an asterisk (*), which means all hosts have `telnet` access.

  A record of all `telnet` and console logins is maintained in the */etc/messages* file on the root volume.

- The `mount_rootonly` option restricts the mount privilege to root using privileged ports (ports 1 through 1,024). By default, the option is enabled. This option accepts the arguments `off` and `on` to disable and enable it. This option is applicable only if your filer runs NFS.

Note

Some PC clients and some older implementations of NFS on UNIX workstations use non-privileged ports to send requests. If you have these clients at your site, disable the `mount_rootonly` option or upgrade the client software.

- The `wafl.root_only_chown` option enables only root to change the owner of a file. When the option is disabled, the owner of a file can change its ownership without being root. By default, this option is enabled. Use the arguments `off` and `on` to disable and enable the option.

  When a non-root user changes the owner of a file, the set-user-id and set-group-id bits are cleared. If a non-root user tries to change the owner of a file but the change causes the file's recipient to exceed his or her quota, the attempt fails. This option is applicable only if your filer runs NFS.

Software licenses

About software licenses

The filer requires software licenses to enable these services and features:

- NFS
- CIFS
- HTTP
- clustering
- SnapMirror
- SnapRestore

Licenses are installed on the filer at the factory per your order, so the initial setup of your filer does not involve entering license codes.

You need to enter license codes only if any of the following conditions applies:

- You purchased a filer with a release earlier than Release 4.0, and you are upgrading it to Release 4.0 or later.
- You want to enable a service not previously licensed for your filer.
- The filer's file system becomes corrupt and must be rebuilt.

Network Appliance provides you with the appropriate license codes when shipping you the software upgrade kit or when giving you instructions for obtaining the software upgrade over the Internet.

Enabling services

To enter a license code to enable a protocol on your filer, use the `license` command with the following syntax:

`license` *`protocol=code`*

The *protocol* field can be one of these values: `nfs`, `cifs`, `http`, `cluster`, `snapmirror`, and `snaprestore`.

Example: If the license code for NFS is ABCDEFG, enter

```
license nfs=ABCDEFG
```

The events that take place after a license command depend on the protocol specified. The following table discusses the events for each service or feature.

Service or feature	Messages

NFS	`nfs license enabled.`
	`nfs enabled.`
	
	The filer also automatically runs the `nfs on` command to start NFS service. However, the filer does not add the `nfs on` command to */etc/rc* as a result of the `license` command. If you want the filer to run NFS service after each reboot, add `nfs on` to */etc/rc*.

CIFS	`CIFS license enabled.`
	`Run cifs setup to enable CIFS.`
	
	To start CIFS service, set up the filer's CIFS configuration by running `cifs setup`. You don't need to run `cifs setup` after the `license` command if you already set up the CIFS configuration.

HTTP	`http license enabled.`
	`Use "options httpd.enable" to enable http.`
	
	To start HTTP service, enter the following command:
	`options httpd.enable on`
	
	Results: The `options` command takes effect immediately. If you want the filer to automatically start HTTP service after each reboot, add `options httpd.enable` to */etc/rc*.

Cluster	`cluster site license enabled.`
	`  Clustered Failover will be enabled upon reboot.`
	`  Make sure that each individual service is licensed`
	`  on both nodes or on neither node. Remember to configure`
	`  the network interfaces for the other node.`

SnapMirror	`snapmirror license enabled` `  snapmirror enabled`
SnapRestore	`snaprestore license enabled` `  snaprestore is enabled`

Example: The following example shows how to activate several protocols with one command:

```
license nfs=ABCDEFG CIFS=HIJKLMN http=PQRSTUV cluster=AMPERVX
```

Displaying current license codes

To display licensing information, enter the `license` command without parameters, as follows:

```
toaster>license
nfs=ABCDEFG
cifs=not licensed
http=PQRSTUV
cluster=not licensed
snaprestore=not licensed
```

Disabling a license

To disable a license, enter DISABLE as the code for the protocol. For example, to disable your filer's license for NFS, enter

```
license nfs=DISABLE
```

After you disable a license, your filer stops service for the corresponding protocol. You can restart the service by reentering the license code.

Replacing license codes

If you misplace a license code, contact Network Appliance Technical Support to obtain a copy.

Disk and File System Management 3

Chapter contents

This chapter provides information about, and how to manage, the filer's disks and file systems. It is divided into the following sections.

Read...	To learn about...
"Disk concepts" on page 72	Principles of RAID groups and the different functions of disks in them, checking for media errors, using disks of various sizes, and handling disk failures.
"Volume concepts" on page 80	Volume concepts.
"Procedures for managing disks and volumes" on page 84	Managing disks and volumes on a filer.
"Volume reversion using SnapRestore" on page 101	Reverting the contents of a volume to those of a selected snapshot.

Disk concepts

Section contents

This section covers the following topics:

- Understanding RAID groups
- About disk addresses
- Use disk scrubbing to protect data from media errors
- Understanding hot spare disks
- Understanding hot swap
- Using disks of various sizes
- Understanding usable space on each disk
- Handling disk failures
- Effects of disk failure on filer operation

Understanding RAID groups

The filer uses RAID Level 4 to ensure data integrity even when one of the disks fails. The file system design, together with the support for RAID, optimizes filer performance and enables you to incrementally expand the filer's disk storage capacity.

In a RAID group, different disks have different functions. Most of the disks in the RAID group are *data disks*. One disk is the *parity disk*, which enables the filer to recover the data on a data disk if one fails.

Multiple RAID groups: The filer supports multiple RAID groups. The factory default filer configuration contains one RAID group. The filer supports up to 32 RAID groups. Each RAID group belongs to only one volume; you cannot assign more than one volume to the same RAID group.

Spare disks are used by RAID groups as needed. They do not have to be in the same disk shelf to be available to a RAID group. The filer automatically assigns disks to RAID groups and creates new RAID groups as each RAID group is filled with its maximum number of disks.

RAID group size: The following characteristics apply to RAID group size:

- The default number of disks in a RAID group (including the parity disk) is 14.
- A RAID group must contain at least two disks.
- The largest RAID group size you can create *manually* is 28 disks.

- The maximum RAID group size is 52 disks.

Occasionally, you might encounter situations in which you want to specify a RAID group size other than the default. For example, you might want to configure a filer with smaller RAID groups for the following reasons:

- Using smaller RAID groups reduces disk reconstruction time if a disk fails.
- Using smaller (therefore, more) RAID groups provides higher reliability by reducing the risk of data loss due to multiple-disk failure.

  Conversely, configuring with larger (therefore, fewer) RAID groups in a filer uses fewer disks for parity, leaving more disks available for data storage.

Parity disks: In each RAID group, the filer assigns the role of parity disk to the largest disk in the RAID group. After a data disk failure, the filer uses the parity disk in conjunction with the other data disks to reconstruct the failed disk's data and optionally write it to a hot spare disk. The parity disk must be at least as large as the largest data disk.

For more information about how the filer recovers from disk failures by using the parity disk or hot spare disk, refer to "Handling disk failures" on page 77.

About disk addresses

You identify a disk by its address, which is listed in the `HA.Disk_ID` column of the output of the `sysconfig -r` command. In this output listing, `HA` refers to the host adapter and `Disk_ID` refers to the disk ID number.

You use the disk address to

- interpret screen messages (for example, command output or error messages) that you see on your display
- quickly locate the disk that the message is referring to

SCSI disk addresses: With SCSI disks, the disk address is a combination of the disk's SCSI adapter number and the disk's SCSI ID number on the adapter. For example, disks attached to SCSI adapter 0 are numbered 0.0 through 0.6, disks attached to SCSI adapter 9a are numbered 9a.0 through 9a.6, and so on.

Fibre channel disk addresses: With fibre channel disks, the disk address is a combination of the disk's adapter number and the disk's fibre channel loop ID. To create the fibre channel loop ID, multiply the shelf ID switch value by eight and add it to the bay number. For example, `ha 8`, `shelf 1, disk 2` has disk ID `8.10`.

For more information: For more information about locating specific drives on disk shelves, refer to the hardware guide for the disk shelf.

Use disk scrubbing to protect data from media errors

The filer uses the RAID disk scrubbing procedure to increase data availability. The filer scans each disk in the RAID group for media errors. If the filer finds media errors, it fixes them by reconstructing the data from parity and rewriting the data.

Disk scrubbing reduces the chance of a multiple disk failure, caused by a disk media error encountered while the system was running in degraded mode.

The filer only scans a RAID group when all the group's disks are operational. Although disk scrubbing slows the filer somewhat, network clients might not notice the change in the filer's performance because disk scrubbing starts automatically at 1:00 A.M. on Sunday, when most systems are lightly loaded.

By default, disk scrubbing is enabled. You might want to disable scrubbing if you have a recurring problem that scrubbing encounters.

Example: For example, there might be an unrecoverable error on a disk that you cannot fix before the next disk scrub. The following commands disable and enable disk scrubbing:

```
options raid.scrub.enable off
options raid.scrub.enable on
```

Commands to start and stop disk scrubbing: You can also manually start and stop disk scrubbing regardless of the current value (On or Off) of the `raid.scrub.enable` option. Following are the commands for starting and stopping scrubbing manually:

```
disk scrub start
disk scrub stop
```

Sample messages logged from scrubbing: Messages from the disk scrubbing process are sent to the system error logging daemon.

Following are sample messages that can appear:

- If the filer finds an inconsistent parity block during scrubbing, it prints the following messages:

  `Inconsistent parity on volume` *`volume_name`*`, RAID group` *`n`*`, stripe #`*`n`*.
  `Rewriting bad parity block on volume` *`volume_name`*`, RAID group` *`n`*`, stripe #`*`n`*.

Note

An `Inconsistent parity` error message might indicate file system corruption. If you get such an error, contact Network Appliance Technical Support for assistance.

- If the filer finds a media error on the parity disk, it prints the following message:

  `Rewriting bad parity block on volume volume_name, RAID group n, stripe #n`

- If the filer finds a media error on a data disk, it prints the following message:

  `Rewriting bad block from parity on disk n, block n`

- If the filer finds more than one bad block, it prints the following message:

  `Multiple bad blocks found on volume volume_name, RAID group n, stripe #n`

The following sample messages appear after disk scrubbing is complete:

```
Scrub found n parity inconsistencies
Scrub found n media errors
Disk scrubbing finished...
```

Understanding hot spare disks

In addition to data disks and parity disks, the filer supports zero or more *hot spare* disks. A hot spare disk is not part of any RAID group and does not contain file system data. After a disk failure, the filer automatically rebuilds data (or parity) onto a hot spare disk, which then replaces the failed disk in the RAID group. This procedure avoids a system shutdown and returns the system to full performance.

A hot spare disk cannot replace a failed disk that is larger than itself. If you use only one hot spare disk in a filer, the hot spare disk must be as large as the largest file system disk. If you have multiple hot spare disks installed, the system uses the smallest hot spare disk needed to replace the failed disk.

Understanding hot swap

The filer enables you to *hot swap* disk drives. Hot swapping a disk drive means installing or removing it from the disk shelf while the filer is running, with minimal interruption to a file system. For example, you might want to hot swap a disk into a filer to replace a disk or to add a hot spare disk.

Using disks of various sizes

You can use disks of various sizes in the same filer. The following table shows a sample configuration.

Data	Data	Data	Data	Parity	Hot Spare
Disk 0 (4 GB)	Disk 1 (4 GB)	Disk 2 (4 GB)	Disk 3 (9 GB)	Disk 4 (9 GB)	Disk 5 (9 GB)

The following sections discuss the guidelines for choosing the size of a disk to add to the filer or to use as a replacement.

Replacing a failed disk: When you replace a failed disk, the physical space of the new disk must be equal to or larger than the *usable* space of the replaced one. The Used column of the `sysconfig -r` output shows the amount of usable space on each disk.

However, if the physical space of the new disk is larger than the usable space of the old disk, the system can use no more space than was available on the old disk. For example, if you replace a 4-GB disk with a 9-GB disk, the extra 5 GB on the new disk is not available to the file system. Therefore, it is most cost-effective to replace a disk with a disk of the same size.

Adding a disk to a file system: When you add to a file system a disk that is larger than the current parity disk, the new disk automatically becomes the parity disk. The disk used as the parity disk before the disk addition becomes a data disk.

Understanding usable space on each disk

A disk's *usable* space can be different from its *physical* space. The information here applies to data, parity, and hot spare disks:

- Disks from different manufacturers might differ slightly in size even though they belong to the same size category.
- To ensure that you can replace one disk with another in the same size category, a filer shipped with Data ONTAP 2.1 or later automatically "right-sizes" its disks. That is, it rounds off the disk capacity and divides disks into 2-GB, 4-GB, and 9-GB classes.
- If you upgraded your filer from a release before 2.1, the filer doesn't right-size the existing disks. Instead, it continues to use all the physical disk space.
- If a disk has never been right-sized, you can only replace it with another disk that has the same number of blocks or more.

The following table shows the usable disk space for each class of disk supported by the filer.

1-GB disks	2-GB disks	4-GB disks	9-GB disks
1,000 MB (2,048,000 blocks)	2,000 MB (4,096,000 blocks)	4,000 MB (8,192,000 blocks)	9,000 MB (17,612,800 blocks)

Although automatic disk right-sizing is not applied to existing disks in an upgraded system, it is applied to disks being added to the filer.

You can use `sysconfig -r` to compare the physical space and the usable space, and to determine whether disks are right-sized.

Handling disk failures

If one block on a data disk fails, the filer uses the parity disk in its RAID group to reconstruct the data on that block. The block is mapped to a new location on disk. If an entire data disk fails, the parity disk for that RAID group prevents any data loss and enables the filer to continue running.

Although the filer can continue to function with a failed disk, if it cannot reconstruct that failed disk on a hot spare, it automatically shuts down after 24 hours to encourage you to replace the failed disk. You can change the amount of time from 24 hours to another value using the `raid.timeout` option to the `options` command.

Effects of disk failure on filer operation

The effects of a disk failure on filer operation depend on whether the filer has a hot spare disk.

Without a hot spare disk: If the filer is not equipped with a hot spare disk, after a disk fails the filer enters a state called “degraded mode.” In this state, the RAID feature enable the filer to continue to run without losing data (although the filer’s performance is affected). Replace the failed disk as soon as possible because a second disk failure in the same RAID group causes the entire file system to be lost.

When a disk fails, the filer logs a warning message in the */etc/messages* file and to the system console every hour, notifying you of the number of hours before the system shuts down.

The shutdown ensures that you notice the disk failure. You can restart the filer without fixing the disk, but the filer continues to shut itself off at the specified intervals until you repair the problem.

By default, the filer shuts down after 24 hours. You can change this time interval using the `option raid.timeout` command; the argument is the time, in hours, that the system runs before automatic shutdown.

The system shuts down after the specified period if it is running in degraded mode. A filer is in degraded mode if either of the following conditions exist:

- One disk in any RAID group has failed.
- The batteries on the NVRAM card are low (if the filer is a PCI-based. system).

With a hot spare disk: If you reserve one or more disks as hot spare disks when you configure your filer, the filer also enters degraded mode after a disk failure. However, the filer immediately begins rebuilding the missing data in the background on the hot spare disk, with minimal interruption to file service.

The filer logs this activity in the */etc/messages* file and does not automatically shut down. If you turn off the filer while it is in degraded mode, it stops data reconstruction. After you turn the filer back on, the filer restarts the data reconstruction process from the beginning.

Except for a loss in performance while data is rebuilt on the hot spare disk, the failure of a single disk is transparent to the user. The filer exits degraded mode and returns to normal operation after it finishes reconstructing the data.

Network Appliance recommends that you replace the failed disk with a new hot spare disk after the filer finishes reconstructing data. This way, the filer continues to have a hot spare disk that it can use in case another disk fails.

The `sysconfig -r` command displays which disk is reserved as the hot spare disk. In addition to disk failure and hot spare disk replacement activity, the */etc/messages* file logs any failure in a periodic check of the hot spare disk.

Command to control RAID data reconstruction speed: You can control the speed of RAID data reconstruction by entering the following command:

`options raid.reconstruct_speed` *`speed`*

where *speed* is a number ranging from 1 (slowest) to 10 (fastest). Because RAID data reconstruction consumes CPU time, sometimes increasing the speed of data reconstruction slows the filer's network operations. The default speed is 4, which means approximately 40% of the CPU time is used for RAID data reconstruction.

When RAID data reconstruction is in progress, use the `sysstat` command to check the system load on the filer. If the load is light, increase the speed of RAID data reconstruction to maximize CPU utilization. For more information about the `sysstat` command, refer to "Displaying filer statistics" on page 548.

Volume concepts

Section contents

This section covers the following topics:

- Understanding volumes
- Determining the number of volumes to use
- Planning a multiple volume configuration
- Installing a foreign volume

Understanding volumes

Data on the filer is organized in *volumes*. A volume is an independent file system with its own RAID groups.

The initial configuration for new filers running Data ONTAP 5.0 or later includes one 2- to 14-disk volume (a root file system). All remaining disks are spares.

Volume naming conventions: You choose the volume names. The names must follow these naming conventions:

- begin with either a letter or an underscore (_)
- contain only letters, digits, and underscores
- contain no more than 255 characters

The root volume: Each filer must have a *root volume* to boot. The root volume of a filer (configured with either a single volume or multiple volumes) is the volume whose */etc* directory is used by the filer for configuration information.

The filer uses two naming conventions to indicate the root volume:

- */vol/vol0*
- /

In the */vol/vol0* convention

/vol	Indicates that the next part of the path, such as *vol0* in this example, is a volume name.
/vol0	Indicates the default name of the root volume for a filer initially installed with Data ONTAP 5.0 or later. You can change this name using the `vol rename` command.

For backward compatibility with filers running releases of Data ONTAP earlier than 5.0, you can use the / naming convention to indicate the root volume on a filer that you upgraded to Data ONTAP 5.1 that contains a previously existing file system.

Mounting volumes: On filers configured with multiple volumes, mounting / is equivalent to mounting */vol/vol0/*, (where *vol0* is the root directory of the root volume). Paths that begin with / (for example, */etc*) refer to directories on the root volume.

Note

/vol is not a directory—it is a special virtual root path under which the filer mounts other directories. You cannot mount */vol* to view all the volumes on the filer; you must mount each filer volume separately.

In mount requests and server commands, prefix the path names of the volumes and directories you want to mount using the convention */vol/volume_name/directory.*

Example: For example, */vol/users/home/cheryl* is a directory called */home/cheryl* in a volume named *users*.

Determining the number of volumes to use

Whether you should use the default single volume configuration or create additional volumes depends mainly on the storage capacity of the filer.

Use of a single volume: If you want the filer to be configured with a single volume, you do not need to do any further volume configuration after you complete the initial setup.

If your filer doesn't have a large number of disks, a single-volume configuration is probably all you need. You can create additional volumes in the future.

Use of multiple volumes: There are several factors to consider before deciding to create and use multiple volumes:

Configuring with multiple volumes aids in the administration of filers that have large storage capacities, enabling you to

- Perform administrative and maintenance tasks, for example, backup and restore, on individual volumes rather than on a single, large file system.
- Set `option` command values, for example, `snap sched`, `raidsize`, `minra`, `no_atime_update`, and so on, differently for individual volumes.

- Take individual volumes off-line, for example, to perform administrative tasks on their file systems or associated RAID groups, while the other volumes remain on-line, without interrupting the availability of the data on them.

Limitations of configuring with multiple volumes include

- The filer's storage space is partitioned.
- Additional administrative overhead is introduced, for example, defining export points.
- You can expand but not concatenate, shrink, or split volumes.
- You cannot perform a local copy of a volume's contents; you must use `dump` and `restore` or `ndmp copy`.

Planning a multiple volume configuration

Before you configure a filer with multiple volumes, you must decide on the number and sizes of volumes you want to configure.

When deciding the number and sizes of volumes you want to configure, keep the following considerations in mind:

- Configuring with more volumes
  - provides more flexible quota and snapshot configuration
  - requires more export points
- Configuring with larger volumes increases the time needed to restore a volume from tape.
- The maximum number of volumes per filer is 23.
- The maximum *recommended* volume size is 250 GB.

For instructions about specific volume configuration procedures, see "Volume management tasks" on page 89.

Installing a foreign volume

You can remove an entire volume from one filer and install it in another, which makes the moved volume a *foreign* volume to the filer.

Example: For example, you might want to move a volume to a different filer to

- replace the volume's disk shelf with one that has a greater storage capacity
- gain access to the files on a dead filer

When a filer detects a foreign volume at boot time, it places the foreign volume off-line. You can then bring the foreign volume on-line. For more information about installing and bringing up a foreign volume on-line, refer to "Adding a foreign volume" on page 91.

Procedures for managing disks and volumes

Section contents

This section provides step-by-step procedures for performing many filer management tasks from the command line on a filer administration host or client. It is organized into two main sections:

- Disk management tasks
- Volume management tasks

Alternatively, you can perform these procedures using the FilerView program, which has a graphical interface.

Disk management tasks

About this section

This section contains procedures for managing the filer's disks and RAID groups.

Setting the size of a volume's RAID groups

To set the RAID group size of a volume when you create it, enter the command

`vol create volume -r n`

where ***`volume`*** is the name of the volume and ***`n`*** is the number of disks you want in each RAID group. Every RAID group must contain at least two disks.

Changing the size of a RAID group after creating it

To change the size of a RAID group, enter the command

`vol volume options raidsize size`

where ***`size`*** is the number of disks you want in the RAID group.

Note

You can only change the size of the last RAID group in a volume. You cannot change the size of RAID groups after they have been filled.

Installing new disks

New disks are ones that have never been used. Complete the following steps to install new disks.

Step	Action
1	Install one or more disks according to the hardware guide for your disk shelf. The system displays a message confirming that one or more disks were installed, then waits 15 seconds as the disk(s) is turned on. The system recognizes the disks as a hot spare disks. **Note** If you added multiple disks, they might require 25–40 seconds to come up to speed as the system checks the device addresses on each adapter and returns to normal operation.
2	Type the following command: **`sysconfig -r`**
3	Check the `sysconfig -r` output to verify that the new disk has been added.

Adding disks to volumes

To add new disks to a volume, follow these steps:

1. Enter the command

   `sysconfig -r`

   to verify that there are spare disks available for you to add.

2. Add the disks to a volume by entering the command

   `vol add` ***`volume ndisks`***

   where ***`volume`*** is the name of the volume and ***`ndisks`*** is the number of disks you want to add to the volume.

Refer to the *na_vol*(1) man page for details about adding disks to volumes.

Removing a failed disk

When a raid disk has failed, you need not enter any commands—just remove the failed disk from the disk shelf.

Removing a hot spare disk

Complete the following steps if you want to swap disks because you want to use a hot spare disk in another filer:

Step	Action
1	Type the following command and use the output to determine the disk number: **`sysconfig -r`**
2	Type the following command to spin down the disk, replacing *disk_name* with the name of the disk from the output in Step 1: **`disk remove`** ***`disk_name`*** After the disk stops spinning, the disk is ready to be removed.
3	Remove the disk from the filer following the instructions in the hardware guide for your filer model. File service resumes 15 seconds after you remove the disk.

Removing an active file system disk

Complete these steps if you want to remove a disk because it is logging excessive errors.

Step	Action
1	Type the following command and use the output to determine the disk number: **`sysconfig -r`**
2	Type the following command to fail the disk, using the disk name from the output in Step 1 in place of *disk_name*: **`disk fail`** ***`disk_name`*** **Note** The `disk fail` command permanently marks a disk as failed. You cannot reuse the disk; you must replace it. After the `disk fail` command, the system operates in degraded mode, which means that a disk is missing from the RAID group.

Step	Action
3	Remove the disk from the filer following the instructions in the hardware guide for your filer model. File service resumes 15 seconds after you remove the disk.

Volume management tasks

Introduction

This section contains procedures for configuring and managing volumes.

Note

Although you can expand volume sizes and the RAID groups assigned to them after you create an initial multiple volume configuration, you cannot split or shrink volumes after you create them.

Creating volumes

You can create up to 23 volumes on a filer. Each volume must contain at least two disks.

To create a new volume, at the system prompt enter

`vol create` ***`newvol n`***

where ***`newvol`*** is the name for the new volume and ***`n`*** is the number of disks to use. You must have at least ***`n`*** spare disks available.

After creating a new volume

After you create a new volume on a CIFS filer, you must create shares that refer to the new volume to enable clients to access it.

After you create a new volume on an NFS filer, you must

1. Update the system */etc/exports* file.
2. Run `exportfs`.
3. Add the appropriate mount point information to the */etc/fstab* or */etc/vfstab* file on clients that mount volumes from the filer.

Adding disks to a volume

To add more disks to an existing volume, enter

`vol add` ***`vol n`***

where ***`vol`*** is the name of the volume and ***`n`*** is the number of disks to be added.

Monitoring volume status

To determine volume status, such as size, options, disk assignments, and so on, enter

`vol status volume`

To view the RAID group and individual disk information for a particular volume, enter

`vol status -r volume`

To view the RAID group and individual disk information for all volumes, enter

`vol status -r`

Setting volume options

To set various volume options, enter

`vol options volume option value`

Converting a mirror into a regular volume

After you use the SnapMirror™ feature to mirror data to a mirror, you can convert the mirror into a regular volume so that clients can write data to the volume. Enter the following command syntax to convert a mirror to a regular volume:

`vol options volume snapmirrored off`

Although you can convert a mirror into a regular volume, you cannot set the `snapmirrored` option to on to convert a regular volume into a mirror. To start using a volume as a mirror, follow the instructions in Chapter 19, "Data Replication Using SnapMirror," to mirror data to the volume.

Caution

After you convert a mirror into a regular volume, the filer stops using it for data replication. If you want to use the volume as a mirror again, you must take the volume off-line and follow the instructions in Chapter 19, "Data Replication Using SnapMirror," to restart the process of replicating data into the volume.

Making a volume inactive

To remove a volume from active use upon next reboot, enter

`vol offline volume`

Reactivating an off-line volume

To reactivate an off-line volume, enter

`vol online volume`

Adding a foreign volume

To add a foreign volume, that is, a volume that was previously installed on another filer, you move the disks that contain the volume from the old filer to the destination filer.

Complete the following steps to add a foreign volume.

Step	Action
1	Follow the instructions in the hardware guide to remove the disks from the old filer.
2	Turn off the destination filer and install the disks in the destination filer's disk shelf.
3	Turn on and boot the destination filer. **Results:** When the destination filer boots, it places the foreign volume off-line. If the foreign volume has the same name as an existing volume on the filer, the filer renames it *volume_name(1)*, where *volume_name* is the original name of the volume. **Caution** If the foreign volume is incomplete, repeat Steps 1 and 2 to add the missing disks. Do not try to add missing disks while online—doing so will cause them to become hot spare disks.

4	If the filer renamed the foreign volume because of a name conflict, type the following command to rename the volume:
	
	`vol rename oldname newname`
	
	Example: The following command renames the volume *vol0(1)* to *vol1*:
	
	`vol rename vol0(1) vol1`

5	Type the following command to bring the volume on-line in the new filer, replacing *volume_name* with the name of the volume: `vol online volume_name`
6	Enter the following command to confirm that the added volume came on-line: `vol status`

Destroying a volume

To destroy a volume, turning its disks back into spare disks, follow these steps.

Step	Action
1	To deactivate the volume, enter **`vol offline volume`**
2	Enter the **`reboot`** command to reboot the filer.
3	To destroy the volume, enter **`vol destroy volume`**

Renaming a volume

To rename a volume, follow these steps.

Step	Action
1	Enter **`vol rename oldvolume newvol`**
2	Update the */etc/exports* file and run `exportfs`.
3	Update any CIFS shares that refer to the volume.

Handling volume failures

A volume might fail because of an *inconsistent* directory or a double-disk failure. If the system does not reboot after a volume failure, take the volume off-line, as follows:

1. Boot the filer from a system boot diskette into maintenance mode.
2. Use the `vol` command to take the failed volume off-line.

If the failed volume was the filer's root volume, you must designate another volume as the new root volume.

File statistics for volumes

How Data ONTAP provides file statistics

The `filestats` command provides you with a quick way to display a summary of file statistics within a volume on a filer, by reading file information from a snapshot that you specify.

Information obtained by the filestats command

The output from the `filestats` command provides you with a list containing the following information about files from a snapshot in a volume:

- Size
- Creation time
- Modification time
- Owner

The filestats command syntax

The `filestats` command has the following syntax:

`filestats [ages` *`ages`*`] [sizes` *`sizes`*`] [timetype {a,m,c,cr}][style {readable,table,html}]volume` *`volume_name`* `snapshot` *`snapshot_name`*

The `volume` and `snapshot` arguments are required.

volume_name is the name of the volume.

snapshot_name is the name of the snapshot.

Use the `ages`, `timetype`, `sizes`, and `style` options when you want to list specific file information from a volume. For more information about `filestats` options, see "Filestats command options" on page 97.

Example with no options specified

The following example shows sample output from the `filestats` command, without any options:

```
tpubs-cf2> filestats volume vol0 snapshot hourly.1
VOL=vol0 SNAPSHOT=hourly.1
INODES=274528 COUNTED_INODES=875 TOTAL_BYTES=458354190
TOTAL_KB=143556
```

```
FILE SIZE         CUMULATIVE COUNT       CUMULATIVE TOTAL KB
1K                465                    1576
10K               832                    3356
100K              853                    3980
1M                856                    4660
10M               864                    32808
100M              875                    143524
1G                875                    143254
MAX               875                    143254

AGE(ATIME)        CUMULATIVE COUNT       CUMULATIVE TOTAL KB
0                 0                      0
30D               841                    132780
60D               850                    132932
90D               859                    143464
120D              875                    143528
MAX               875                    143528

UID               COUNT                  TOTAL KB
#0                873                    143528
#20041            2                      0

GID               COUNT                  TOTAL KB
#0                851                    41556
#30               21                     1972
#1                3                      0
```

Use file statistics for snapshot management

File statistics help you determine when to schedule snapshots by enabling you to see when most file activity takes place on a volume. Using the `filestats` command also helps you determine snapshot disk consumption.

Example with ages option specified

Enter the `filestats` command with the `ages` option to display a daily breakdown of file changes in a volume, as shown in the following example:

```
filestats ages 1D,2D,3D,4D,5D,6D,7D,8D,9D,10D,11D,12D,
13D,14D volume vol0 snapshot hourly.0
```

- Use the daily age breakdown displayed in the CUMULATIVE TOTAL KB column of the AGE output to determine the average change in data per day.
- Divide the amount of disk space you want to reserve for snapshots by the daily change average. For example, if you find that the average daily change rate is 3 GB and you have a 200 GB volume, 40 GB (or 20 percent) of which

you want to reserve for snapshots, divide 40 by 3 to determine the number of daily snapshots you can have before exceeding your space limit. In this example, 13 daily snapshots is your limit.

Example to determine volume capacity

You can also use the `filestats` command to determine when most activity occurs on a volume during a given day so that you can effectively schedule hourly snapshots.

The following example shows how you can use the `filestats` command to determine when most file changes occur in a volume within a 24-hour period:

```
filestats ages 1H,2H,3H,4H,5H,6H,7H,8H,9H,10H,11H,12H,
13H,14H,15H,16H,17H,18H,19H,20H,21H,22H,23H,24H volume
vol0 snapshot hourly.0
```

If *hourly.0* was taken at 8 a.m., and most file changes took place between 7H and 9H, which corresponds to 3 p.m. and 5 p.m. in this example, you can schedule more snapshots during these hours, and fewer throughout the rest of the day. Scheduling more snapshots before or during increased file activity decreases the time between file changes and snapshots.

For information about managing snapshots, refer to "Managing snapshot disk consumption" on page 403.

Getting a file statistics summary

Description

Use the output from the `filestats` command for a summary of the following information about files in a volume:

- Size
- Creation time
- Modification time
- Owner

Restrictions

The following two arguments are required when using the `filestats` command:

- `volume`
- `snapshot`

Step

To use the `filestats` command, complete the following step.

Step	Action
1	Enter the following command: **`filestats volume`** ***`volume_name`*** **`snapshot`** ***`snapshot_name`*** *volume_name* is the name of the volume. *snapshot_name* is the name of the snapshot.

Filestats command options

Options to use with the filestats command

You can use the following options with the `filestats` command:

- `ages`
- `timetype`
- `sizes`
- `style`
- `expr`

About the ages option

The `ages` option of the `filestats` command enables you to see when files have been accessed. You can specify file ages in seconds, hours, and days, using a comma to separate each value. By default, file ages are broken down by days, in 30-day increments.

Example of the ages option

For example, to display files with ages under 900 seconds (15 minutes), 4 hours, and seven days, respectively, enter the following command:

```
filestats ages 900,4H,7D volume vol0 snapshot hourly.1
```

The output looks like the following:

```
AGE(ATIME)       CUMULATIVE COUNT       CUMULATIVE TOTAL KB
900              0                      0
4H               0                      0
7D               785                    21568
MAX              882                    146000
```

About the timetype option

The `timetype` option enables you to specify the time types that you want to list in the age comparison.

The following table describes the valid `timetype` values you can use with the `timetype` option.

Value	Definition
a	Access time

Value	Definition
`m`	Modification time
`c`	File change time (last size/status change)
`cr`	File creation time

About the sizes option

The `sizes` option enables you to specify the breakdown of sizes, using a comma to separate each value. Default values are in bytes, but you can also use the following three suffixes at the end of a number you specify:

- `K` (kilobytes)
- `M` (megabytes)
- `G` (gigabytes)
- `*` (a special value you use to list all unique file sizes, one line per unique size. Using this command can result in output of several thousands of lines.)

Example of the sizes option

For example, to display file sizes in four categories—files with less than 500 kilobytes, files with less than 2 megabytes, files with less than 1 gigabyte, and all other files, enter the following command:

```
filestats sizes 500K,2M,1G volume vol0 snapshot
hourly.1
```

The output looks like the following:

```
FILE SIZE        CUMULATIVE COUNT       CUMULATIVE TOTAL KB
500K             862                    4969
2M               866                    10748
1G               882                    146000
MAX              882                    146000
```

About the style option

The `style` option controls the output style. The three `style` option arguments are as follows:

- `readable`. This is the default and is what you see when you use the `filestats` command with no `style` option.

- `table`. Use the `table` argument when you the `filestats` output will be used by processing programs.
- `html`. Use the `html` argument for output that will be read by a Web browser.

About the expr option

The `expr` option is an advanced option of the `filestats` command. The `expr` option enables you to specify a Boolean expression that the `filestats` command refers to for each file it encounters.

The following table lists valid file attributes that you can use with the `expr` option.

Attribute	Definition
`tid`	Tree ID
`type`	File type
`perm`	Permissions
`flags`	Additional flags
`nlink`	Count of hard links
`uid`	Numeric user ID of file owner
`gid`	Numeric group ID of file owner
`size`	Size, in bytes
`blkcnt`	Size, in blocks
`gen`	Generation number
`atime`	Time of last read or write, in seconds
`mtime`	Time of last write, in seconds
`ctime`	Time of last size/status change
`crtime`	Time of file creation
`atimeage`	Access time, in seconds
`mtimeage`	Modification age
`ctimeage`	Size/status time age

Attribute	Definition
`crtimeage`	File creation age

Boolean expressions to use with the expr option

You can use the following standard Boolean expressions as arguments to the `expr` option.

Argument	Argument
&&	>
||	<
==	>=
!=	>=
!	

Note

Enclose the entire Boolean string following the `expr` option in quotation marks. Attributes you use in the string must be enclosed in braces {}.

Example of the expr option

For example, if you want to restrict the `filestats` command output to include only files whose size is greater than 10,000 bytes and whose UID is 5150, enter the following command:

```
filestats expr “{size}>10000&&{uid}==5150” volume vol0
snapshot hourly.0
```

Volume reversion using SnapRestore

About SnapRestore

The filer's SnapRestore™ feature enables you to revert a volume to the state it was in when a particular snapshot was taken. Without this feature, you would need to use one of the following methods to restore data to a volume:

- Restoring files from tape
- Copying files from a snapshot to the active file system

Using either of these methods takes a longer time than reverting the volume. This is because with SnapRestore, no data needs to be copied; the file system is just put back in its earlier state.

SnapRestore is a licensed feature. You must purchase and install the license code before you can use it.

How SnapRestore works

After you select a snapshot for reversion, the filer reboots with the volume containing the same data and timestamps as it did when the snapshot was taken. All data that exists before you initiate the reversion is overwritten.

Caution

You cannot undo the reversion to change the volume back to the state it was in before the reversion.

What SnapRestore reverts

SnapRestore reverts only the file contents. It does not revert attributes of a volume. For example, the snapshot schedule, volume option settings, RAID group size, and maximum number of files remain unchanged after the reversion.

Option settings applicable to the entire filer, however, might be reverted. This is because the option settings are stored in a registry in the */etc* directory. If you revert the root directory, the registry is reverted to the version that was in use at the snapshot creation time. For more information about how reverting a root volume works, refer to "Effects of reverting a root volume" on page 103.

You can revert a volume to a snapshot taken when the filer was running a different Data ONTAP version.

Files that SnapRestore cannot recover

You cannot revert a volume to recover a deleted snapshot. For example, if you delete the *hourly.2* snapshot and revert the volume to the *hourly.1* snapshot, you cannot find the *hourly.2* snapshot after the reversion. Although the *hourly.2* snapshot existed at the creation time of the *hourly.1* snapshot, SnapRestore cannot revert the contents of the *hourly.2* snapshot because you already deleted it.

How SnapRestore affects recent snapshots

After you revert a volume to a particular snapshot, you lose the snapshots that are more recent than the snapshot used for the volume reversion. For example, after you revert the volume to the *hourly.1* snapshot, you no longer have access to more recent snapshots, such as the *hourly.0* snapshot. This is because at the creation time of the *hourly.1* snapshot, the *hourly.0* snapshot did not exist.

Typical applications of SnapRestore

If a client application corrupts data files in a volume, you can revert the volume to a snapshot taken before the data corruption. The following examples illustrate some situations in which you can apply SnapRestore to recover from corrupted data.

Example: A messaging application or a database application stores user data in one or two files that can grow to several hundred GB in a volume. If, for some reason, this application corrupts the files, you can revert the volume to a snapshot taken before the data corruption.

Example: You can revert a volume used as a test environment to its original state after each test.

Considerations before using SnapRestore

This section describes two considerations you must make before deciding whether you should use SnapRestore to revert a volume.

Time required for data recovery: If the amount of corrupted data is small, it is easier to copy the files from a snapshot or restore the files from tape than to use SnapRestore for these reasons:

- You can preserve the data in other files in the same volume.
- The filer does not need to reboot.

If the amount of data to be recovered is large, it takes a long time to copy the files from a snapshot or to restore from tape. In this case, SnapRestore is the preferred method for recovering from data corruption.

Free space required for data recovery: If a file to be recovered needs more space than the amount of free space in the active file system, you cannot copy the file from the snapshot to the active file system. For example, if a 10-GB file is corrupted and only 5 GB of free space exists in the active file system, you cannot copy the file from a snapshot to recover the file. In this case, SnapRestore can quickly recover the file; you do not have to spend time making the additional space available in the active file system.

How SnapRestore works with SnapMirror

The following list describes how SnapRestore and SnapMirror interact with each other:

- You can revert a volume that is the source volume for data replication. You cannot, however, revert a volume that is currently the mirror for data replication.
- You can revert to any snapshot that is displayed by the `snap list` command. That is, you can revert to a regular snapshot or a snapshot created by SnapMirror for data replication. The snapshots created by SnapMirror have a different naming convention than the regular snapshots, as explained in "Snapshots created during data replication" on page 522.
- If you have both regular and SnapMirror snapshots in the volume, avoid reverting to a snapshot taken before the SnapMirror snapshot. If you must revert to a snapshot taken before the SnapMirror snapshot, after the reversion, the volume contains no SnapMirror snapshot that is essential for the incremental update of the mirror. The filer must re-create the base-line version of the mirror.

  For example, *vol1* is the source volume for data replication and it contains two snapshots: *hourly.0* and *filerA_vol1.2*. If *hourly.0* was taken earlier than *filerA_vol1.2* and you revert *vol1* to *hourly.0*, you cannot find *filerA_vol1.2* after the reversion. As a result, the filer cannot start an incremental update of the mirror. It must re-create the baseline version of the mirror.

Effects of reverting a root volume

Because the */etc* directory of the root volume contains configuration information about the filer, reverting the root volume might change the filer configuration. The following list describes the effects of reverting the root volume:

- The volume loses the changes that were made to the */etc* directory after the snapshot creation time. Suppose you change the IP address of an interface on the filer after the *hourly.0* snapshot was taken. If you revert the root volume to the *hourly.0* snapshot, the filer reboots with the old IP address for the interface.

- The options used for the entire filer are reverted to the settings that were in effect when the snapshot was taken.

Recommendations: Avoid reverting the configuration files. To avoid reverting the configuration files, follow one of these steps:

- Store all data that might need to be reverted in a volume other than the root volume. This ensures that you never need to revert the root volume.
- If the data you want to revert resides in the root volume, back up the */etc* directory to another volume or another filer before using SnapRestore. After you revert the volume, restore the */etc* directory and reboot the filer.

  If you back up the */etc* directory to another volume, you can use the `vol options` *`volume`* `root` command to make the filer reboot with that volume as the root volume. In this way, when the filer reboots during the reversion, it can use the correct settings in the */etc* directory.

Effects of SnapRestore on filer backup and recovery

Because all files in a reverted volume have timestamps that are the same as those when the snapshot was created, the filer's `dump` and `restore` commands might be affected. Incremental backup and restore operations can no longer rely on the timestamps to determine what data needs to be backed up or restored.

Recommendation: After you revert a volume, perform a level-0 backup of the volume. After the level-0 backup is finished, the filer can perform subsequent incremental backups correctly. Also, if you need to restore data from tape to this volume, use only the backups created after the volume reversion.

Reverting a volume to a selected snapshot

Description

Use SnapRestore to revert a volume to a snapshot. You can use this feature at any time. After you enter the command for reverting a volume, the filer reboots with the volume containing the same data as it did when the snapshot was taken.

Prerequisites

You must meet these prerequisites before using SnapRestore:

- SnapRestore is a licensed feature. You must enter the `snaprestore` license code before you can revert a volume to a snapshot.
- Snapshots must exist on the filer so that you can select one snapshot for the reversion.
- The volume to be reverted must be on-line.
- The volume to be reverted must not be a mirror used for data replication.

Cautions

Be sure that you understand the following rules before reverting a volume:

- SnapRestore overwrites all data in the volume. After you use SnapRestore to revert to a selected snapshot, you cannot undo the reversion.
- When you revert a source volume for data replication, try not to select a snapshot taken before the SnapMirror snapshot. If you must revert to a snapshot taken before the SnapMirror snapshot, the filer can no longer perform an incremental update of the mirror; it must recreate the base-line version of the mirror.
- Snapshot deletions are irrevocable. If you delete a snapshot, you cannot recover the snapshot by using SnapRestore.
- After you revert a volume to a selected snapshot, you lose all the snapshots that were taken after the selected snapshot.
- Reverting the root volume causes the filer to reboot with configuration files that were in effect when the snapshot was taken.

Steps

Follow these steps to revert one or more volumes. At any time before you enter **y** in the last step, if for some reason, you do not want to proceed with the volume reversion, enter Ctrl-C.

Step	Action
1	Notify network users that you are going to revert a volume so that they know the current data in the volume will be replaced by that of the selected snapshot. NFS users can unmount the files and directories in the volume before the reversion. If they do not unmount the files and directories, they might see the “Stale file handle” error message after the volume reversion.
2	If you know the name of the snapshot for reverting the volume, go to Step 5. If you want to review the list of snapshots available for volume reversion, enter the following command: `vol snaprestore` *`volume ...`* *volume* is the name of the volume to be reverted. Enter the name only, not the complete path. You can enter multiple volume names in the command, separated by spaces. **Result:** The filer displays a warning message and prompts you to confirm your decision for reverting the volume. Between the time you enter the `vol snaprestore` command and the time when reversion is completed, the filer stops deleting and creating snapshots.
3	Enter **y** to confirm that you want to revert the volume. **Result:** The filer displays a list of snapshots for you to choose from.
4	Enter the name of the snapshot for reverting the volume and go to Step 7. **Result:** The filer displays the name of the volume and the name of the snapshot for the reversion. Then it asks whether you want to reboot the filer to proceed with the reversion.

Step	Action
5	Enter the following command: `vol snaprestore` *`volume`* `-s` *`snapshot`* `[`*`volume`* `-s` *`snapshot`* `...]` *volume* is the name of the volume to be reverted. Enter the name only, not the complete path. *snapshot* is the name of the snapshot for volume reversion. You can enter multiple pairs of volume name and snapshot name in the command. **Result:** The filer displays a warning message and prompts you to confirm your decision for reverting the volume. Between the time you enter the `vol snaprestore` command and the time when reversion is completed, the filer stops deleting and creating snapshots.
6	Enter **y** to confirm that you want to revert the volume. **Result:** The filer displays the name of the volume and the name of the snapshot for the reversion. Then it asks whether you want to reboot the filer to proceed with the reversion.
7	Enter **y** to confirm that you want to continue with the reversion. **Result:** The filer reboots.

Network Administration 4

About this chapter

Chapter contents

This chapter discusses the topics listed in the following table.

Read...	To learn about...
"Working with large files" on page 110	How the filer supports large files.
"Using SNMP" on page 111	Using SNMP on the filer from a UNIX host, including information about the Network Appliance custom Management Information Base (MIB).
"Host name resolution" on page 117	How the filer resolves names.
"Routing" on page 123	How the filer routes packets using its routing table.
"Using ifconfig to configure an interface" on page 127	How to configure a network interface.
"EtherChannel trunking" on page 132	How to configure trunks or virtual interfaces.
"Database file protection" on page 154	How to enable database file protection.

Working with large files

About large files

You can have large files on your filer. The maximum size of a large file is determined by the smaller of the following two values:

- Maximum file size supported by your server
- Maximum file size supported by your clients

However, if the maximum file size on your client is larger than the maximum file size on your server and you enter a command to display the volume size on your server, the file size is displayed incorrectly. In this case, the maximum file size appears as 2 GB, even though the file on the client is larger than 2 GB.

Caution

If you create a file that exceeds 64 GB, you cannot revert the filer to Data ONTAP 5.0.1 or earlier and access that file.

Software requirements

To use large files, ensure that your system meets the following requirements.

For...	Make sure you that have...
Filer operating system	Data ONTAP 5.1 or later
NFS version 3	enabled
UNIX host and clients	Solaris 2.6 or later
Windows NT host and clients	4.0 or later

How to enable NFS

To enable NFS version 3, use the `options` command to set the `nfs.v3.enable` option to On.

Using SNMP

About SNMP

You use SNMP (Simple Network Management Protocol) to direct a process, called an agent, on the filer to perform network management tasks such as gathering status and diagnostic information. The information is sent to network management stations, which are client workstations on a network. The network management stations use third-party applications to process the information.

The information that is exchanged to perform these tasks is described in ASCII files called Management Information Bases (MIBs).

Data SNMP provides

For diagnostic and other network management services, the filer supports the SNMP MIB-II specification. Based on SNMP version 1, this specification provides data about the following MIB-II groups:

- system
- interfaces
- address translation
- IP
- ICMP
- TCP
- UDP
- SNMP MIB-II

SNMP commands enable users to specify up to eight communities and trap notifications for up to eight management stations.

Command to configure the SNMP agent

To use SNMP, configure the SNMP agent using the `snmp` command. A typical set of SNMP commands in the */etc/rc* file in the root volume is as follows:

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

SNMP commands supported by Network Appliance

SNMP commands that Network Appliance supports are described in detail in the *na_snmp*(1) man page. The following paragraphs provide brief explanations of these commands:

- `snmp contact 'jdoe@abc.com 415-555-1212'`

  Sets the email address and telephone number of the person responsible for the filer. You can include the person's full name, but an email address and a telephone number enable Network Appliance Technical Support to contact the right person after receiving an automatic email message through autosupport from your site concerning your filer status.

- `snmp location 'ABC corporation, engineering lab'`

  Sets the physical location of the filer. This value is returned by the SNMP agent.

- `snmp community add ro private`

  Creates a read-only community called `private`. The SNMP manager uses the community name as a password to communicate with the filer's SNMP agent.

- `snmp traphost add snmp-mgr1`

  Makes the system `snmp-mgr1` the recipient of all SNMP traps from the filer.

- `snmp init 1`

  Initializes the SNMP daemon with the values specified with `snmp` commands. This command also sends the SNMP cold start and links up or down traps, as appropriate, to any trap hosts that were previously registered using the `snmp traphost` command. This command should be the last SNMP command in the filer's */etc/rc* file.

- `snmp traps`

  Generates asynchronous notification of events, called traps.

About the Network Appliance custom MIB

The Network Appliance custom MIB provides detailed information about many aspects of filer operation. The current version of the Network Appliance custom MIB is 1.1.2 It contains objects that help you manage multivolume features.

Where to get the MIB: You can obtain the custom MIB from the Data ONTAP CD, or from the NOW Web page at *http://www.now.netapp.com*. To locate the MIB on the CD, read the *intro.txt* file.

Installing the MIB: Install the MIB file on your network management workstation according to the installation procedure for your workstation, so that your workstation can obtain information from the filer about the objects that are part of the MIB.

Using deprecated single-volume objects: *Single-volume objects* are *deprecated*, but you can use them for single-volume systems as before. If you are using a single-volume filer, you do not need to make changes to use the new MIB.

Finding multivolume objects: The descriptions of the deprecated single-volume objects contain the names of the corresponding new *multivolume objects*. For multivolume objects, use the new objects rather than the deprecated ones.

About MIB group contents

The top-level groups in the custom MIB and the information they contain are described in the following table.

Group name	Contents
`cifs`	Statistics like those displayed by the `cifs stat` command.
`filesys`	Information related to the file system, including the equivalent of the `maxfiles` and `df` commands, and some of the information from the `snap list` command. For more information about `maxfiles`, see "Increasing the maximum number of files" on page 437. For more information about the `df` command, see "The df command" on page 438. For more information about snapshots, see Chapter 11, "Snapshots."
`nfs`	Statistics like those displayed by the `nfsstat` command, including statistics for each client if per-client statistics are enabled. The per-client statistics are indexed by client IP addresses.
`product`	Product-level information, such as the software version string and the system ID.

Group name	Contents
`quota`	Information related to disk quotas, including the output of the `quota report` command. To access quota information, quotas must be turned On. For more information about quotas, see "Restricting or tracking disk usage by using disk quotas" on page 426.
`raid`	Information about RAID equivalents to the `sysconfig -r` output. For more information about `sysconfig`, see Chapter 20, "System Information and Performance."
`sysstat`	System-level statistics, such as CPU uptime, idle time, and the number of kilobytes transmitted and received on all network interfaces.

About traps

You use the `snmp traps` command to inspect the value of MIB variables periodically and send an SNMP trap to the machines on the traphost list whenever that value meets the conditions you specify. The traphost list specifies network management stations that receive trap information.

You use the SNMP third-party applications on your network management station to process the trap information. For example, you can set a trap to monitor the fans on a filer and have the SNMP application on your network management station put a flashing message on your console that tells you that a fan has ceased operating.

You can set traps on any numeric variable in the filer's MIB.

How to define traps

You define or change a user-specified trap using the following `snmp traps` command:

`traps` *`trapname.parameter value`*

`trapname` is the name of the trap.

`value` is the value that you assign to the definition.

`parameter` must be a parameter listed in the following table.

Parameter	Description
`var`	A trap's variable is the MIB variable that is queried to determine the trap's value. All MIB variables must be specified in the form `snmp.`*`oid,`* where oid is an OID (Object Identifier). A list of OIDs in the NetApp MIB is in the `traps.dat` file in the same directory as the MIB.
`trigger`	A trap's trigger is a piece of code that determines whether the trap should send data. The following triggers are available: ◆ `single-edge-trigger`—sends data when the trap's target MIB variable's value crosses an edge—a value that you specify. ◆ `double-edge-trigger`—enables you to have the trap send data when an edge is crossed in either direction (the edges can be different for each direction. ◆ `level-trigger`—sends data whenever the trap's value exceeds a certain level.
`edge-1` `edge-2`	A trap's edges are the threshold values that are compared against during evaluation to determine whether to send data.
`edge-1-direction` `edge-2-direction`	Edge-triggered traps only send data when the edges are crossed in one direction. By default this is UP for the first edge and DOWN for the second edge. The direction arguments let you change this default.
`interval`	The interval is the number of seconds between evaluation of the trap. A trap can only send data as often as it is evaluated.
`interval-offset`	The interval offset is the amount of time in seconds until the first trap evaluation, and is zero by default. You can set it to a non-zero value to prevent too many traps from being evaluated at once (at system startup, for example).

Parameter	Description
`backoff-calculator`	After a trap sends data, you might not want it to be evaluated so often anymore. 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 that it is still full. There are two kinds of backoff calculators: stepwise and exponential.
`backoff-step`	The number of seconds to increase the evaluation interval if you are using a step backoff. If a trap's interval is 10 and its backoff-step is 3590, the trap is evaluated every 10 seconds until it sends data, and once an hour thereafter.
`backoff-multiplier`	The value by which to multiply a trap's evaluation interval each time it fires. If you set the backoff calculator to exponential-backoff and the backoff multiplier to 2, the interval doubles each time the trap fires.

It is important to distinguish between the kind of user-specified traps you can set using the `snmp traps` command, and the built-in support for traps defined in RFC 1213, such as cold-start. Built-in traps such as cold start are automatically sent to the hosts on the traphosts list when some event (a reboot in the case of a cold start) occurs. User-specified traps only exist after they are defined by a series of `snmp traps` commands.

Traps are persistent. After you set a trap, it remains across reboots until you specifically remove it.

Where to get further information

Refer to the following RFCs (Requests For Comments) for more information:

- RCF 1155—defines and describes the SNMP MIB I specification
- RFC-1157—defines and describes SNMP
- RFC-1213—defines and describes the SNMP MIB-II specification

Host name resolution

How the filer resolves host names

The filer resolves host names by searching maps or databases for services to use. The filer tries name resolution services in a default order or in the order that you specify in the */etc/nsswitch.conf* file in the root volume.

Name resolution search

By default, first the filer tries to resolve host names locally by searching the */etc/hosts* file in the root volume and in the */etc/nsswitch.conf* file in the root volume. If it cannot resolve the host name, the filer tries NIS, if NIS is enabled. If the filer still cannot resolve the host name, the filer requests services from a DNS server, if DNS is enabled. You can specify any or all of the resolution methods.

Default search order

The following table shows the default search order for each map.

Map	Services in search order
`hosts`	root files, NIS, DNS
`passwd`	root files, NIS
`netgroup`	root files, NIS
`group`	root files, NIS
`shadow`	root files, NIS

Specifying a search order

To specify a different order in which the filer contacts host name services, create the */etc/nsswitch.conf* file in the root volume. Each line must have the following format:

```
map: service ...
```

Parameter	Description
`map`	One of the following maps or databases: `hosts`, `passwd`, `netgroup`, `group`, or `shadow`.

Parameter	Description
`service`	One or more of the following: ◆ files for local files in the */etc* directory in the root volume ◆ dns for DNS ◆ nis for NIS

Example search order

You can list services in the order in which you want the filer to contact the services. For example, the following file instructs the filer to contact first NIS for hosts, then DNS, and finally local files in */etc* in the root volume. For passwords, the contact order is NIS, then local files in */etc* in the root volume.

```
hosts: nis dns files
passwd: nis files
```

When the filer resolves a host, the search stops.

Note

When performing CIFS operations, the filer can use WINS servers for host name service. However, filer commands unrelated to CIFS use the */etc/hosts* file, DNS, or NIS to resolve host names, as described in the following sections.

Using the /etc/hosts file for host name resolution

The filer can use the */etc/hosts* file in the root volume to resolve host names used in the */etc/rc*, */etc/syslog.conf*, */etc/dgateways*, */etc/exports*, */etc/netgroup*, and */etc/hosts.equiv* root volume files. If you do not use any host names, you do not need */etc/hosts*.

By default, the filer reads the */etc/hosts* file in the root volume whenever it needs to resolve host names, so changes to the file take effect immediately. If you change the IP address for an interface, the new IP address doesn't take effect until you reboot the filer and execute the appropriate `ifconfig` command from */etc/rc*. It is safest to reboot the filer after changing any */etc* file to make sure that the new configuration works correctly.

If you use DNS and you put files first in the */etc/nsswitch.conf* file in the root volume: The filer looks up host addresses in */etc/hosts* in the root volume before sending queries to the DNS server. Therefore, it is important that each entry in */etc/hosts* contains accurate information.

To reduce the need for updating information in */etc/hosts* in the root volume, when using DNS, keep only a minimum number of entries in */etc/hosts*.

If you use NIS: You can do one of the following actions:

- Modify the Makefile of the NIS master to copy the NIS master's */etc/hosts* file from the root volume to the filer when it is changed.
- Have the filer use NIS directly.

Example

You can put the following line at the end of the NIS Makefile section for *hosts.time*:

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

Substitute the name of each filer in the "for filer in..." list in place of toaster1, toaster2, and so on.

Using DNS

The filer includes DNS client capabilities to query DNS servers for host-name-to-IP-address and IP-address-to-host-name translation services. With DNS enabled, you no longer have to update the filer's */etc/hosts* file in the root volume every time you add a new host to the network. (If you use the default search order or put files before DNS in the */etc/nsswitch.conf* file, you still have to update the */etc/hosts* file if one of its entries changes before the filer tries to resolve host names.)

Note

To prevent naming inconsistencies, Network Appliance recommends that when you enable DNS, you use only the default */etc/hosts* file in the root volume.

Enabling DNS during setup: At `setup`, if you enter **y** in response to the following prompt, `setup` prompts you for a DNS domain name, as follows:

```
Do you want to run DNS resolver [n]: y
Please enter DNS domain name []:
```

After you enter a DNS domain name, `setup` prompts you for the IP addresses for up to three DNS name servers. Based on the IP addresses you enter, `setup` generates the */etc/resolv.conf* file in the root volume. Entries in */etc/resolv.conf* file consist of the word “nameserver” followed by an IP address, as follows:

```
nameserver ip_address
```

For details about name server query policies, see the *na_resolv.conf*(5) man page.

Enabling DNS without using setup: If you didn’t start DNS during `setup`, you can start DNS using the following procedure.

Step	Action

1	Create a */etc/resolv.conf* file in the root volume. The file consists of up to three lines, each specifying a name server host in the following format:
	`nameserver ip_address`
	For example:
	`nameserver 192.9.200.10`
	`nameserver 192.9.200.20`
	`nameserver 192.9.200.30`

2	Edit the */etc/rc* file in the root volume to make sure that the option specifying the DNS domain name is set and that the option to enable DNS is set to On.
	For example:
	`options dns.domainname netapp.com`
	`options dns.enable on`

3	Reboot the filer or enter the commands at the filer prompt. **Result:** DNS is now enabled. You no longer have to update the filer’s */etc/hosts* file in the root volume every time you add a new host to the network, unless you specify files first in the */etc/nsswitch.conf* file in the root volume.

Disabling DNS

To disable DNS, enter the following command or put the command in the */etc/rc* file in the root volume to make the change permanent:

```
options dns.enable off
```

Changing your DNS domain name: To change your DNS domain name, enter the following command or put the command in the */etc/rc* file in the root volume to make the change permanent:

```
options dns.domainname domainname
```

Using NIS

The filer includes NIS client capabilities to query NIS servers for host-name-to-IP-address and IP-address-to-host-name translation services.

Note

Because the `nsswitch.conf` file already enables you to specify the order in which the filer finds password information, you do not need to use + or - entries in the filer's */etc/passwd* file in the root volume. Any existing + or - entries are ignored.

NIS maps the filer uses

The filer uses the following NIS maps:

```
hosts.byname
hosts.byaddr
passwd.byname
passwd.byuid
passwd.adjunct
group.byname
group.bygid
netgroup.byhost
```

Shadow password information is obtained from the `passwd.byname` map.

Enabling NIS during setup: During `setup`, the following prompt appears:

```
Do you want to run NIS Client [n]:
```

If you enter **y**, `setup` prompts you for an NIS domain name, as follows:

```
Please enter NIS domain name [] :
```

Enter an NIS domain name.

Enabling NIS without using setup: If you didn't start NIS during `setup`, you can start NIS using the following procedure.

Step	Action

1	Edit the */etc/rc* file in the root volume to make sure that the option specifying the NIS domain name is set and the option to enable NIS is On. Insert lines similar to the following:
	`options nis.domainname netapp.com`
	`options nis.enable on`

2	Enter the `options` commands from the command line or reboot the filer. **Result:** Options entered only on the command line are not saved if you reboot. NIS is now enabled. You no longer have to update the filer's */etc/hosts* file in the root volume every time you add a new host to the network, unless you specify files first in the */etc/nsswitch.conf* file in the root volume.

Disabling NIS: To disable NIS on a running system, enter the following command, or put the command in the */etc/rc* file in the root volume to make the change permanent:

```
options nis.enable off
```

Changing your NIS domain name: To change the domain name on a running system, enter the following command or put the command in the */etc/rc* file in the root volume to make the change permanent:

```
options nis.domainname new.domain
```

Displaying the NIS server name: To display the NIS server name, enter the following command:

```
ypwhich
```

The `ypwhich` command has no options.

Routing

About filer routing

Even though the filer can have multiple network interfaces, it does not function as a router; that is, the filer does not route packets between its interfaces on behalf of other network hosts. It can, however, route its own packets.

Routing table on the filer

For routing its own packets, the filer relies on the default route and explicit routes. Typically, the filer learns explicit routes through `icmp` redirect messages received from the default router; you do not need to enter explicit routes in the filer's routing table. To display the filer's current routing table, use the `netstat -r` command. For example:

```
netstat -r

Routing tables
Internet:
Destination      Gateway            Flags Refs   Use  Interface
default           nwo-               UG    1      138  e0
298.295.227       link#1             UC    0        0  e0
298.295.227.1     link#1             UHL   1       24  e0
nwo-              8:0:20:79:f9:79    UHL   1        0  e0
298.295.227.255   link#1             UHL   1     3696  e0
```

If you must enter explicit information into the filer's routing table, use the `route` command. See the *na_route*(1) man page about how to add or modify information in the routing table.

If the filer cannot find an explicit route in the routing table for a particular destination, it uses the default route. This means that the filer sends the traffic to the default router, which is specified in the */etc/dgateways* file in the root volume.

Specifying default routers

One default router is specified during `setup`. You can, however, add potential default routers at any time to the */etc/dgateways* file in the root volume. For each added router, you should also specify the metric, which is a number indicating the route preference for the router. The highest preference is 1, which is also the default preference for the router specified during `setup`. The lowest preference is 15. For information about the format of the */etc/dgateways* file, refer to "The /etc/dgateways file" on page 125.

Using the routed daemon to manage multiple routers

To help manage multiple routers and to enable you to create redundant routing schemes, the filer runs the `routed` daemon, a simple routing daemon, which is enabled at boot time. This daemon "listens" for Routing Information Protocol (RIP) packets being exchanged between routers on the network to determine which routers are alive. From the routers that are alive, the `routed` daemon selects the one with the highest preference to use as the filer's default router.

However, the filer doesn't rely on the `routed` daemon to construct the routing table. The function of the filer's `routed` daemon is to check the status of the default router. Refer to the *na_route*(1), *na_routed*(1), and *na_dgateways*(1) man pages to learn more about routing on the filer.

You turn the `routed` daemon Off and On with the `routed` command.

To turn the `routed` daemon Off, enter

`routed off`

To turn the `routed` daemon back On, enter

`routed on`

Note

If you turn the `routed` daemon Off by editing */etc/rc* in the root volume, manually designate a default router in */etc/rc*.

Displaying routing status

To display the status of the default gateway list, use the `routed status` command. The `-n` option forces the command to display numeric values for gateway names. An example of the `routed status` display is as follows:

```
routed status
  RIP snooping is on
  Gateway   Metric State   Time Last Heard
  karl           1  ALIVE  Wed Mar 9 03:38:41 GMT 1994
  groucho        1  ALIVE  Wed Mar 9 03:38:41 GMT 1994
  292.0.266.366  1  ALIVE  Wed Mar 9 03:38:41 GMT 1994
  292.0.266.377  1  ALIVE  Wed Mar 9 03:38:41 GMT 1994
  june           2  ALIVE  Wed Mar 9 03:38:41 GMT 1994
  292.0.266.332  2  ALIVE  Wed Mar 9 03:38:41 GMT 1994
  292.0.266.352  3  ALIVE  Wed Mar 9 03:38:41 GMT 1994
  292.0.266.351  4  ALIVE  Wed Mar 9 03:38:41 GMT 1994
  292.0.266.350  5  ALIVE  Wed Mar 9 03:38:41 GMT 1994
```

The `routed status` display shows the following information:

- whether RIP snooping is active (On or Off)

- the current list of default gateways
- the metrics of the default gateways (1 through 15)
- the state of the gateways (ALIVE or DEAD)
- the last time each gateway was heard from

The /etc/dgateways file

The */etc/dgateways* file in the root volume is the configuration database for the `routed` daemon. From the routers that `routed` has determined to be alive, `routed` selects the one with the highest preference to be the default router. When the filer cannot find an explicit route for a packet, it routes the packet to the default router.

The file consists of lines with the following format:

```
gateway metric
```

where *gateway* is the name or the IP address of a default router and *metric* is a preference indicator, which ranges from 1 (highest) to 15 (lowest).

#Gateway	Metric
192.9.200.10	1
eng_gateway	2

Note

Each entry for such a default router must have an IP address that belongs to the IP subnet of one of the interfaces configured for the filer.

How the filer replies to requests

The following list describes how the filer uses its interfaces to respond to different types of packets.

NFS-over-UDP requests: The filer does not use the conventional IP routing mechanisms to reply to NFS-over-UDP requests. The filer sends the response on the network interface on which the request was received to the same address that generated the request. For example, the filer named toaster uses the *toaster-e1* interface to send packets in response to NFS requests received on the *toaster-e1* interface.

This way of handling NFS-over-UDP requests enables you to attach multiple interfaces of the filer to networks with the same IP subnetwork number while keeping NFS-over-UDP traffic isolated to the appropriate physical networks.

Because of this scheme, it is possible that NFS-over-UDP responses might be returned through a different path than you might expect from an examination of the IP routing table using `netstat -r`. This scheme generally works well, although it can result in different routes than expected if your environment contains one-way routes. For example, the IP packets might not be routed as you intended if you configured the network so that the IP traffic from host1 to host2 is routed through router1 and the IP traffic from host2 to host1 is routed through router2.

NFS-over-TCP, -CIFS, and -HTTP requests: The filer tries to return NFS-over-TCP and -HTTP traffic over the interface on which the traffic was received. However, there are exceptions. For example, if the filer experiences excessive queuing to a response or experiences a time-out followed by a retransmit, the filer routes the traffic by using conventional IP routing table lookups. If the filer has multiple interfaces attached to networks with the same IP network number, the filer uses the first interface it finds with that number to send the responses.

Note

If the NFS-over-UDP, -TCP, -CIFS, or -HTTP traffic takes place on an ATM interface, the filer always uses conventional IP routing table lookups and routing for that traffic.

IP-based traffic other than NFS and HTTP requests: For other types of traffic, for example, traffic generated by `telnet`, `rsh`, and `ping`, the filer uses IP routing table lookups and routing. If the filer has multiple interfaces attached to networks with the same IP network number, the filer uses the first interface it finds with that number to send the responses.

Using ifconfig to configure an interface

About the ifconfig command

The */etc/rc* file in the root volume contains `ifconfig` commands to configure network interfaces, including virtual interfaces, at system boot. You can also manually use the `ifconfig` command when the system is operating.

The ifconfig command syntax

The `ifconfig` command syntax is as follows:

```
ifconfig <interface>
    [ [ alias | -alias ] <address> ] [ up | down ]
    [ netmask <mask> ] [ broadcast <address> ]
    [ mtusize <size> ] [ mediatype <type> ]
    [ trusted | untrusted ] [ wins | -wins ]
    [ partner | -partner <address> ]
```

Reasons to use the ifconfig command

You use the `ifconfig` command for the following purposes:

- changing the interface's IP address, network mask, or broadcast address
- setting the media type on an Ethernet interface
- setting the maximum transmission unit (MTU)
- configuring the interface up or down
- specifying the partner IP address in a cluster

Changing the interface's IP address, network mask, or broadcast address

The following examples show configuring an Ethernet interface and a virtual interface on a filer:

```
ifconfig e1 292.9.299.6
ifconfig e2 netmask 255.255.0.0
ifconfig e3 broadcast 292.9.299.255
```

The interface names on your system might be different, depending on your filer model. For information about interface naming conventions, refer to "Naming conventions for network interfaces" on page 43.

Setting the media type on an Ethernet interface

The following example show configuring an Ethernet interface on a filer:

```
ifconfig e1 mediatype tp
```

The media types you can use depend on the type of Ethernet card. The possible types you can enter in the `ifconfig` command are the same as those you can select when running `setup`. They are described in the following table.

Media type	Description
100tx	100Base-T
100tx-fd	100Base-T, full-duplex
tp	100Base-T
tp-fd	10Base-T, full-duplex
auto	100Base-T/100Base-TX Ethernet using Auto-Negotiation
1000fx	Gigabit Ethernet

Setting the maximum transmission unit (MTU)

The following example shows setting the MTU for an FDDI interface:

```
ifconfig f0 mtusize 1500
```

The following table lists the default MTU sizes.

Interface	Default MTU size
FDDI	4352
Ethernet	1500
Gigabit Ethernet	1500
ATM	9188

Use a smaller MTU value for an interface if a bridge or router on the attached network cannot break large packets into fragments.

To view the current setting of the MTU value, use the following command:

```
netstat -i
```

Configuring the interface up or down

The following example illustrates how to configure interfaces up and down:

```
ifconfig e1a up
ifconfig f1 down
ifconfig virt_interface up
```

Specifying the partner IP address in a cluster

Within a cluster, each filer must have one or more interfaces configured with its partner's IP address. The filer uses the interface configured with its partner IP address to serve data while in takeover mode. The following terms relate to IP addresses in clusters.

Term	Definition
partner IP address	One or more IP addresses assigned to a clustered filer's partner, if the interface is a physical interface.
partner name	One or more partner names assigned to a clustered filer's partner, if the interface is a virtual interface.
dedicated standby interface	An interface in a clustered filer that is dedicated to a partner IP address. This interface is idle when the cluster is operating in normal mode. It becomes active when the filer in which it is installed takes over for its partner.
shared interface	An interface in a clustered filer that is shared by the IP addresses of the filer in which it is installed and a partner IP address. The partner IP address is idle when the cluster is operating in normal mode. It becomes active when the filer in which it is installed takes over for its partner.

Partner IP addresses can be configured to use dedicated standby interfaces, or they can be configured to shared interfaces.

Partner IP addresses must be configured to use interfaces on the same subnetwork as the partner.

Dedicated interface example: The following example shows how to configure two filers to assume each other's identities during takeovers using dedicated interfaces.

Filer 1	Filer 2
Name: Toaster1	Name Toaster2
Interface e1 address: 198.9.200.28	Interface e1 address: 198.9.200.38
Interface e2 partner address: 198.9.200.38	Interface e2 partner address: 198.9.200.28

The following commands entered on the toaster1 console do the following tasks:

- Configure interface e1 on toaster1 to use 198.9.200.28 as its own IP address
- Configure interface e2 on toaster 1 to use 198.9.200.38 as the partner IP address for toaster2's e1 interface during takeover:

```
ifconfig e1 198.9.200.28
ifconfig e2 partner 198.9.200.38
```

The following commands entered on the toaster2 console do the following tasks:

- Configure interface e1 on toaster2 to use 198.9.200.38 as its own IP address
- Configure interface e2 on toaster2 to use 198.9.200.28 as the partner IP address for toaster1's e1 interface during takeover:

```
ifconfig e1 198.9.200.38
ifconfig e2 partner 198.9.200.28
```

Shared interface example: The following example shows how to configure two filers to assume each other's identities during takeovers using shared interfaces.

Filer 1	Filer 2
Name: Toaster1	Name Toaster2
Interface e1 address: 198.9.200.28	Interface e1 address: 198.9.200.38
Interface e1 partner address: 198.9.200.38	Interface e1 partner address: 198.9.200.28

The following commands entered on the toaster1 console do the following tasks:

- Configure interface e1 on toaster1 to use 198.9.200.28 as its own IP address.
- Configure interface e1 on toaster 1 to use 198.9.200.38 as the partner IP address for toaster2's e1 interface during takeover:

```
ifconfig e1 198.9.200.28
ifconfig e1 partner 198.9.200.38
```

The following commands entered on the toaster2 console do the following tasks:

- Configure interface e1 on toaster2 to use 198.9.200.38 as its own IP address
- Configure interface e1 on toaster2 to use 198.9.200.28 as the partner IP address for toaster1's e1 interface during takeover:

```
ifconfig e1 198.9.200.38
ifconfig e1 partner 198.9.200.28
```

Edit /etc/rc file to make changes persistent after reboot

If you want changes made with `ifconfig` to remain in effect after a reboot, include the `ifconfig` commands in the */etc/rc* file.

Viewing interface configuration information

The following table illustrates examples of how to use the `ifconfig` command to view interface configuration information.

Description	Syntax
Show the current configurations of all network interfaces.	`ifconfig -a`
Show the current configuration of a specific network interface.	`ifconfig` *`interface`* **Example:** Enter the following command to show the current configuration of interface e0: `ifconfig e0`

na_ifconfig man page has more information

For more information about the `ifconfig` command, refer to the *na_ifconfig*(1) man page.

EtherChannel trunking

Trunks are a logical group of interfaces

To get the security of failover or the throughput that multiple interfaces working as one interface can provide, you can group up to four Ethernet interfaces. You group them into a logical interface unit known as a trunk. A trunk is composed of links, each of which is an interface. Commands that work on physical interfaces also work on trunks, except that the `ifstat` command returns limited trunk information.

Synonyms for trunks

Trunks are also referred to by the following terms:

- Virtual aggregations
- Link aggregations
- EtherChannel virtual interfaces.

This document uses the term trunk.

Interfaces before trunking

The following figure shows four separate interfaces, *e3a*, *e3b*, *e3c*, and *e3d*, before trunking.

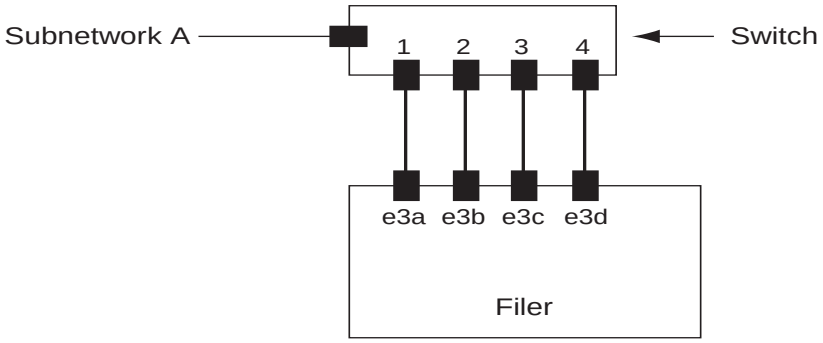

Interfaces after trunking

The following figure shows the four interfaces after trunking into a multiple-mode trunk called *Trunk1*.

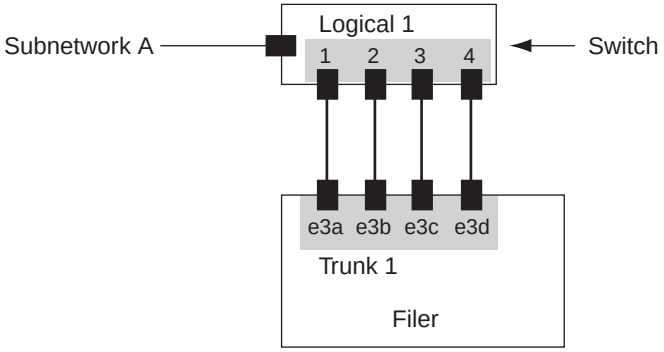

Kinds of trunks

Two kinds of trunks

There are two kinds of trunks:

- Single-mode trunks enable one link to fail over to another link.
- Multiple-mode trunks enable faster throughput by having links share network loads.

Single-mode trunks

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

In the following figure, *e0* and *e1* are part of the SingleTrunk1 single-mode trunk. The active interface, *e0*, fails. Failure means that the link status of the interface is down, which signals that the interface has lost connection with the switch. The *e1* interface takes over and maintains the connection. The interface *e1* also takes over the MAC address of the *e2* interface.

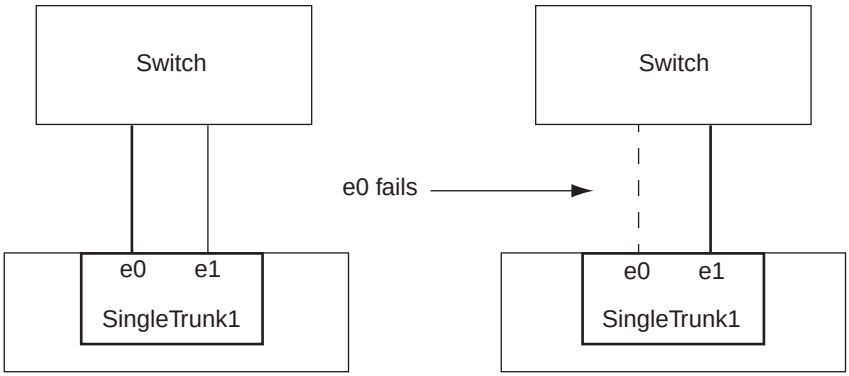

With single-mode trunks, the filer performs takeover based on the absence of a link.

Multiple-mode trunks

In a multiple-mode trunk, all the interfaces are active. This provides greater speed than a single interface.

A multiple-mode trunk requires a switch that supports manually configurable trunking. The switch determines how the load is balanced among the interfaces.

In the following figure, *e0*, *e1*, *e2*, and *e3* are part of the MultiTrunk1multiple-mode trunk. All four interfaces in the MultiTrunk1multiple-mode trunk are active.

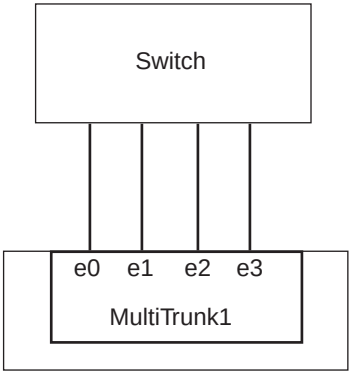

Hardware requirements for trunks

To use a multiple-mode trunk, you need a switch that supports manually configurable trunking over multiple port connections. The switch determines how to forward incoming packets to the filer, so you configure the switch so that all the port connections are part of a single logical port. For information about configuring the switch, see the switch documentation.

Filer network interfaces that are part of the same trunk do not have to be on the same network card, but some Ethernet switches and routers require that all members of the trunk be either half duplex or full duplex.

Virtual interfaces

Trunking supported by virtual interface feature

The feature that supports trunking is known as a virtual network interface, or virtual interface. The filer treats a virtual interface in the same way as a physical interface, except for the `ifstat` command, which does not show virtual interface information. The `vif` command is customized for virtual interfaces and can provide information that `ifconfig` can't. For example, to get status information about a virtual interface, use the `vif status` command rather than `ifstat`.

Naming virtual interfaces

The name of a virtual interface is a string that is no longer than 15 characters that meets the following criteria:

- It must begin with a letter.
- It must not contain a space.
- It must not already be in use for a virtual interface.

Virtual interface names are case-sensitive.

Trunking virtual interfaces

You can trunk virtual interfaces

You create a trunk of virtual interfaces to eliminate a switch as a single point of failure. A trunk of virtual interfaces is known as a second-level virtual interface.

Note

With second-level interfaces, if a switch fails and there is a failover to another switch, it might take a few minutes for the spanning tree relay to be reconfigured.

Second-level interface configurations

You can use second-level virtual interfaces on a single filer or in a cluster.

Second-level virtual interfaces on a single filer

Why use second-level virtual interfaces on a single filer

You use second-level virtual interfaces on a single filer to maintain service even if a switch fails.

Example of a second-level virtual interface on a single filer

A subnetwork has two switches that are capable of trunking over multiple port connections. The filer has a two-link multiple-mode trunk to one switch and a two-link multiple-mode trunk to the second switch. You can create a second-level trunking single mode that contains both of the multiple-mode trunks. When the second-level trunk is configured using the `na_ifconfig`(1) command, only one of the two multiple-mode trunks is brought up as the active link. If all the underlying interfaces in the active trunk fail, the second-level trunk activates the link corresponding to the other trunk.

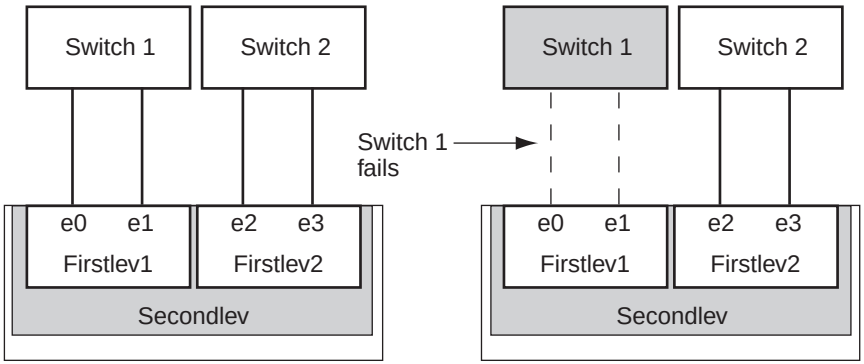

In the preceding figure, Secondlev is the second-level virtual interface and is composed of the two virtual interfaces *Firstlev1* and *Firstlev2*. *Firstlev1* is initially the active interface; if Switch 1 drops both links, Switch2 and *Firstlev2* take over and maintain the connection to the network.

Second-level virtual interfaces in a cluster

Second-level virtual interfaces in a cluster can provide additional protection

In a cluster, second-level virtual interfaces can provide additional protection against switch failures. Normally, in a cluster, network connectivity is maintained even if one filer fails. With second-level virtual interfaces, even if a switch fails, network connections are maintained. Network connections are maintained even in the case of a filer failure and a switch failure.

Normal cluster operation with second-level virtual interfaces

The following figure shows second-level trunks in a cluster. Both filers are in operations, and Filstlev1 in Secondlev1 connects Filer 1 to the network through Switch 1 and Firstlev 4 in Secondlev2 connects Filer 2 to the network through Switch 2.

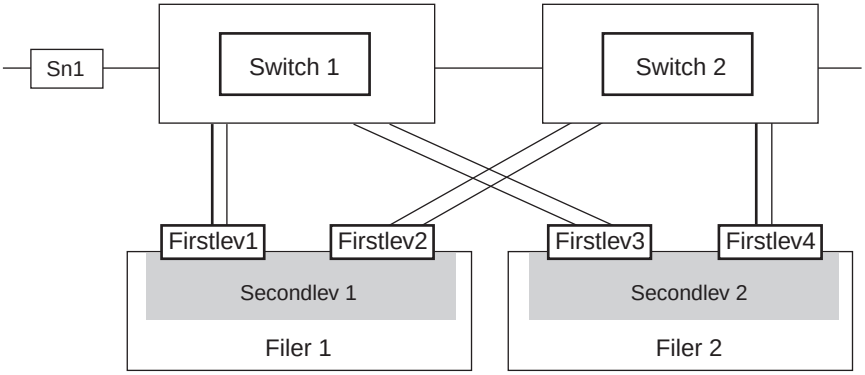

Switch failure in a cluster with second-level virtual interfaces

If one of the switches fails, the following happens:

- If Switch1 fails, Firstlev2 and Firstlev 4 maintain the connection for their filers through Switch 2.
- If Switch2 fails, Firstlev1 and Firstlev 3 maintain the connection for their filers through Switch 1.

In the following figure, Switch 1 fails in a cluster. Firstlev2 takes over the connectivity of Firstlev 1 through Switch 2.

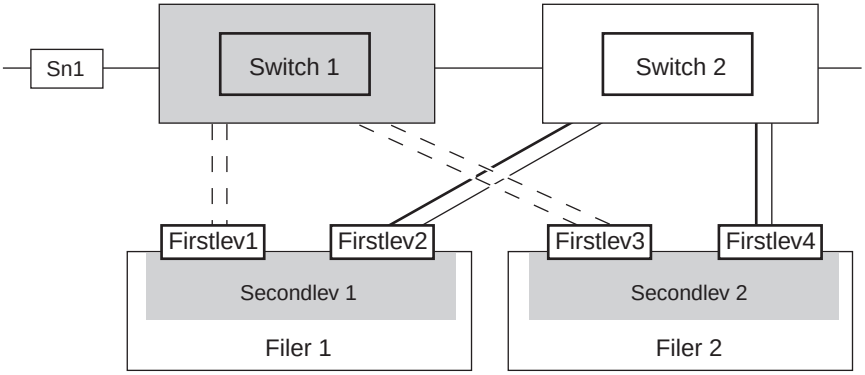

Virtual interface management

Use the vif commands to manage virtual interfaces

You manage virtual interfaces with the `vif` command. This command enables you to perform the following actions:

- Create either a single-mode or a multiple-mode trunk.
- Specify a preferred link to activate in a single-mode trunk.
- Add physical interfaces to a trunk.
- Display the status of a virtual interface.
- Display virtual interface statistics, such as the number of packets received and transmitted on each link that makes up a virtual interface. You can specify the time interval, in seconds, at which the statistics are displayed.
- Destroy a trunk.

Put these vif commands in /etc/rc

The following commands are not persistent:

- `vif add`
- `vif create`
- `vif favor`

Put these commands in the */etc/rc* file to make them persistent across reboots.

Creating a single-mode trunk

Description

Use this procedure to create a trunk in which only one interface is active at a time, thereby ensuring access to a network. After you complete this procedure, the interfaces you specify are combined into a trunk in which one interface is active and the others are ready to take over in case of failure of the active interface.

Prerequisites

You need the following items to complete the procedure:

- A name for the trunk that meets the following criteria:
  - It must begin with a letter.
  - It must not contain a space.
  - It must not already be in use for a virtual interface.

    Trunk names are case-sensitive.
- A list of the interfaces you want to combine into the trunk.

Step

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

Step	Action
1	Enter the following command or put it in the in the */etc/rc* file: **`vif create single`** ***`trunk interfaces`*** ***`trunk`*** is the name of the trunk. ***`interfaces`*** is a list of the interfaces you want the trunk to consist of.

Example

The single-mode trunk shown on "Single-mode trunks" on page 134 was created with the following command:

```
vif creatre single SingleTrunk1 e0 e1
```

Specifying a preferred link in a single-mode trunk

Description

Use this procedure to specify an interface you prefer to be active in a single-mode trunk. For example, if you are replacing or upgrading a physical interface and want the new interface to be the active link, you make the new interface the preferred link. After you complete this procedure, the link you specify is the active link and the other links are not active.

Note

If no links are preferred, the active link is selected randomly.

Step

To specify an interface you prefer to be the active link in a single-mode trunk, complete the following step.

Step	Action
1	Enter the following command or put it in the in the */etc/rc* file: **`vif favor`** ***`interface`*** ***`interface`*** is the name of an interface.

Removing a link from preferred status in a single-mode trunk

Description

Use this procedure to remove a link from preferred status in a single-mode trunk. A preferred link has priority for being the active link. You might want to perform this procedure if you are replacing or upgrading a physical interface that is currently the preferred link and want to stop using it.

Note

There can only be one preferred link in a trunk.

After you complete this procedure, the currently preferred link is not the preferred link, and you can use the `vif favor` command to make another link the preferred link.

Step

To remove a link from preferred status in a single-mode trunk, complete the following step.

Step	Action
1	Enter the following command: **`vif nofavor`** ***`interface`*** ***`interface`*** is the name of an interface.

Creating a multiple-mode trunk

Description

Use this procedure to create a trunk in which all interfaces are active at the same time, thereby increasing throughput. After you complete this procedure, you have increased throughput compared to using only a single interface. You can perform this procedure any time you have a configured switch and the physical interfaces available.

Prerequisites

You need the following items to complete the procedure:

- A switch that supports manually configurable trunking configured according to the manufacturer's instructions for a multiple-mode trunk.
- A name for the trunk that meets the following criteria:
  - It must begin with a letter.
  - It must not contain a space.
  - It must not already be in use for a virtual interface.

    Trunk names are case-sensitive.
- A list of the interfaces you want the trunk to consist of.

Step

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

Step	Action
1	Enter the following command or put it in the in the */etc/rc* file: **`vif create multi`** ***`trunk interfaces`*** ***`trunk`*** is the name of the trunk. ***`interfaces`*** is a list of the interfaces that make up the trunk.

Example

The multiple-mode trunk shown on "Single-mode trunks" on page 134 was created with the following command:

```
vif creatre multi MultiTrunk1 e0 e1 e2 e3
```

Creating a second-level virtual interface on a single filer

Description

You use this procedure to create a second-level virtual interface on a single filer. You follow this procedure when you have available two switches configured for multiple-port connections and four or more interfaces. After you complete this procedure, all interfaces are active at once until there is a switch failure, in which case connectivity is maintained by the links to one switch even if the other switch fails.

Steps

To create a second-level virtual interface on a single filer, complete the following steps.

Step	Action
1	Enter the following commands to create two multiple-mode interfaces: **`vif create multi`** ***`trunk1 if1 if2`*** **`vif create multi`** ***`trunk2 if3 if4`*** ***`trunk1`*** and ***`trunk2`*** are the names of the trunks. ***`if1`***, ***`if2`***, ***`if3`***, and ***`if4`*** are interfaces.
2	Enter the following command to create a single-mode interface from the multiple-mode interfaces: **`vif create single`** ***`secondlev trunk1 trunk2`*** ***`secondlev`*** is the name of a second-level virtual interface. ***`trunk1`*** and ***`trunk2`*** are the names of the trunks.

Example

The following commands create the second-level virtual interface shown in "Example of a second-level virtual interface on a single filer" on page 139.

```
vif create multi Firstlev1 e0 e1
vif create multi Firstlev2 e2 e3
vif create single Secondlev Firstlev1 Firstlev2
```

Creating second-level virtual interfaces in a cluster

Description

You use this procedure to create a second-level virtual interface in a cluster. You follow this procedure when you have available two switches configured for multiple-port connections and eight or more interfaces. After you complete this procedure, connectivity for both filers is maintained even if one switch fails.

Steps

To create a second-level virtual interface in a cluster, complete the following steps.

Step	Action
1	Enter the following commands to create two multiple-mode interfaces on one filer: **`vif create multi`** ***`trunk1 if1 if2`*** **`vif create multi`** ***`trunk2 if3 if4`*** ***`trunk1`*** and ***`trunk2`*** are the names of the trunks. ***`if1`***, ***`if2`***, ***`if3`***, and ***`if4`*** are interfaces.
2	Enter the following command on the same filer as in Step 1 to create a single-mode interface from the multiple-mode interfaces: **`vif create single`** ***`secondlev1 trunk1 trunk2`*** ***`secondlev1`*** is the name of a second-level virtual interfacer ***`trunk1`*** and ***`trunk2`*** are the names of the trunks.
3	Enter the following commands on the other filer to create two multiple-mode interfaces: **`vif create multi`** ***`trunk3 if5 if6`*** **`vif create multi`** ***`trunk4 if7 if8`*** ***`trunk3`*** and ***`trunk4`*** are the names of the trunks. ***`if5`***, ***`if6`***, ***`if7`***, and ***`if8`*** are interfaces.

Step	Action
4	Enter the following command on the same filer as in Step 3 to create a single-mode interface from the multiple-mode interfaces: **`vif create single`** ***`secondlev2 trunk3 trunk4`*** ***`secondlev2`*** is the name of a second-level virtual interface. ***`trunk3`*** and ***`trunk`*** 4 are the names of the trunks.
5	Enter the following commands to configure the second-level virtual interfaces for takeover: **`ifconfig`** ***`secondlev1`*** **`partner`** ***`secondlev2`*** **`ifconfig`** ***`secondlev2`*** **`partner`** ***`secondlev1`*** ***`secondlev1`*** and ***`secondlev2`*** are the names of second-level virtual interfaces.

Example

The following commands create the second-level virtual interface in the cluster shown in "Normal cluster operation with second-level virtual interfaces" on page 140.

```
vif create multi Firstlev1 e1 e2
vif create multi Firstlev2 e3 e4
vif create single Secondlev1 Firstlev1 Firstlev2
```

```
vif create multi Firstlev3 e5 e6
vif create multi Firstlev4 e7 e8
vif create single Secondlev2 Firstlev3 Firstlev4
```

```
ifconfig Secondlev1 partner Secondlev2
ifconfig Secondlev2 partner Secondlev1
```

Adding physical interfaces to a trunk

Description

Use this procedure to add one or more physical interfaces to a trunk. After you complete this procedure, you get the following results, depending on the type of trunk you are configuring:

- For a single-mode trunk, you get improved reliability by providing additional interfaces to fall back on.
- For a multiple-mode trunk, you get improved throughput.

You can perform this procedure any time after the trunk you are adding interfaces to has been created. You must also configure the switch for the additional ports and physical interfaces available for the additions to take effect.

Step

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

Step	Action
1	Enter the following command or put it in the in the */etc/rc* file: **`vif add`** ***`trunk interfaces`*** ***`trunk`*** is the name of a previously configured virtual interface. ***`interfaces`*** is a list of the physical interfaces you want to add to the trunk.

Displaying the status of a trunk

Description

Use this procedure to display the status of a specified trunk. After you complete this procedure, you get information that, for example, is useful in troubleshooting trunk problems. You can use this procedure any time.

Step

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

Step	Action
1	Enter the following command: **`vif status`** ***`trunk`*** ***`trunk`*** is the name of the trunk. If you don't specify a trunk, the status of all trunks is displayed.

Sample output

```
vif status
default: transmit 'multi', fail 'log'
vif0: 2 links, transmit 'multi', fail 'default'
    up:
        e3a: state up, since 28Jan1999 23:57:26 (00:00:07)
                mediatype: 100tx-fd
                flags: enabled next-link
                input packets 52, input bytes 3178
                output packets 6, output bytes 553
                up indications 6, broken indications 0
                indication: up at 28Jan1999 23:57:28
                        consecutive 6, transitions 0
        e3b: state up, since 28Jan1999 23:57:24 (00:00:05)
                mediatype: 100tx-fd
                flags: enabled address-set-1
                input packets 34, input bytes 2915
                output packets 5, output bytes 210
                up indications 4, broken indications 0
                indication: up at 28Jan1999 23:57:28
                        consecutive 4, transitions 0
```

Displaying trunk statistics

Description

Use this procedure to display statistics for a specified trunk over a specified period of time. After you complete this procedure, you get information that, for example, is useful in troubleshooting trunk problems. You can use this procedure any time.

Step

To display statistics, complete the following step.

Step	Action
1	Enter the following command: **`vif stat`** ***`trunk interval`*** ***`trunk`*** is the name of the trunk. If you don't specify a trunk, the status of all trunks is displayed. ***`interval`*** is the interval, in seconds. The default is one second.

Sample output

This is sample output for the `stat` option.

```
vif stat vif0
Virtual interface (trunk) vif0
    e3a                       e3b
In      Out               In      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
```

Destroying a trunk

Description

Use this procedure to destroys, or deletes, a trunk. You destroy a trunk when you cease needing it or want to use the interfaces for other purposes than a trunk. After you complete this procedure, the interfaces in the trunk act individually rather than as links in a trunk.

Prerequisites

Before you destroy a trunk, you must configure it down using the `ifconfig down` command.

Step

To destroy a trunk, complete the following step.

Step	Action
1	Enter the following command: **`vif destroy`** ***`trunk`*** ***`trunk`*** is the name of the trunk.

Database file protection

How Data ONTAP provides database file protection

Data ONTAP provides database file protection through the `nvfail` option of the `vol options` command. The `nvfail` option enables a filer to detect NVRAM inconsistencies at boot time. Use it to warn database administrators of NVRAM problems that could compromise database validity. If the filer finds any problems, database instances hang or shut down, and the filer sends error messages to the console to alert you to check the state of the database.

How to provide additional protection for database files

You can provide additional protection to specific database files by adding them to an optional file you create called */etc/nvfail_rename*. When you enable `nvfail` and the filer detects NVRAM errors at boot up, the filer renames any database files specified in the *nvfail_rename* file by appending *.nvfail* to the original file name. When the filer renames a database file, the database cannot restart automatically. This gives you the opportunity to examine the file for inconsistencies before you remove the *.nvfail* extension and make the file accessible again.

How nvfail works

When you enable `nvfail`, the following process takes place during boot-up.

If...	**Then...**
The filer detects no NVRAM errors	File service starts normally.
The filer detects NVRAM errors and you use the optional *nvfail_rename* file	a) The filer returns a stale filehandles (ESTALE) error to NFS clients trying to access the database, causing the application to hang, crash, or shut down, and sends an error message to the filer console and log file. b) The filer renames database files specified in the *nvfail_rename* file by appending *.nvfail* to the original file name, making those files unavailable to both CIFS and NFS clients.

If...	Then...
The filer detects NVRAM errors and you do not use the optional *nvfail_rename* file	a) The filer returns a stale filehandles (ESTALE) error to NFS clients trying to access the database, causing the application to hang, crash, or shut down, and sends an error message to the filer console and log file. b) No database files are renamed. When the application restarts, files are available to both CIFS and NFS clients, even if you have not verified that they are valid.

Where to look for database file verification instructions

See the documentation for your specific database software for instructions about examining database file validity.

Error message example

When you enable `nvfail` and the filer encounters NVRAM errors, the message sent to the console and the */etc/messages* file looks like the following:

```
All filehandles have been invalidated due to previous NVRAM
failure.
All filesystems must be remounted by the client(s).
```

If you created an *nvfail_rename* file, you receive an additional message:

```
Renaming files in /etc/nvfail_rename.
nvfail_rename: old_file_name: new_file_name
```

Enabling and disabling database file protection with nvfail

Description

Use the `nvfail` option with the `vol options` command to provide database file protection by turning NVRAM error processing On or Off. The default setting is Off.

Step to enable nvfail

To enable the `nvfail` option, complete the following step.

Step	Action
1	Enter the following command: **`vol options volume_name nvfail on`** *volume_name* is the name of the volume.

Step to disable nvfail

To disable the `nvfail` option, complete the following step.

Step	Action
1	Enter the following command: **`vol options volume_name nvfail off`** *volume_name* is the name of the volume.

Using the nvfail_rename file for additional database protection

Description

Use the optional *nvfail_rename* file when you want to rename database files after the filer detects NVRAM errors. This enables you to examine the files for consistency before clients can access them.

Restrictions

You can only list one database file name per line in the *nvfail_rename* file, but you can list as many files as you want.

Steps

To create the *nvfail_rename* file, complete the following steps.

Step	Action
1	Use an editor to create (or modify) the *nvfail_rename* file in the filer's */etc* directory.
2	List the path name and file name, one file per line, within the *nvfail_rename* file; for example: */vol/vol1/home/dbs/oracle-WG73.dbf*
3	Save the file.

ATM LANE and SPANS Administration 5

Overview

About this chapter

This chapter outlines the features and concepts of ATM and describes the key components of an Emulated LAN (ELAN), including the LAN Emulation Client (LEC), LAN Emulation Server (LES), LAN Emulation Configuration Server (LECS), and Broadcast and Unknown Server (BUS). It also describes how to use FORE/IP over SPANS.

Chapter contents

This chapter discusses the following topics.

- "Overview of ATM" on page 160 describes the key features and concepts of ATM and ATM LANE.
- "ATM LANE" on page 161 and subsequent sections describes the key components of an ELAN (Emulated Local Area Network), including the LEC, LES, LECS, and BUS, and how to configure an ELAN.
- "FORE/IP over SPANS" on page 183 describes how to configure ATM PVCs (Permanent Virtual Circuits).

Overview of ATM

Introducing ATM

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

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

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

How you can use ATM on a filer

You can use ATM in two ways on a filer:

- ATM LANE (LAN Emulation), which provides the services of an Ethernet LAN to higher-level network application software
- FORE/IP over SPANS (Simple Protocol for ATM Network Signaling) with PVCs

The filer can simultaneously support Fore/IP over SPANS as well as LANE over UNI 3.0 or 3.1 on the same physical interface.

ATM LANE

Understanding LANE

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

Differences between LANs and ATM:

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

- The shared medium approach of LANs make them ideal for broadcast and multicast messages.
- The MAC addresses used to identify a network interface are typically based on the manufacturer's serial numbers.

  This means that the address is constant, independent of the network to which the interface is connected.

Services provided by LANE

LANE is an ATM service that offers the following services:

- The capability to run LAN-based application software on an ATM network.
- The capability for ATM networks to be interconnected with conventional LANs by means of existing bridging methods.

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

  This means that users cannot communicate between ELANs without the appropriate routing or bridging methods.

What is an ELAN?

An Emulated LAN consists of the following components:

- A single LAN Emulation Service (LANE Service), which itself consists of a LANE Configuration Server (LECS), a LANE Server (LES), and a

Broadcast and Unknown Server (BUS). Each of these components are discussed in more detail later in this chapter.

- A set of LAN Emulation Clients (LECs).

  The LEC communicates with other LECs and with the LANE Service using Virtual Channel Connections (VCCs) in an ATM Switched Virtual Circuit (SVC) environment.

An ELAN comprises a group of ATM-attached devices, logically analogous to a group of LAN stations attached to an Ethernet/802.3 LAN segment. Several ELANs can be configured within an ATM network, and membership to an ELAN is independent of where the end station is physically connected. The end station can belong to multiple ELANs.

Addressing and address resolution

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

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

LAN Emulation Client (LEC)

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

LECs communicate with other clients using the LANE Service, and represent users by their MAC addresses. A LEC employs separate VCCs for data and control communication, including LE_ARP requests for address resolution. User data, intended for another end station, is encapsulated in IEEE 802.3 frames.

LAN Emulation Service (LANE Service)

The LANE Service, consisting of a LES, BUS, and LECS, can be implemented as part of one or more end systems or as part of the ATM switch. In cases where the service is implemented in a distributed fashion over multiple devices, benefits include parallel operation as well as better error recovery through redundancy.

Within the LANE Service, the LES is responsible for coordinating the control functions, while the LECS serves network clients by supplying ELAN configuration information. Finally, the BUS forwards broadcast and multicast frames, and handles unresolved unicast frames.

LAN Emulation Server (LES)

The LES performs a number of control services within an ELAN, including registering MAC addresses and resolving these addresses to corresponding ATM addresses when requested. LECs can also request the LES to resolve a route descriptor to an ATM address. In either case, the LES responds to the client or forwards the request to other clients, which might be in a better position to service the query.

The LES also coordinates the process of a LEC joining an ELAN. There is a single LES per ELAN, and each LES has a unique ATM address.

LAN Emulation Configuration Server (LECS)

The LECS maintains information concerning all the ELANs in an administrative domain, and supplies the LEC with the ATM address for the LES in the domain. Before joining an ELAN, the LEC first exchanges configuration information with the LECS.

Upon successfully communicating with the LECS, the LEC receives a list of ELANs that are available to join. There is a single LECS per administrative domain for ELANs within a domain.

Broadcast and Unknown Server (BUS)

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

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

Although there might be multiple BUSes defined within an ELAN, each LEC is associated with only a single BUS per ELAN.

UNI load balancing

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

With respect to load balancing, both incoming and outgoing connections are directed to the least used adapter in a load balancing group. In cases when an adapter in a load balancing group fails, the connections on that adapter are automatically transferred to another adapter.

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

LANE standards supported in this release

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

- ATM Forum LANE Version 1.0 LEC Support
- UNI 3.0 and 3.1
- ILMI Address Registration
- ILMI MIB extensions for LANE

The software works with the FORE OC3 ATM network interface and on Network Appliance F520, F630, and F760 filers. The software provides Ethernet LANE services, with the capability to configure multiple ELANs on each available network interface.

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

Configuring the ATM adapter for LANE

Overview

This section and the following sections describe how to configure the ATM adapter in the filer to communicate using ATM LAN Emulation (LANE). The sections also describe how to verify that an ATM adapter is properly configured and operating on one or more Emulated LANs (ELANs).

Overview of the procedure for configuring the ATM adapter

Configuring the ATM adapter consists of the following steps.

Step	Action
1	Complete the steps described in "Getting started with LANE configuration" on page 167.
2	Configure the LECS address, as described in "Configuring the LECS address" on page 171.
3	Configure the ATM adapter, as described in "Configuring the ATM adapter for an ELAN" on page 174.
4	Check and save the configuration, as described in "Checking and completing the ELAN configuration" on page 177.

Before you begin

Before you start configuring the ATM adapters in the filer, check the following table.

Prerequisite	Explanation
Complete the normal setup procedure for your filer, run automatically when you first install the filer.	You need an ATM switch with one or more ELANs already configured on the switch (with the corresponding configurations for the LECS, LES, and BUS)
Know the LECS address for the ELAN that you want the filer to join.	In most cases, the LECS has been configured to use the “well-known” address; however, this might be different at your site.
If your site has multiple ELANs, know the ATM address of the LECS for each ELAN you want a client to join.	You can configure each ATM adapter in the filer to communicate over multiple ELANs on the network.

Note

If you need more information about creating an ELAN or configuring the LECS, LES, and BUS, refer to the documentation that came with your switch.

Getting started with LANE configuration

Steps for starting the configuration

To start the configuration, complete the following steps.

Step	Action
1	Check that the ATM adapter is functioning. Use the following commands: ◆ `adinfo`—to check that the ATM adapter is installed and functioning properly ◆ `adstat`—to check that the ATM adapter has a working physical connection to the ATM network
2	Check that the User Network Interface (UNI) is operational. The UNI provides an interface between ATM clients and the ATM switch.

Using the adinfo command

To access an ELAN, the ATM adapter in your filer must be installed and functioning properly. Use the `adinfo` command to verify this.

Command syntax: The `adinfo` command has the following syntax:

```
adinfo
```

Example: The sample output from this command is as follows:

```
a1:unit 1:      PCA-200E Media=OC3-MM-SC HW=2.0.1 FW=4.2.0
Serial=4201600 Slot=1

        MAC=00:20:48:40:1C:80

a2:unit 2:      PCA-200E Media=OC3-MM-SC HW=2.0.1 FW=4.2.0
Serial=4201621 Slot=2

        MAC=00:20:48:40:1C:95

a3:unit 3:      PCA-200E Media=OC3-MM-SC HW=2.0.1 FW=4.2.0
Serial=4201955 Slot=3

        MAC=00:20:48:40:1D:E3
```

Interpreting the output: You should see lines for each of the ATM adapters in your filer that is functioning properly. The presence of the line indicates that the adapter has passed its self-test procedure and that the filer initialized the adapter. However, it doesn't verify that the adapter is sending or receiving packets. The `adinfo` command also displays the device name for each of the installed adapters, as well as the unit number, at the beginning of each line.

The unit number uniquely identifies the ATM adapter in the filer, and there is a one-to-one mapping between the device names and unit numbers. The device name consists of the prefix 'a' followed by the physical slot number. The unit number is simply the slot number.

Some commands, such as `uniconfig` and `elconfig`, require the unit number as a parameter. Other commands, including `adinfo` and `adstat`, require the device name.

You might want to note the unit number for any ATM adapters you plan to configure in the near future. You can also perform the `adinfo` command again at any time to determine the correct unit number.

Using the adstat command

The ATM adapter in your filer should be properly connected using the appropriate cabling. To check that your filer has a working physical connection to your ATM network, use the `adstat` command.

Command syntax: The `adstat` command has the following syntax:

```
adstat -d device
```

device is the name of the particular ATM adapter whose connection you want to check.

Example: The following command displays statistics about the ATM adapter in slot 1 of your filer:

```
adstat -d a1
```

The sample output from this command is as follows:

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

Interpreting the output: Verify that the `Carrier` is `ON` in the output. If the `Carrier` is not `ON`, your cabling is either incorrectly connected or faulty, or your ATM network is malfunctioning or misconfigured.

Note

If you need information about connecting the cabling to your filer's ATM adapter, refer to the appropriate section in the *Hardware Guide* that came with your filer.

Checking that the UNI is operational

To check whether the UNI is operational, use the `uniconfig` command.

Because the filer's ATM address is automatically registered with the switch, you only use the `uniconfig` command to display configuration parameters, check the UNI version number, and ensure that the UNI is operational.

Command syntax: The `uniconfig` command has the following syntax:

```
uniconfig show
```

Example: The abbreviated sample output (showing unit 3 only) from this command is as follows:

```
UNI parameters for unit 3
=========================
VPI/VCI                 : 0/5
AAL type                : 5
QoS                     : UBR
UNI configured version  : 3.1
UNI operating version   : 3.1
```

```
SSCOP operational state : operational

Primary ATM address     :
47.0005.80.ffe100.0000.f21a.4d19.002048401de3.00

UNI failover configuration

==========================

Status: dynamic

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

Interpreting the output: To verify that the ATM interface is operational, check the following items:

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

Configuring the LECS address

About configuring the LECS address

For the filer to participate in an ELAN, you must configure the ATM adapter with the address of the LECS. The LEC joins an ELAN by first exchanging configuration information with the LECS. The LECS then supplies the client with the ATM address for the LES (LANE Server).

Procedure for configuring the LECS address

To configure the LECS address, complete the following steps.

Step	Action
1	Set the LECS address.
2	Check that the adapter is communicating with the LECS.

Setting the LECS address

To configure the LECS address for your ATM adapter, use the `elconfig set` command.

Command syntax: The `elconfig set` command has the following syntax:

`elconfig set -lecs` *`ATM_address`* `| -wellknown -unit` *`unit_number`*

ATM_address is the ATM address of the LECS on the network.

unit_number is the unit designator for the ATM adapter in the filer.

Example: The following command sets the LECS address to the well-known ATM address for the adapter in slot 2:

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

If the LECS on your network does not use the well-known address, specify this address in place of “wellknown,” as shown in the following command:

`elconfig set -lecs 47.0079.00.000000.0000.0000.0000.00a03e000001.00 -unit` *`unit_number`*

Checking that the adapter is communicating with the LECS

Knowing the LECS address, the system can now determine all existing ELANs, as well as the ATM address of the LES. However, the LAN Type remains unknown until you configure the adapter to join the ELAN, which is discussed in "Configuring the ATM adapter for an ELAN" on page 174.

To check the state of all adapters in the filer, use the `elconfig show` command.

Command syntax: The `elconfig show` command has the following syntax:

```
elconfig show -all
```

Example: The abbreviated sample output (showing ELANs available through the LECS for Adapter 2 only) from this command is as follows:

```
ELANs on Adapter 2

==================

LECS (current): 47.0079.00.000000.0000.0000.0000.00a03e000001.00

   ELAN           LAN Type    LES ATM Address

   ====           ========    ===============

   eighteenKMTU   Unknown
47.0005.80.ffe100.0000.f21a.4d19.0020481a4d19.14

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

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

Interpreting the output: Verify that each of the ELANs on the network appears in the output, and that the LAN Type and LES ATM addresses appear as expected. There should be a separate line for each configured ELAN on the network, grouped and arranged for each ATM adapter that is installed in the filer.

Note

If a separate line fails to appear for each configured ELAN, grouped by adapter, you need to go back to the ATM switch and make sure that each of the ELANs is properly configured. The documentation that came with your switch should help you with this.

An arrow appearing to the left of an ELAN signifies that the ATM adapter has been configured to operate on that ELAN. The lack of an arrow along with a LAN Type of "Unknown" signifies that the adapter has not yet joined the ELAN. Information about configuring adapters to operate on an ELAN is in "Configuring the ATM adapter for an ELAN" on page 174.

Check that the LES ATM address appears valid. If this is the case, you know that the ATM adapter in the filer is communicating properly with the switch. If the LES ATM address is all zeros, it might mean that the cable connection is not working, or that something else is improperly configured at the switch.

Finally, ensure that the LECS ATM address matches the address that you specified earlier in "Setting the LECS address" on page 171.

Configuring the ATM adapter for an ELAN

About configuring the ATM adapter for an ELAN

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

Before doing so, however, you must configure the LECS ATM address for each adapter, as described in "Configuring the LECS address" on page 171. In most cases, the LECS is set to the well-known address, although this might be different for your network. Also, the ELAN you specify must already have been configured at the switch.

Procedure for configuring the ATM adapter

To configure the LECS address, complete the following steps.

Step	Action
1	Add an ELAN to the adapter.
2	Configure the emulated Ethernet interface.
3	Add additional ELANs to the adapter, if required.
4	Correct ELAN addition mistakes, if necessary.

Adding an ELAN to the adapter

To add the ELAN to the adapter, use the `elconfig add` command.

Command syntax: The `elconfig add` command has the following syntax:

```
elconfig add ELAN -if interface -type ethernet -unit unit_number
```

ELAN is the Emulated LAN that you want the adapter to join.

interface is the logical network interface.

unit_number is the unit designator for the ATM adapter in the filer.

Example: The following command adds the adapter Unit 2 to the nineKMTU ELAN of type Ethernet, using interface *el1*:

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

Interpreting the output: This specific example assumes that the nineKMTU ELAN already exists on the switch. Also, if there is only a single ATM adapter in the filer, you don't need to specify the unit number in the command. The `elconfig` command sets it automatically.

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

Note

To add an ELAN to an adapter using FilerView, enter the `elconfig` command using FilerView's command-line interface.

Configuring the logical Ethernet interface

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

To configure a network interface, use the `ifconfig` command.

Command syntax: The `ifconfig` command has the following syntax:

`ifconfig` *`interface address`* `netmask` *`mask`* `up`

interface is the name of the logical network interface.

address is the IP address associated with the interface.

mask is a dot notation Internet address that is selected according to the class of the IP address.

For more information about the netmask parameter and the `ifconfig` command, refer to the *Command Reference*.

Example: The following command configures the *el0* logical network interface, assigning a corresponding IP address and netmask:

```
ifconfig el0 172.20.12.19 netmask 255.255.252.0 up
```

Adding more than one ELAN to an adapter

Adding more than one ELAN to an ATM adapter enables the filer to communicate over more ELANs than the number of physical adapters installed in the filer. To do so, simply repeat "Adding an ELAN to the adapter" and "Configuring the logical Ethernet interface" for each ELAN you want to add to the adapter.

Correcting ELAN addition mistakes

If you make a mistake using the `elconfig add` command, you can delete an ELAN from an adapter using the `elconfig delete` command.

Command syntax: The `elconfig delete` command has the following syntax:

```
elconfig delete ELAN -unit unit_number
```

ELAN is the Emulated LAN that you want to delete from the adapter.

unit_number is the unit designator for the ATM adapter in the filer.

Example: The following command deletes the nineKMTU ELAN for the adapter designated by unit 2:

```
elconfig delete nineKMTU -unit 2
```

Correcting ELAN addition mistakes after an IP address is assigned

If an IP address has already been assigned to a logical Ethernet interface on the adapter using `ifconfig`, you should see a message stating “elconfig ELAN in use” displayed as output from the `elconfig delete` command.

In this case, you must mark the interface as “down” by performing the `ifconfig` command before deleting the ELAN from the adapter:

Command syntax: The `ifconfig` command has the following syntax:

```
ifconfig interface down
```

interface is the name of the logical network interface.

Example: The following command marks the *el1* logical network interface as “down”:

```
ifconfig el1 down
```

Checking and completing the ELAN configuration

Procedure for checking the configuration

To check and complete the configuration, complete the following steps.

Step	Action
1	Verify the communication link.
2	Check the configuration settings.
3	Check the other elements of the ELAN.
4	Modify the load balancing setting, if desired.
5	Save the configuration in the */etc/rc* and */etc/hosts* files.

Verifying the communication link using ping

After you add an ELAN to an adapter and configure the interface, check to ensure that the filer can communicate with other clients through the ELAN. The easiest way to do this is to ping another LEC on the ELAN to ensure that information is traveling out through the ATM adapter and back again.

To ping a LEC, or any other client, use the `ping` command.

Command syntax: The `ping` command has the following syntax:

```
ping host
```

host is the IP address of the computer to which you want to send an ICMP ECHO_REQUEST datagram.

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

```
ping 204.125.14.45
```

If the host responds, `ping` prints "*host* is alive." Otherwise, `ping` resends the ECHO_REQUEST once a second. If the host does not respond after 20 seconds, `ping` prints "no answer from *host*."

Note

From a diagnostics point of view, it's often useful to have a LEC configured on the switch.

Checking the configuration settings

Double-check the state of the adapters in the filer to ensure that the configuration is as you expect.

To check the configuration setting, use the `elconfig show` command.

Syntax: The `elconfig show` command has the following syntax:

```
elconfig show -all
```

Example: The abbreviated sample output (showing ELANs on Adapter 2 only) from this command is as follows:

```
ELANs on Adapter 2

==================

LECS (current): 47.0079.00.000000.0000.0000.0000.00a03e000001.00

   ELAN           LAN Type    LES ATM Address

   ====           ========    ===============

   eighteenKMTU   Ethernet
47.0005.80.ffe100.0000.f21a.4d19.0020481a4d19.14

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

=> nineKMTU       Ethernet
47.0005.80.ffe100.0000.f21a.4d19.0020481a4d19.11
```

Interpreting the output: Verify the following for newly configured ATM adapters:

- There is an arrow preceding the ELAN name, which indicates that the adapter has joined the ELAN.
- The LAN Type is Ethernet.
- The LES ATM Address is a valid ATM address. If the address is all zeros, it indicates a configuration problem at the switch.

Checking the other elements of the ELAN

You should also check to ensure that the other elements of the ELAN are configured and operating as expected. For example, you should verify that the LES, LECS, and BUS ATM addresses are all valid, and that the state of the ELAN is operational.

In addition, you should check the MTU of each ELAN to verify that the setting matches the configuration created on the switch, by using the `elconfig show` command.

Syntax: The `elconfig show` command has the following syntax:

```
elconfig show -configured
```

Example: The abbreviated sample output (showing information related to the eighteenKMTU ELAN on adapter 2 only) from this command is as follows:

```
ELAN Name                   : eighteenKMTU
Interface                   : el1
Configured Unit             : 2
MAC Address                 : 00:20:48:40:1c:95
LEC Address                 :
47.0005.80.ffe100.0000.f21a.4d19.002048401c95.00
LECS Address                :
47.0079.00.000000.0000.0000.0000.00a03e000001.00
  Configuration Direct VCC :   unit=1 vpi/vci=0/17
LES Address                 :
47.0079.00.000000.0000.0000.0000.00a03e000001.00
  Control Direct VCC       :   unit=- vpi/vci=-/-
  Control Distribute VCC   :   unit=- vpi/vci=-/-
LAN Type                    : Ethernet/IEEE 802.3
Maximum Frame Size          : 18190
State                       : operational
MPOA                        : disabled
LECID                       : 8
```

Modifying load balancing [optional]

To modify load balancing groups, use the `uniconfig set failover` command.

Syntax: The `uniconfig set failover` command has the following syntax:

```
uniconfig set failover [ -state off | static | dynamic ] [-group
unit ... ]
```

`-state static` puts UNI load balancing in a static mode.

`-group unit...` specifies the load balancing group membership.

`-state off` disables load balancing and failover.

The option `-state dynamic` is the default.

Example: The following command disables load balancing and failover:

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

The arguments to the *-group* parameter specify the ATM adapters (units) that should be logically assigned to a load balancing and failover group. If you specify a unit that has already been assigned to another group, the unit is automatically removed from the original group before being assigned to the new group.

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

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

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

Note

By default, all ATM adapters are dynamically assigned to load balancing groups. In this case, no configuration is necessary to activate the load balancing and failover features.

Saving the ATM configuration information

After you successfully configure each of the ATM adapters in the filer and verify their proper operation, you should do the following:

- Save the commands in the */etc/rc* file for automatic execution at boot time.
- Save the host and IP address information in the */etc/hosts* file for each logical ELAN interface.

This saves you from having to reconfigure the adapters manually each time the filer is restarted.

Saving the configuration in the /etc/rc file

Save the commands in the */etc/rc* file for automatic execution at boot time. To do this, mount the root file system and edit the file using a text editor such as vi.

Example: The following is a sample portion of an */etc/rc* file containing configuration commands for three ATM adapters:

```
# unit 1
elconfig set -lecs -wellknown -unit 1
elconfig add default -if el0 -type ethernet -unit 1
ifconfig el0 172.20.12.19 netmask 255.255.252.0 up
```

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

Note that in this example, both adapter 2 and adapter 3 are on the nineKMTU ELAN, and that adapter 3 is on two ELANs.

Configuring the /etc/hosts file

Save the host and IP address information of the ELAN configuration in the */etc/hosts* files.

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

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

If you have a filer that has only ATM adapters, the */etc/hosts* file must contain an entry for the filer's host name. For example, the first line in the example might be as follows:

```
172.20.12.19        myfiler myfiler-el0
```

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

Configuring additional ATM adapters

Repeat the complete instructions in this chapter to configure additional ATM adapters in the filer.

FORE/IP over SPANS

Types of ATM connections

ATM supports two types of SPANS connections: SVCs (switched virtual circuits) and PVCs (Permanent Virtual Circuits).

Use of SVCs

The filer dynamically assigns SVCs when interoperating with ATM hosts and with switches that support the FORE IP/SPANS protocols.

Use of PVCs

You use ATM PVCs to interoperate with ATM hosts and with switches that do not support FORE IP/SPANS. For example, if you are not using a FORE systems switch, PVCs can connect FORE equipment at each end through non-FORE switches.

Note

Data ONTAP currently supports only FORE/IP 5.3. The following differences from previous versions are significant:

- The name of a SPANS interface in Data ONTAP 5.3 is fa*n* (where n is a number); in previous versions of data ONTAP it was a*n*. This means that you must change the interface names in your */etc/rc* and */etc/hosts* files.
- In Data ONTAP 5.3, you use the new command `atmconfig` to configure PVCs.

For further information about the differences between FORE/IP 5.3 and older versions, consult FORE documentation.

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

For each destination that needs to establish a PVC with the filer, you must establish an outgoing and an incoming PVC in at least three places:

- On the filer
- On the destination ATM host
- On all interconnecting ATM switches

SPANS interfaces

For each physical ATM interface, the filer creates a FORE/IP interface, called *fa*, at boot time. The *fa* interface supports FORE/IP on top of SPANS signaling. FORE/IP allows communication as follows:

- Using AAL4 or AAL5 ATM adaptation layer types with no encapsulation
- Using a broadcast ARP (Address Resolution Protocol) for SPANS address resolution
- Using direct communication of all hosts on a physical ATM network without the use of IP routers

Note

NetApp doesn't support FORE/IP load balancing or failover options.

Establishing an outgoing FORE/IP PVC on the filer

To establish an outgoing FORE/IP PVC on the filer, use the following filer command:

`atmarp -s` *`hostname`* `pvc` *`iface vpi vci`* `[`*`aal`* `[`*`encap`* `[`*`addr`*`]]]`

Variable	Description
`hostname`	Name or IP address of the remote host.
`iface`	Name of the ATM interface This is usually fa*n*, where *n* is a number.
`vpi`	VPI (virtual path identifier); this must be 0.
`vci`	VCI (virtual channel identifier); this number must have the following properties: ◆ It must not be in use on your filer. ◆ It must be less than 1,024. ◆ It must obey the limits of the destination host and interconnecting devices.
`aal`	ATM adaptation layer type. It must be 4 or 5, and should be the AAL type supported by the destination host, which is typically 5. The default is 5.

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

Establishing an incoming FORE/IP PVC on the filer

To establish an incoming FORE/IP PVC on the filer, use the following command at the filer:

`atmarp -l` *`pvc iface vpi vci`* [*`aal`* [*`encap`*]]

Use the variables as described in the previous section.

Establishing a FORE/IP PVC on the remote ATM host

Set up a FORE/IP PVC on the remote ATM host according to the documentation for that host.

Establishing a FORE/IP PVC on interconnecting ATM switches

On the interconnecting ATM switches, assign virtual channels corresponding to the VPI and VCI entries made on the filer and the remote ATM host according to the documentation for those switches.

Displaying information about a FORE/IP PVC

To display information about all FORE/IP PVCs and other interfaces on a host, use the following command:

`atmarp` *`hostname`*

hostname is the name of a host.

To display information about all the current FORE/IP PVCs, use the following command:

```
atmarp -a
```

A display similar to the following appears:

```
atmarp -a

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

Deleting a FORE/IP PVC

You can delete outgoing and incoming FORE/IP PVCs.

Deleting an outgoing FORE/IP PVC: To delete an outgoing FORE/IP PVC entry, use the following command:

```
atmarp -d hostname
```

hostname is the name or IP address of the remote host.

Deleting an incoming FORE/IP PVC: To delete an incoming FORE/IP PVC for a remote host, use the following command:

`atmarp -x` *`iface vpi vci`*

iface is the name of the interface

vpi and *vci* have the values of the VPI and VCI of the FORE/IP PVC to be deleted for the specified interface.

File Sharing between NFS and CIFS Users 6

About this chapter

About file sharing

This chapter describes how the filer works with NFS and CIFS clients simultaneously. Because these clients interact with a file server differently, you need to understand how the read and write operations performed by one client affect the operations performed by the other client.

What this chapter discusses

This chapter discusses the following topics.

Read...	To learn about...
"File-locking interactions" on page 191	CIFS and NFS file locks.
"Managing symbolic links for CIFS access" on page 192	How the filer supports CIFS access to symbolic links.
"NFS and CIFS use of the read-only bit" on page 195	How NFS and CIFS use the per-file read-only bit.
"Naming files used by both NFS and CIFS" on page 197	How the filer manages names of files used by both NFS and CIFS users.
"Languages and character sets" on page 199	How to select a character set and how to control conversion to Unicode directories.
"Directory formats" on page 210	How to decide when to change directory formats.
"CIFS file name case" on page 216	How to manage CIFS file name case conversion.
"Directory conversion time" on page 219	How to manage how long it takes to convert directories to Unicode format.
"How to manage UNIX access to NTFS files" on page 221	How to control UID-SID mapping and the WAFL credential cache.

What this chapter does not discuss

This chapter does not discuss

- The differences between procedures that take place on network clients

  For example, the way in which users gain access to files and directories depends on the NFS and CIFS protocols: NFS users *mount* the directories that were exported, and CIFS users *share* directories to which they were given access in the Access Control List (ACL).

  For more information about these procedures, refer to the documentation for the clients' operating systems.

- Filer administration procedures for individual file-sharing protocols

  Procedures such as rebooting the filer, recovering files from snapshots, and backing up the filer are different depending on whether there are CIFS clients at your site. The procedures are described in the applicable chapters of this guide.

File-locking interactions

About this section

This section describes in general what happens when a client program using one protocol tries to read or write a file that is currently used by a client program using a different protocol in an environment consisting of CIFS, (PC)NFS, and NFS users. For details about locking in a particular protocol, consult the documentation for that protocol.

Types of clients

There are CIFS clients and two kinds of NFS clients: UNIX-based NFS clients and (PC)NFS clients. The (PC)NFS clients and CIFS clients behave in the same general way.

Types of locks

NFS locks are advisory, while CIFS locks are mandatory. CIFS applications depend on locking to behave properly. Because NFS locks are only advisory, file-manipulation operations, such as `rm`, `rmdir`, and `mv`, by a UNIX-based NFS client on a file opened by a (PC)NFS application can cause the application to crash.

Reads by UNIX-based NFS clients always succeed.

A deny-write operation causes a file-manipulation operation, such as `rm`, `rmdir`, and `mv`, on a CIFS-accessed file by a UNIX-based NFS client to fail.

Byte-range locks work on portions of a file. Byte-range operations other than reads by a UNIX-based NFS client fail if the attempted operation is forbidden by the lock. As is appropriate for NFS, a UNIX-based NFS application might be forbidden to access a byte-range that is locked by CIFS.

Note

There is one exception to the enforcement of locks set by CIFS clients on the filer. When the filer runs the `dump` command, it ignores the file lock set by a CIFS client that prevents read access to parts of a file or the entire file. Ignoring the “read” file lock allows the filer to back up all files.

Managing symbolic links for CIFS access

About symbolic links

CIFS clients can follow symbolic links, which are created by NFS clients. A symbolic link is a special file that points to another file or directory. A symbolic link is, in some respects, like a shortcut in the Windows environment.

There are two kinds of symbolic links: absolute and relative:

- Absolute symbolic links begin with a slash (/) and are treated as a path relative to the root of the file system.
- Relative symbolic links begin with a character other than a slash (/) and are treated as a path relative to the parent directory of the symbolic link.

Controlling access to symbolic links

You can control CIFS access to symbolic links in three ways:

- Enabling or disabling symbolic links
- Redirecting absolute symbolic links
- Preventing or allowing the following of symbolic links that can refer to a directory higher in the same tree

Enabling symbolic links

When the symbolic links for CIFS feature is enabled, which is the default setting, if the object being accessed by a CIFS client is an absolute or relative symbolic link, the filer follows the link under the following conditions:

- The ultimate target is in the same share as the symbolic link.
- A symbolic link encountered in any path component other than the final one is always followed.
- The final component of a symbolic link is followed only if the operation is to open an existing file.
- Other operations, such as deleting and renaming, result in deleting or renaming the symbolic link itself rather than the target of the symbolic link.

CIFS client applications often perform operations such as writing to a temporary file, renaming the original file to a backup name, then renaming the temporary file to the original name. Therefore, take care in using symbolic links whose ultimate target is a file, as opposed to a directory. If the original file were targeted directly by a symbolic link, this sequence of operations would have the result—unintended by the application—of the file being stored in the directory where the

symbolic link was, and the renamed symbolic link pointing at the original file rather than to the updated file. For symbolic links to directories, this type of situation does not arise.

Because many PC applications work as described previously, if there are symbolic links that point to files, a PC could encounter such symbolic links. It is best to disable symbolic links for CIFS when there are symbolic links that point to files.

If you expect many files to be changed by applications that update files as described, you might want to disable symbolic links for CIFS.

How to enable and disable symbolic links

You enable and disable symbolic links with the `cifs.symlinks.enable` option. The option is On by default.

To disable symbolic links for CIFS, use

```
options cifs.symlinks.enable off
```

To reenable symbolic links for CIFS, use

```
options cifs.symlinks.enable on
```

How to redirect absolute symbolic links

In a UNIX environment, the NFS client interprets the file system location represented by an absolute symbolic link. The CIFS client cannot do this. In a CIFS environment, the filer enables you to redirect absolute symbolic links on the filer.

For example, you might want to redirect symbolic links pointing at the */u/users/charlie* directory to the */home/charlie* directory on the filer. You do so by specifying symbolic link redirection mappings in a text file named */etc/symlink.translations*.

The format of the */etc/symlink.translations* file is

```
Map link target
```

where both *link* and *target* are absolute symbolic link path names.

For example, the entry

```
Map /u/users/charlie/* /home/charlie/*
```

makes symbolic links pointing at the */u/users/charlie* directory point to the */home/charlie* directory.

How to prevent symbolic link cycling

You can create directory structures that are cyclic by creating a symbolic link that refers to a directory higher in the same tree, through use of a symbolic link having a “dot” or “dot-dot” component. Therefore, a simple recursive descent of the tree goes deeper and deeper until the maximum path length is reached. At that point an error is returned. For example, if you used Windows Explorer to search for files in such a cyclic directory, the same files show up repeatedly.

The `cifs.symlinks.cycleguard` option controls whether symbolic links that might include a directory higher in the same tree are followed.

To eliminate the possibility of cyclic directory structures, make sure that the `cifs.symlinks.cycleguard` option is On, which is the default, with the following command:

```
options cifs.symlinks.cycleguard on
```

If you use symbolic links having dot or dot-dot components and want the filer to follow the links, set the `cifs.symlinks.cycleguard` option to Off with the following command:

```
options cifs.symlinks.cycleguard off
```

When you list the contents of a directory, symbolic links that are valid references to files or directories are listed as if the target of the symbolic link existed in the directory. If the symbolic link cannot be expanded, it still looks like a file in a directory listing; however, any attempt by an application to open the link results in an access error.

NFS and CIFS use of the read-only bit

About read-only bits

The filer, along with MS-DOS and Windows, supports a per-file read-only bit that reflects whether a file is writable or read-only. This bit applies only to files and not to directories. NFS has no protocol operations that know about the per-file read-only bit. However, some software, when used both by NFS clients on UNIX systems and by CIFS clients on Windows systems, requires that the read-only bit reflects whether the file is writable.

How NFS treats the read-only bit

The following list describes how NFS treats the read-only bit:

- Any file with the read-only bit turned on is treated, for all NFS operations, as if it had no write permission bits turned on.
- If a file has at least one write permission bit turned on and an NFS client turns off all write permission bits, the filer turns on the read-only bit for that file. As described in the preceding paragraph, a file with the read-only bit turned on appears to an NFS client not to have any write permission bits turned on.
- If a file has no write permission bits turned on and an NFS client turns on any write permission bit, the filer turns off the read-only bit for that file.
- If a file's read-only bit is turned on and an NFS client attempts to find out the permission bits for the file, the actual permission bits for the file are not sent to the NFS client; instead, the filer sends the permission bits to the NFS client with the write permission bits masked off.
- If a file's read-only bit is turned on and a CIFS client turns the read-only bit off, the filer turns on the owner's write permission bit for the file.
- Files with the read-only bit turned on are writable only by the superuser.

How the filer tracks the NFS or CIFS client read-only bit

Whenever the read-only bit is turned on by a client, even if it was already on before the client did so, the filer tracks whether the client that turned on the bit was an NFS or CIFS client, as follows:

- If the bit was turned on by a CIFS client, renaming the file is not allowed. This is because file systems on MS-DOS and Windows systems do not allow renaming a file whose read-only bit has been set. NFS deletes follow NFS conventions. That is, deletes are allowed if the user has write permission in the parent directory.

- If the bit was turned on by an NFS client that turned off all write permission bits, removing or renaming the file is allowed if the user has sufficient permission to do so. This is because file systems on UNIX systems allow removing or renaming a file that has no write permission bits set.
- If the filer is using UNIX-style security, CIFS clients are also allowed to delete a file with the read-only bit set. This is required for compatibility with standard UNIX source control programs, such as RCS.
- If the filer is configured with PC-style security, the read-only bit is enforced.

Naming files used by both NFS and CIFS

About file naming conventions

File naming conventions depend on both the network clients' operating systems and the file-sharing protocols. For example, file names are case-sensitive for clients running UNIX and are case-insensitive, but case-preserving, for clients running Windows operating systems.

Maximum length of file names

On the filer, the maximum length of a file name is 255 characters for NFS clients and CIFS clients that support the PC's long file name format. Some CIFS clients, such as MS-DOS and Windows 3.*x* clients, support only file names in the 8.3 format (8 characters for the file name and 3 characters for the file name extension). In any directory that has access from a CIFS client, the filer creates and maintains two names: the original long name and an additional short name in 8.3 format.

How the filer generates short 8.3 file names

The filer generates an 8.3 file name as follows:

1. It truncates the file name to six characters.
2. It appends a tilde (~) and a number or letter to the name. If it runs out of letters and numbers because there are too many similar names, it creates a unique file name that bears no relation to the original file name.
3. It truncates the file name extension to three characters.

Note

The number or letter appended to the short name ensures that the file name is unique. It is not for showing the order of file creation.

For example, if an NFS client creates a file named *specifications.html*, the short name created by the filer is *specif~0.htm*. If this short name already exists, the filer uses a different number at the end of the file name. For example, if the UNIX client creates another file named *specifications_new.html*, the short version of *specifications_new.html* is *specif~1.htm*.

Which clients support short file names

The short names appear on clients that support only the 8.3 format. The short names are not visible to NFS clients. On Windows 95 and Windows NT clients, you can choose to display the short name or the long name by using File Properties.

Note

Under some circumstances, an application running on a client that uses names in 8.3 format can “lose” the file’s original long-format name. This can occur as a consequence of the way an application saves a file that it has edited. Some applications rename the original file, then save the edited file as if it were newly created. The filer thus receives instructions to delete the original file and create a new one. When the client supports only 8.3 names, this new name no longer has an equivalent in long format.

Legal characters used in file names

The characters that you can use in file names depend on the client operating systems. Because restrictions on legal characters vary from one operating system to another, refer to the documentation for your client’s operating system for more information about prohibited characters.

When you name a file to be shared by users on different operating systems, it helps to use only characters that are common to both. For example, if you use UNIX to create a file and use a colon (:) as its file name, an MS-DOS user sees the name displayed as ~0 because the colon is an illegal character in an MS-DOS file name.

Case-sensitivity in file names

Uppercase and lowercase characters are significant to NFS clients but not to CIFS clients. For example, if a file named *specifications* already exists, an NFS user can still save another file under the name *Specifications*, but a CIFS user is instructed by the application to choose another file name. This section describes how both NFS and CIFS users can use file names that differ only by case.

When a client creates a file name, the filer preserves the case. For example, if a CIFS client creates *Spec.txt*, the file name is displayed by both CIFS and NFS clients as *Spec.txt*. If an NFS user later creates a file named *spec.txt*, NFS and CIFS clients see the file names as follows:

- On NFS clients, one file is displayed as *Spec.txt* and the other is displayed as *spec.txt*. That is, the file names are displayed in the same way as they were created.
- On CIFS clients, even those that support long names, one file is displayed as *Spec.txt* and the other is displayed as *Spec~0.txt*.

Languages and character sets

File names, languages, and character sets

File names in languages other than English can use characters that have diacritic marks, or accents, or can use characters that are not even in the Roman alphabet. UNIX and Windows systems can create file names with no restrictions, provided that they do not violate the naming conventions of the operating system that they were created in. The language you select affects multiprotocol behavior and the code page for WIndows 98 and older clients.

File names use character sets

To make sure that file names stored on the filer are usable by both UNIX and Windows applications, you must choose a character set that contains the characters in the language that your clients use. To do so, you choose a language, and the filer uses a character set that is appropriate to the language.

Every volume has a language

Every volume uses a language, and therefore a character set, that you specify for file names. The root volume determines the code page for PCs and the console character set.

Language selection

The language you specify controls the name translation between UNIX and Windows names. If you installed a new filer, you selected a default language when you set up the filer. If you upgraded to Data ONTAP 5.3, you must use the `vol lang` command to set the language.

Note

If a filer is licensed for CIFS, you must set a language for every volume on the filer.

What a language applies to

The filer uses a character set appropriate to a given language. The language you select determines which character set the filer uses for the following names:

- User names
- Share names
- System and domain names

Note

The following must be in ASCII:

- Qtree names
- Snapshot names
- Volume names

Kinds of character sets supported

The filer supports the following types of character sets. The following table shows what protocols and operating systems use a particular character set. You use this information to determine what directory format you use.

Character set type	Description	Used by
ASCII	A 7-bit character set used by most computers. Does not allow letters with accents; that is, diacritics.	NFS, console and log files, qtree names, snapshot names, and volume names.
Unicode	A 16-bit character encoding system. It includes all major languages.	WIN 32 applications and file names in the following systems: ◆ Windows NT ◆ Windows 98 ◆ Windows 95
UNIX	A variety of character sets used by UNIX. Can include single-byte and multibyte characters.	NFS, console, and log messages
UTF-8	An ASCII-compatible multibyte Unicode encoding.	Some Solaris clients

Character set type	Description	Used by
OEM	Single or multiple encodings for older clients	◆ DOS ◆ Windows 98 (except for file names) ◆ Windows 95 (except for file names) ◆ NT system and domain name

Languages supported

The filer supports the languages shown in the following table. The language code for each language appears next to the language. To override the normal UNIX character set so that a UNIX system uses UTF-8 code, add .UTF-8 to the language code.

You use a language code to specify a language with the `vol lang` command, as described in "Setting the language of a volume" on page 207.

Language	Language Code
Danish	da
Dutch	nl
English	en
English (US)	en_US
Finnish	fi
French	fr
German	de
Hebrew	he
Italian	it
Japanese euc-j	ja
Japanese PCK(sjis)	ja_JP.PCK
Norwegian	no
Portuguese	pt
POSIX (ASCII only for DOS and ISO-Latin1 for NFS. Equivalent to the Sun locale.)	C
Spanish	es
Swedish	sv

How to choose a language

The following flow chart shows how to choose a language.

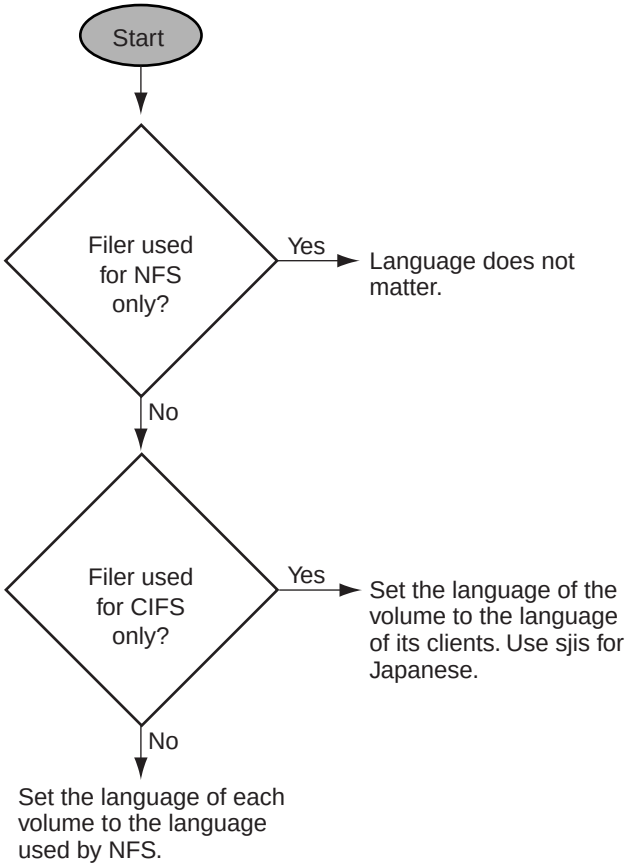

Language procedures

You can perform the following operations to query and specify the language that a volume uses.

Operation	Section
Show what languages the filer supports.	"Displaying a list of supported languages" on page 205
Specify a language for the console.	"Setting the console encoding" on page 206
Change the language, and therefore the character set, that a volume uses for file names.	"Setting the language of a volume" on page 207
Create a volume that uses a specified language other than the default.	"Creating a volume that uses a specified language" on page 208
Show each volume with the language that it uses.	"Displaying which volume uses which language" on page 209

Displaying a list of supported languages

Description

This procedure displays a list of languages for which the filer supports character sets. You can perform this procedure at any time to determine what language to specify for a volume if you do not already know what languages are supported.

Step

To display a list of the languages for which the filer supports character sets, complete this step.

Step	Action
1	Enter the following command: `vol lang`

Setting the console encoding

Description

Use this procedure to set how the console displays non-ASCII information. You would change the encoding if your console does not accept or display characters properly using the current encoding. The default encoding is nfs.

Step

To set the encoding of the filer console, complete this step.

Step	Action
1	Enter the following command: `options console.encoding` ***`encoding`*** ***`encoding`*** is one of the following encoding styles. You can use both NFS extended (greater than 0x7F) and SGML characters for input. ◆ nfs—The NFS character set of the root volume. You can use both NFS extended (greater than 0x7F) and SGML characters for input. ◆ sgml—SGML character format. You can use both NFS extended (greater than 0x7F) and SGML characters for input. ◆ utf8—UTF-8 character sets. For input, any character greater than 0x7F is treated as the beginning of a UTF-8 encoding.

Setting the language of a volume

Description

This procedure sets the language that a volume uses to store file names. You complete this procedure when you want the files names stored in a volume to use a different language than the default language. You should complete this procedure before any files are created in the volume so that all file names use the same language.

After you complete the procedure, the filer stores file names in the selected volume using a character set that best fits the language you select.

Prerequisites

To complete this procedure, you need the language code of the language you want to use. You can get the language code from the table in "Languages supported" on page 202.

Caution

Changing the language after the volume contains files can cause some of the NFS encodings to be invalid and is not recommended.

Step

To set the language and character set that a volume uses to store file names, complete the following step.

Step	Action
1	Enter the following command: `vol lang` ***`volume langcode`*** ***`volume`*** is the name of the volume whose language you want to change. ***`langcode`*** is the code for the language you want the volume to use.

Creating a volume that uses a specified language

Description

This procedure creates a volume with a specified language. You use this procedure when you are creating a volume but want to use a language different from the language of the root volume, which is the default language. You use this procedure any time you want to create a volume.

After you complete the procedure, the filer stores file names in the volume you created using a character set that best fits the language you selected.

Prerequisites

To complete this procedure, you need the language code of the language you want to use. You can get the language code from the table in "Languages supported" on page 202.

Caution

Changing the language of a volume after the volume contains files can cause some of the NFS encodings to be invalid and is not recommended.

Step

To set the language and character set that a volume uses to store file names, complete the following step.

Step	Action
1	Enter the following command: `vol create volume -l langcode` ***volume*** is the name of the volume whose language you want to change. ***langcode*** is the code for the language you want the volume to use.

Displaying which volume uses which language

Description

This procedure displays a list of volumes with the language each volume is configured to use. You perform this procedure when you want to know which volume uses which language. This is useful in matching clients with languages or deciding whether to create a volume to accommodate clients that use a language that you might not have a volume for.

Step

To display a list of volumes with the language each volume is configured to use, complete the following step.

Step	Action
1	Enter the following command: **`vol status -l`**

Sample output

Each row of the list displays the name of the volume, the language code, then the language.

```
Volume Language
vol0   ja (Japanese euc-j)
```

Directory formats

Three directory formats

The filer supports three directory formats. Each format supports a particular set of file names, as shown in the following table.

	Directory Format:		
	Pre-Data ONTAP 4.0	Data ONTAP 4.0	Unicode
Versions used in	Data ONTAP 3 and earlier	Data ONTAP 4.0 through 5.2	Data ONTAP 5.3
Types of file names supported	NFS file names only	Multiprotocol file names in ASCII, with limited non-ASCII support for CIFS names	File names that use non-ASCII characters

Directory sizes

Pre-Data ONTAP 4.0 and Data ONTAP 4.0 are slightly smaller than Unicode directories. If you are concerned about space on disks or in volumes, you might want to avoid Unicode directories.

Directory creation and formats

By default, the filer creates NFS directories in pre-4.0 format. When accessed by NFS, directories remain in pre-4.0 format, and the first CIFS access of a directory converts it to Unicode format.

CIFS directories are created in Unicode format.

Directory conversion

About conversion

The filer converts directories to Unicode whenever the directory is accessed by CIFS. The conversion process can take a few minutes if the directory contains many files.

Conversion to Unicode format

By default, directory conversion to Unicode format occurs as described in the following table.

Access from	Directory format	
	Pre-Data ONTAP 4.0	**Data ONTAP 4.0**
CIFS	Converted to Unicode	Converted to Unicode
NFS	No conversion	Converted to Unicode

Directory format conversion procedures

You can perform the following operations to control directory format conversion.

Operation	Section
Make directories into Unicode directories whenever they are accessed.	"Converting directories to Unicode format upon any access" on page 214
Force all DOS file names to lowercase.	"Forcing CIFS file names to lowercase" on page 217
Make all newly created directories Unicode directories.	"Directory conversion time" on page 219

Forcing Unicode default directory creation

Description

Use this procedure to have all directories created in Unicode format. This eliminates time lags due to conversion when a directory populated by UNIX users is then accessed by Windows users. It also allows new directories created by UNIX to be accessible to CIFS when they appear in a snapshot. You should do this immediately after installing Data ONTAP 5.3. After you complete this procedure, all directories that you create are in Unicode format, and their files are accessible to CIFS clients.

Step

To force all directories to be created in Unicode format, complete the following step.

Step	Action
1	Enter the following command: `options wafl.create_ucode on`

Returning to pre-Data ONTAP 4.0 default directory creation

Description

Use this procedure to return to the creation of directories in pre-Data ONTAP 4.0 format. You can do this at any time, perhaps to save space because pre-4.0 format uses slightly less space than Unicode format.

After you complete this procedure, the filer returns to creating all directories in pre-4.0 format when created by NFS, and only the first CIFS access converts a directory to Unicode format. CIFS still creates directories in Unicode as usual.

Caution

This procedure does not change the format of existing directories.

Step

To return to the default creation of pre-Data ONTAP 4.0 format directories, complete the following step.

Step	Action
1	Enter the following command: `options wafl.create_ucode off`

Converting directories to Unicode format upon any access

Description

Use this procedure to have the filer convert pre-4.0 directories to Unicode format when accessed from either CIFS or NFS, as opposed to only from CIFS. You perform this procedure to facilitate CIFS file-sharing. After you complete this procedure, pre-4.0 directories are converted to Unicode format even when accessed by NFS.

Step

To convert directories to Unicode format when they are accessed from either CIFS or NFS, complete the following step.

Step	Action
1	Enter the following command: `options wafl.convert_ucode on`

Returning to default Unicode format directory conversion

Description

Use this procedure to return to the original default directory conversion behavior. After you complete this procedure, access from CIFS causes conversion of pre-4.0 and 4.0 format directories to Unicode format, but accesses from NFS do not. You can perform this procedure at any time. You use this procedure to save a little space, because pre-4.0 format directories are smaller than Unicode directories.

Step

To return to the original default directory conversion behavior, complete this step.

Step	Action
1	Enter the following command: `options wafl.convert_ucode off`

CIFS file name case

Case preservation

By default, the filer preserves the case of CIFS file names.

Case conversion procedures

You can force all PC-created file names to be stored on the filer in lowercase or return to having the filer preserve case.

To force all CIFS file names to be stored on the filer in lowercase, follow the procedure in "Forcing CIFS file names to lowercase" on page 217. You might want to do this to avoid the following problems:

- UNIX lowercase names might contain uppercase characters when converted to CIFS names.
- Some CIFS clients changing the case of NFS file names on the filer, making them inaccessible from NFS.

To change back to having the filer preserve the case of CIFS file names, follow the procedure in "Preserving the case of CIFS file names" on page 218.

Forcing CIFS file names to lowercase

Description

This procedure enables you to force all CIFS file names to be stored on the filer in lowercase. You should do this if you have CIFS clients, so that they do not change the case of NFS file names on the filer that were originally lowercase.

If you perform this procedure, you should do so before Windows files appear on the filer.

After you complete this procedure, the filer stores all CIFS names in lowercase and does not preserve case.

Step

To force all CIFS file names created by CIFS to be lowercase, complete the following step.

Step	Action
1	Enter the following command: `options cifs.save_case off`

Preserving the case of CIFS file names

Description

This procedure enables you to return from forcing all file names to be stored in lowercase to the default behavior of preserving the case of CIFS file names. You perform this procedure when you want to return to the default behavior.

Caution

File names that were converted to lowercase are not changed.

Step

To return to preserving the case of CIFS file names, complete the following step.

Step	Action
1	Enter the following command: `options cifs.save_case on`

Directory conversion time

Directory conversion can take a lot of time

Although the conversion process to Unicode directory format is automatic, the initial conversion of a directory can take considerable time, especially if the directory contains a large number of files.

For example, converting a directory of 100,000 files takes the filer about 9 minutes on a NetApp F540. It is important to take these conversion times into consideration when deciding when to convert the directories. While a directory is being converted, the filer might not be able to perform any other file system or network operations until the conversion is complete.

When there is no need to convert

If there are portions of the directory tree that will never have CIFS access, there is no need to convert them. However, any future CIFS access to an unconverted directory immediately triggers its conversion.

How to speed up directory conversion

If you have a directory that contains more than 50,000 files, before triggering a conversion, you can use an NFS client to distribute files among a greater number of subdirectories. This speeds up the conversion process and avoids a possible crash.

Speeding up conversion time by renaming NFS directories

Description

Use this procedure to convert a directory to Unicode format if you have access to a Windows NT client connected to the filer. After you complete this procedure, you have a Unicode directory containing files that were in a non-Unicode directory, and its files are accessible to CIFS clients.

Step

To convert directories to Unicode format quickly, complete the following step.

Step	Action
1	Create a new CIFS directory from a Windows NT client on the same volume in the same qtree as the directory you want to convert.
2	With the NFS `mv` command, rename files from the directory you want to convert into the directory you just created.
3	Optionally remove the old directory.
4	Optionally rename the new directory.

How to manage UNIX access to NTFS files

UNIX users need Windows NT credentials to access NTFS files

If a user is using UNIX and tries to open a file with an ACL in a mixed or NTFS (Windows NT security style) qtree, the filer uses NTFS security semantics to determine whether the user has access to the file.

The filer does this by converting the UNIX UID (User ID) into a Windows NT credential, which is also known as a WAFL (Write Anywhere File Layout) credential. Windows uses the credential to verify that a user has access rights to the file. A UNIX user can have both a UNIX name and a Windows name.

As part of the process of creating a WAFL credential, the filer contacts an NT domain controller to look up a user's SID (Security ID) and groups.

WAFL credential caching

Windows verifies each request for access to files. Each access usually takes between 10 and 500 milliseconds, and can take longer because contacting a domain controller to create a WAFL credential is time-consuming. To reduce the time spent in contacting a domain controller, the filer can cache the WAFL credential in a credential cache. This cache is called the WAFL credential cache.

How to manage the WAFL credential cache

The default configuration

The default configuration stores each WAFL credential for 20 minutes. This is generally sufficient for most situations.

Two ways to manage the WAFL credential cache

You can manage the cache globally with options and specifically with the `wcc` command. This enables you to add entries to and remove entries from the WAFL credential cache and to troubleshoot file access problems that might be caused by mapping problems.

Note

If you make changes to the WAFL credential cache and an NFS client has cached information, it can take a noticeable amount of time for the changes to become visible.

Global cache management options

If the default configuration does not work well for your site, you can use an option to set how long each WAFL credential cache entry is valid. You can also trace CIFS logins as an aid to debugging mapping problems.

- To set how long each WAFL credential cache entry is valid, use the `wafl.wcc_minutes_valid` option, as described in "Setting how long each WAFL credential cache entry is valid" on page 225.
- To display information about every CIFS login attempt, use the CIFS login tracing feature, as described in "Toggling CIFS login tracing" on page 232.

When to use the wcc command

You use the `wcc` command to perform the following tasks:

- Add names to or remove names from the WAFL credential cache, as described in "Adding an entry to the WAFL credential cache" on page 226 and "Deleting entries from the WAFL credential cache" on page 227
- Monitor the WAFL credential cache by displaying statistics about it, as described in "Displaying WAFL credential cache statistics" on page 228.
- Troubleshoot file access and other problems by displaying what name mappings would be, as described in "Displaying a mapping result for a UNIX name" on page 230 and "Displaying a mapping result for a Windows name" on page 231.

The wcc command syntax

The `wcc` command has five basic types, each with its own function. The type is determined by the first option, as follows:

- `wcc -a` adds a name to the WAFL credential cache.
- `wcc -d` displays statistics about the WAFL credential cache.
- `wcc -s` displays what the mapping of a Windows NT name would be.
- `wcc -u` displays what the mapping of a UNIX name would be.
- `wcc -x` removes entries from the WAFL credential cache.

The wcc command options

Each `wcc` command type has a different set of options, as shown in the following table.

Type	Other options	Function
wcc -a	-u *uname* -i *ip-addr* [-v]	Adds the name to the WAFL credential cache. The `-v` option displays Windows NT groups.
wcc -d	[-v] ...	Displays the following statistics about the WAFL credential cache: ◆ Number of entries in the cache ◆ Age of the oldest entry ◆ Number of Administrator-privileged entries The `-v` option adds mappings for every user.
wcc -s	*ntname* [-i *ip-addr*] [-v]	Displays what the current mapping of the specified name would result in, but does not change the WAFL credential cache. The `-v` option displays numeric SIDs.
wcc -u	*uname* [-i *ip-addr*] [-v]	

Type	Other options	Function
wcc -x	[-f] [-v]	Removes all entries from the WAFL credential cache. The `-v` option displays how many entries have been removed.
	-i *ip-addr* [-v]	Removes all entries from the WAFL credential cache with the same IP address. The `-v` option displays how many entries have been removed.
	-s *ntname* [-i *ip-addr*] [-v]	Removes the entries with the given Windows NT name and optional qualifying IP address. If the Windows NT name is the name of a group, removes all members of that group from the WAFL credential cache. The `-v` option displays how many entries have been removed.
	-u *uname* [-i *ip-addr*] [-v]	Removes the entries with the given UNIX name and optional qualifying IP address. The `-v` option displays how many entries have been removed.

The `-v` option increases the level of detail of information. The kind of information varies with each command. You can have up to three instances of the -v option (-vvv) per command. Each repetition of the option increases the level of detail; three instances provide statistics that are only of interest to Network Appliance Technical Support.

Setting how long each WAFL credential cache entry is valid

Description

This procedure enables you to change how long each WAFL credential cache entry is valid. If you need to see security updates as they occur, you might want to use a smaller value than the default. However, access verifications are more frequent than with a greater value. This means that with a smaller value, users might experience slower performance.

Step

To change how long each WAFL credential cache entry is valid, complete this step.

Step	Action
1	Enter the following command: **`options wafl.wcc_minutes_valid`** ***`n`*** *n* is the number of minutes you want each entry to be valid. It can range from 1 through 20160. The default value is 20.

Adding an entry to the WAFL credential cache

Description

This procedure adds an entry to the WAFL credential cache. You can use this procedure in a script to load the WAFL credential cache at boot time with entries rather than wait for those entries to be created in the course of accessing a file. You can perform this procedure at any time.

Prerequisites

You must have the names and IP addresses of those you want to add to the WAFL credential cache.

Cautions

If you add more entries than the maximum number of entries allowed, the older entries are deleted.

Step

To add an entry to the WAFL credential cache, complete this step.

Step	Action
1	Enter the following command: `wcc -a -u` ***`uname`*** `-i` ***`ipaddress`*** ***`uname`*** is the UNIX name of a user. ***`ipaddress`*** is the IP address of the host that the user is on.

Deleting entries from the WAFL credential cache

Description

This procedure deletes entries from the WAFL credential cache. You do this to force the lookup of UIDs the next time they are used, but you don't want to wait until the entries time out automatically.

Prerequisites

You must have the name and optionally the IP address of the person or group you want to remove from the WAFL credential cache.

Caution

If the Windows NT name is the name of a group, this procedure removes all members of that group from the WAFL credential cache.

Step

To remove an entry from the WAFL credential cache, complete this step.

Step	Action
1	Enter the following command: `wcc -x name` *name* can be one of the following specifications: ❖ `-s` followed by a Windows user name or group name ❖ `-u` followed by a UNIX name You can further narrow the specification of a user by adding `-i`, followed by the IP address of the host that the user is on.

Displaying WAFL credential cache statistics

Description

This procedure displays WAFL credential cache statistics in detail. You can do this at any time to monitor the WAFL credential cache.

Step

To display statistics about the WAFL credential cache, complete this step.

Step	Action
1	Enter the following command: **`wcc -d`** You can get more detailed information by appending `-v` to the command line.

Sample output

The following sample output shows the output of statistics with the `-d` option.

```
wcc -d
mday (UID 10050) from 10.100.4.41 => NT-DOMAIN\mday*

Total WCC entries: 3; oldest is 127 sec.
Total Administrator-privileged entries: 1
 * indicates members of "BUILTIN\Administrators" group
```

The following sample output shows the output of statistics with the `-v` option used twice.

```
cc -dvv
mmm (UID 1321) from 10.100.4.41 => NT-DOMAIN\mmm
        ***************
        UNIX uid = 1321

        NT membership
                NT-DOMAIN\mmm
                NT-DOMAIN\Domain Users
                NT-DOMAIN\SU Users
                NT-DOMAIN\Installers
                NT-DOMAIN\tglob
                NT-DOMAIN\Engineering
                BUILTIN\Users
```

```
        User is also a member of Everyone, Network Users,
        Authenticated Users
        ***************
mday (UID 10050) from 10.100.4.41 => NT-DOMAIN\mday*
        ***************
        UNIX uid = 10050

        NT membership
                NT-DOMAIN\mday
                NT-DOMAIN\Domain Users
                NT-DOMAIN\Domain Admins
                NT-DOMAIN\SU Users
                NT-DOMAIN\Installers
                BUILTIN\Users
                BUILTIN\Administrators
        User is also a member of Everyone, Network Users,
        Authenticated Users
        ***************
hawleyr (UID 1129) from 10.100.4.41 => NT-DOMAIN\hawleyr
        ***************
        UNIX uid = 1129

        NT membership
                NT-DOMAIN\hawleyr
                NT-DOMAIN\Domain Users
                NT-DOMAIN\Installers
                NT-DOMAIN\SU Users
                BUILTIN\Users
        User is also a member of Everyone, Network Users,
        Authenticated Users
        ***************

Total WCC entries: 3; oldest is 156 sec.
Total Administrator-privileged entries: 1
 * indicates members of "BUILTIN\Administrators" group
```

Displaying a mapping result for a UNIX name

Description

This procedure displays what the current mapping of the specified UNIX name of a UNIX user would result in, but does not change the WAFL credential cache itself. You use this procedure if a UNIX user cannot access a file that the user should be able to access, and you suspect that mapping problems might be part of the problem.

Step

To display what the current mapping of a UNIX name would result in, but not change the WAFL credential cache, complete this step.

Step	Action
1	Enter the following command: **`wcc -u`** ***`uname`*** ***`uname`*** is the UNIX name of a user. ◆ You can further narrow the specification of the user by adding `-i`, followed by the IP address of the host that the user is on. ◆ You can get more detailed information by appending `-v` to the command line.

Sample output

The following example shows the mapping of a UNIX name.

```
wcc -u fuser001
(NT - UNIX) account name(s):  (NT-DOMAIN\fuser001 - fuser001)
        ***************
        UNIX uid = 1172

        NT membership
                NT-DOMAIN\fuser001
                NT-DOMAIN\Domain Users
                BUILTIN\Users
                TFILER\Test
        User is also a member of Everyone, Network Users,
        Authenticated Users
        ***************
```

Displaying a mapping result for a Windows name

Description

This procedure displays what the current mapping of the specified UNIX name of a Windows NT account would result in, but does not change the WAFL credential cache itself. You use this procedure if a Windows NT user cannot access a file that the user should be able to access, and you suspect that mapping problems might be part of the problem.

Step

To display what the current mapping of a Windows name would result in, but not change the WAFL credential cache, complete this step.

Step	Action
1	Enter the following command: **`wcc -s`** ***`uname`*** ***`uname`*** is a Wndows NT account. ◆ You can further narrow the specification of the user by adding `-i`, followed by the IP address of the host that the user is on. ◆ You can get more detailed information by appending `-v` to the command line.

Sample output

The following example shows the mapping of a Windows name.

```
wcc -s bluebottle
(NT - UNIX) account name(s):  (NT-DOMAIN\bluebottle - pcuser)
        ***************
        UNIX uid = 65534

        NT membership
                NT-DOMAIN\bluebottle
                NT-DOMAIN\Domain Users
                BUILTIN\Users
        User is also a member of Everyone, Network Users,
        Authenticated Users
        ***************
```

Toggling CIFS login tracing

Description

You can enable login information to be displayed after every CIFS login. You can do this any time you need to troubleshoot mapping problems. Login problems are often the cause of users being denied access to files they should have access to. CIFS login tracing is especially useful in producing verbose messages when a login attempt fails.

Caution

Turning on CIFS login tracing should be used carefully because every CIFS login is traced and results in console messages. Constant use of CIFS login tracing can result in many console and log messages.

Step to turn on CIFS login tracing

To turn on CIFS login tracing, complete this step.

Step	Action
1	Enter the following command: **`options cifs.trace_login on`**

Sample output

If user jdoe attempts to log in, messages like the following appear on the console:

```
Mon Jan  4 15:21:38 PST [CIFSAuthen]: Login attempt by NETAPP\jdoe
from SMITH-PC
Mon Jan  4 15:21:38 PST [CIFSAuthen]: User authenticated by DC
Mon Jan  4 15:21:38 PST [CIFSAuthen]: PC user name maps to UNIX user
smith
Mon Jan  4 15:21:38 PST [CIFSAuthen]: Unix user set to root by
wafl.nt_admin_priv_map_to_root
Mon Jan  4 15:21:38 PST [CIFSAuthen]: Login accepted
```

Step to turn off CIFS login tracing

To turn off CIFS login tracing, complete this step.

Step	Action
1	Enter the following command: `options cifs.trace_login off`

NFS Administration 7

NFS administration covers four areas

This chapter has four sections

This chapter is divided into four sections that contain the topics described in the following table.

Read...	To learn about...
"Section A: Managing NFS Exports" on page 236	Exporting filer volumes and directories to clients so that users can access them.
"Section B: Configuring a filer for (PC)NFS" on page 251	Setting the filer up to support (PC)NFS clients so that users on those clients can access filer volumes and directories.
"Section C: Configuring a filer for WebNFS" on page 258	Setting up the filer to serve Web pages to browsers that support the NFS transport.
"Section D: Displaying NFS statistics" on page 262	Displaying and interpreting NFS statistics.

Section A: Managing NFS Exports

Section contents

This section describes how to manage NFS exports. It includes the topics listed in the following table.

Topic	Description
"Introducing the /etc/exports file" on page 237	Describes the format for entries in the */etc/exports* file.
"Rules for exporting volumes and directories" on page 239	Describes rules you should follow when creating entries in the */etc/exports* file.
"Default /etc/exports entries" on page 242	Describes the default entries that `setup` places in the */etc/exports* file.
"Restricting access to volumes and directories" on page 243	Describes how to restrict access to exported volumes and directories, including ways to use the `-access`, `-root`, `-ro`, and `-rw` options with */etc/exports* file entries.
"The exportfs command" on page 245	Describes the syntax for the `exportfs` command and how to use the command to activate and deactivate exports.
"The /etc/netgroup file" on page 247	Describes how to define groups of hosts, subnets, and other groups to make entries in the */etc/exports* file easier to manage.
"Exporting to subnets" on page 249	Describes how to make entries in the */etc/exports* file more manageable by exporting volumes and directories to subnets rather than to individual clients.

Introducing the /etc/exports file

/etc/exports controls client access to directories

The */etc/exports* file controls how NFS clients access filer directories. You add entries to the */etc/exports* file for all the directories you want to export.

Format for /etc/exports entries

The format of the entries in the */etc/exports* file is as follows:

```
filer_directory_path export_specification
```

Filer directory path format

The filer directory path specifies which directory is made available to clients. The format is as follows:

`/vol/`*`volume_name/directory`*

Examples: The following lines show examples of filer directory paths:

```
/vol/vol0/home
/vol/users/local/jarnold
```

Export specification determines access privileges

The export specification specifies the privileges that clients have to mount and access the filer directory path. The format is as follows:

`-root=`*`list`*`,-access=`*`list`*`,-ro=`*`list`*`,-rw=`*`list`*

One keyword is required: The keywords `-root`, `-access`, `-ro`, and `-rw` are all optional; however, you must include at least one keyword in an export entry.

What the list variable represents: The *list* variable represents a list that includes one or more

- host names
- netgroup names
- subnets

You can combine elements in an entry: You can combine host names, netgroup names, and subnets in an entry, as shown in the following example:

```
-root=adminhost:administrators,-rw=blender:pcusers:123.45.67.0/24
```

Example 1: exporting default filer volume to administration host:

The following line exports the root directory of the default filer volume to the administration host with root privileges. The administration host can mount the directory, change permissions and ownerships, and create and delete directories and files.

```
/vol/vol0 -root=adminhost
```

Example 2: exporting home directory to administration host and clients: The following line exports the home directory from the default filer volume to the administration host and two clients. The administration host has root privileges and can mount the directory, change permissions and ownerships, and create and delete directories and files. The clients can mount the directory and create and delete directories.

```
/vol/vol0/home -root=adminhost,-rw=blender:mixer
```

Related information

The following table describes related information that you might want to read.

Read...	To learn about...
"Rules for exporting volumes and directories" on page 239	Rules for exporting volumes and directories.
"Default /etc/exports entries" on page 242	The default entries that `setup` places in the */etc/exports* file.
"Restricting access to volumes and directories" on page 243	Ways that you can restrict access to volumes and directories.
"The exportfs command" on page 245	How to activate the exports listed in the */etc/exports* file.

Rules for exporting volumes and directories

Export each volume separately

If the filer has multiple volumes, you must export each volume separately; you cannot export all volumes by specifying `/vol` as the exported directory.

Example: The following lines show how to export the volumes on a filer that has three volumes:

```
/vol/vol0 -root=toaster, -access=mixer:blender
/vol/users -root=toaster, -access=mixer:blender
/vol/builds -root=toaster, -access=mixer:blender
```

Nonexample: The following line shows an invalid entry for the */etc/exports* directory for a filer that has three volumes; */vol* cannot be used by itself as a path:

```
/vol -root=toaster, -access=blender:mixer
```

Filer must resolve host names

To export directories to hosts, the filer must be able to resolve host names into IP addresses. Host name resolution can take place

- through DNS name resolution
- by using the */etc/hosts* file on the filer root volume

Refer to "Host name resolution" on page 117 for more information about host-name-to-IP-address resolution.

Cannot restrict access by host

Neither the filer nor NFS provide a way to specify hosts that *cannot* mount the filer. To restrict access to exports, you must export volumes and directories in a manner that specifically includes those hosts that should be allowed to mount them. Exclusion occurs when a host is not specifically included in the list of hosts authorized to mount exported volumes and directories.

You can export ancestors and descendants

The filer permits directories that have exported ancestors to be exported. In many implementations of the UNIX operating system, you cannot export a directory that has an exported ancestor in the same file system.

Example: The following lines show exports that the filer allows:

```
/vol/vol0 -access=adminhost,-root=adminhost
/vol/vol0/home -access=blender:mixer
```

Nonexample: The following lines show exports that are not allowed on some UNIX systems:

```
/home -root=adminhost,-access=blender:mixer
/home/local -root=adminhost,-access=blender:mixer
```

Filer determines permissions by matching longest prefix

The filer uses the longest matching prefix in determining permissions.

In the preceding example

- A client mounting */vol/vol0/home/user1* gets permissions for */vol/vol0/home* because */vol/vol0/home* is the longest matching prefix.
- A client mounting */vol/vol0* gets `-access=adminhost, -root=adminhost` permissions.

Note

Because of the way the filer determines permissions, it makes little sense to give a client greater permissions at a higher level in the file system.

Example: The following lines show an */etc/exports* file that creates a security breech by enabling any host to mount the */vol/vol0* directory while restricting specific hosts to mounting the */vol/vol0/home* directory. In this example, any host can gain access to the */vol/vol0/home* directory by mounting the */vol/vol0* directory:

```
/vol/vol0
/vol/vol0/home -access=bashful:dopey:sleepy
```

Edit /etc/exports after changing volume names

If you rename a volume using the `vol rename` command

- Entries in the */etc/exports* file that refer to the volume become incorrect.
- The in-memory information about active exports gets updated automatically, and clients continue to access the exports without problems.
- The filer displays an error message when the `exportfs -a` command is entered, or when the filer is rebooted.
- Clients display the error message, “Stale NFS file handle,” after the filer is rebooted.

Caution

To ensure that the entries in the */etc/exports* file remain valid, always edit the entries in the file to reflect volume name changes immediately after renaming volumes.

Default /etc/exports entries

/vol/vol0 and /vol/vol0 home are exported by default

By default, the root volume (*/vol/vol0*) and the */vol/vol0/home* directory are exported to the administration host when you run `setup`.

Example of default exportfs file

The default */etc/exports* file contains the following lines; in this example, the name of the administration host is toaster:

```
#Auto-generated by setup Mon Oct 27 14:15:36 PST 1997
/vol/vol0 -access=toaster,root=toaster
/vol/vol0/home -root=toaster
```

The following table explains the contents of the default */etc/exports* file:

/etc/exports entry	Explanation
`/vol/vol0 -access=adminhost, -root=adminhost`	Only the administration host (adminhost) can mount the root directory and modify files in the directory. **Note** On filers with a single volume, you can refer to the root directory without the */vol* prefix.
`/vol/vol0/home -root=adminhost`	The administration host (adminhost) can mount the home directory as root. All other clients can access the home directory to read and write files.

Restricting access to volumes and directories

Use export options to restrict directory access

You can use the export options to restrict access to directories in various ways.

Restricting access to /home

You can restrict access to the */home* directory to particular groups by

- using the `chmod` command to change access modes for the directory
- using the `-rw` or `-access` options in the */etc/exports* file to limit write privilege to specific hosts

The -access option

The `-access` option lists the hosts that can mount exported directories. When you use the `-access` option, only the hosts listed can mount the associated directory.

Syntax: The syntax for the `-access` option is as follows:

`-access=`*`hostname`*`[:...:`*`hostname`*`]`

Limits: There is no limit to the number of host names you can specify with the `-access` option. However, the length of a line in the */etc/exports* file cannot exceed 1,024 characters.

Note

If you cannot fit all the host names in a 1,024-character line, you can use netgroups in place of host names.

The -root option

The `-root` option lists the hosts that can mount exported directories as root. Hosts that mount exported directories as root have full control over the directories and can perform the following operations:

- Create and delete directories and files.
- Change ownership and group associations of directories and files.
- Set access permissions for directories and files.

Syntax: The syntax for the `-root` option is as follows:

`-root=`*`hostname`*`[:...:`*`hostname`*`]`

Limits: You can specify 1 to 256 host names with the `-root` option.

Restrictions: You cannot use netgroup names with the `-root` option.

The -rw option

The `-rw` option lists the hosts that can modify the exported directories; hosts not listed by the `-rw` option have read privilege only.

Syntax: The syntax for the `-rw` option is as follows:

`-rw=`*`hostname`*`[:...:`*`hostname`*`]`

Limits: You can specify 1 to 256 host names with the `-rw` option.

Restrictions: You cannot use netgroup names with the `-rw` option.

The -ro option

The `-ro` option lists the clients that cannot modify the exported directories.

Syntax: The syntax for the `-ro` option is as follows:

`-ro=hostname[:...:hostname]`

Related information

The following table lists additional information that you might find helpful.

Read...	To learn about...
The *na_exports(1)* man page.	More about the */etc/exports* file.
"The /etc/netgroup file" on page 247	How to create and maintain netgroups.

The exportfs command

Using the exportfs command

You use the `exportfs` command to export and unexport volumes and directories. Depending on which options you use, the command exports the volumes and directories listed in the */etc/exports* file or a specific volume or directory.

Syntax

The syntax for the `exportfs` command is as follows:

```
exportfs [ -aiuv ] [ -o options ][ filer_directory_path ]
```

The following table describes the options.

Option	Description
`-a`	Exports all the entries listed in the */etc/exports* file.
`-u`	Unexports all the entries listed in the */etc/exports* file. **Note** When you use the `-u` option with the `-a` option, all exports are unexported regardless of whether they were created from the */etc/exports* file or with the `-o` option.
`-i`	Exports all the export entries listed in the */etc/exports* file, but ignores the options specified for the entries.
`-v`	Prints each path name as it is exported or unexported.

Option	Description
`-o`	Specifies the options for a volume or directory that you include in the command line. **Example:** To export the */vol/vol/home/terry* directory to the host "mixer" with read/write access, you enter the following command: **`exportfs -o -rw=mixer /vol/vol0/home/terry`** **Note** Volumes and directories that are exported directly rather than through the */etc/exports* file remain exported until canceled using the `-au` option or until the filer is rebooted.

Canceling all exports

To cancel all exports, enter the following command:

`exportfs -au`

Updating exports through /etc/exports

When you make changes to the */etc/exports* file, take one of the following actions to make the changes take effect:

- Run the `exportfs -a` command.

Or

- Reboot the filer.

If you delete entries from the */etc/exports* file when you make changes, take one of the following actions to activate the changes and ensure that deleted exports are deactivated:

- Run the following commands:

```
exportfs -au
exportfs -a
```

Or

- Reboot the filer.

The /etc/netgroup file

The /etc/netgroup file defines groups of clients

The filer */etc/netgroup* file defines groups of clients that the filer uses for checking access permission while processing a mount request.

Syntax

The following line shows the syntax for each line in the */etc/netgroup* file:

```
groupname member-list
```

Limits: Each line in the */etc/netgroup* file is limited to 4,096 characters.

Member-list syntax: Each element in `member-list` is

- another group name
- an entry in the following form:

  (*hostname, username, domainname*)

  - An element in an entry can be blank, but the commas must be present.
  - When group names are used in the */etc/exports* file, the `username` and `domainname` fields are ignored.
  - When domain names are used, they must be DNS names; NIS names cannot be used.

Restrictions

You cannot use netgroup names with the `-rw` and `-root` options.

Changes take effect immediately

Changes made to the netgroup file take effect immediately.

Example of /etc/netgroups

The following lines show an example of a */etc/netgroups* file:

```
trusted-hosts (adminhost,,)
untrusted-hosts (bashful,,) (dopey,,) (sleepy,,)
all-hosts trusted-hosts untrusted-hosts
```

Example of /etc/exports using netgroups

The following lines show an example of an */etc/exports* file that uses netgroup group names:

```
/vol/vol0 -access=trusted-hosts,root=adminhost
/vol/vol0/home -access=all-hosts,root=adminhost
```

Copy /etc/netgroup when filer doesn't use NIS

If your filer is not configured as an NIS client, the network groups on the filer are not linked with NIS.

Must copy NIS netgroup file: You must copy an existing NIS network group over to */vol/vol0/etc/netgroup* on the filer before the `exportfs` command can use it.

Automating copying with a Makefile: You can modify the Makefile of the NIS master to copy the NIS master's */etc/netgroup* file to the filer when it is changed.

Example Makefile: The following lines of code in the NIS Makefile section for `netgroup.time` copy the */etc/netgroup* file to filers named toaster1, toaster2, and toaster3; substitute the name of your filers in the "for" list, in place of *toaster1*, *toaster2*, and *toaster3*, and add any other filer names to which you want the file copied:

```
@mntdir=/tmp/nac_etc_mnt_$$$$;\
if [ ! -d $$mntdir ]; then rm -f $$mntdir; mkdir $$mntdir; fi;\
for filer in toaster1 toaster2 toaster3 ; do \
mount $$filer:/vol/vol0/etc $$mntdir;\
mv $$mntdir/netgroup $$mntdir/netgroup.bak;\
cp /etc/netgroup $$mntdir/netgroup;\
umount $$mntdir;\
done;\
rmdir $$mntdir
```

Exporting to subnets

About exporting to subnets

You can export a directory to clients on a subnet rather than to individual clients.

Valid export options for subnets

The valid export options in the */etc/exports* file for exporting to subnets are as follows:

`-ro=`*`subnet_address`*`[:`*`subnet_address`*`]...`

`-rw=`*`subnet_address`*`[:`*`subnet_address`*`]...`

`-root=`*`subnet_address`*`[:`*`subnet_address`*`]...`

Format for IP subnet addresses

The subnet address is a dotted IP subnet address and a mask written in the following format:

`dotted_ip/num_bits`

dotted_ip can be

- an IP address (“a.b.c.d”)

or

- an IP subnet
  - a for a class A network
  - a.b for a class B network
  - a.b.c for a class C network

The size of the subnet is specified by the number of leading bits of the netmask, *num_bits*.

Export to a subnet as you do to a client

You export a directory to a subnet as you do to an individual client, except that you specify a subnet address rather than a full IP address in an export option.

Example 1: root access: To export */vol/vol0/home* on the filer for root access to a client named tooting and all addresses of the form 123.45.67.x with a netmask 255.255.255.0 (24 leading bits), place the following entry in the */etc/exports* file:

```
/vol/vol0/home -root=tooting:123.45.67.0/24
```

Example 2: read/write access: To export */vol/vol0/home* for read and write access to all addresses of the form 123.45.x.y with a 16-bit netmask (255.255.0.0), place the following entry in the */etc/exports* file:

```
/vol/vol0/home -rw=123.45.0.0/16
```

Example 3: equivalent methods for exporting: The following entries in the */etc/exports* file are equivalent. They export */vol/vol0/home* to a client named host1, the specified subnet, and a client named host2.

```
/vol/vol0/home -rw=host1:123.45.67.8/24:host2
```

```
/vol/vol0/home -rw=host1:123.45.67/24:host2
```

Section B: Configuring a filer for (PC)NFS

About configuring a filer for (PC)NFS

Section contents

This section describes how to use the (PC)NFS daemon that is supplied with the NFS protocol. The daemon enables DOS, Windows, Windows NT, OS/2, and Macintosh clients that run (PC)NFS client software to access filer volumes and directories. It includes the topics listed in the following table.

Topic	Describes...
"Introducing the (PC)NFS daemon" on page 252	◆ What the daemon is. ◆ What the daemon does.
"Enabling the daemon" on page 253	◆ How to start the *pcnfsd* daemon. ◆ How to have the filer start the daemon automatically whenever the filer is rebooted.
"Setting up (PC)NFS user entries" on page 254	◆ The requirement that users have entries in the */etc/password* file and, optionally, the */etc/shadow* file on a UNIX host. ◆ How to copy the */etc/passwd* and */etc/shadow* files from a UNIX host to the filer.
"Setting the default umask" on page 256	◆ What the umask is. ◆ What the umask does. ◆ What values you can use in the umask. ◆ The syntax for the `options.pcnfsd umask` command.

Introducing the (PC)NFS daemon

What the daemon does

The filer `pcnfsd` authentication daemon enables clients running a (PC)NFS client to access a volumes and directories on the filer.

(PC)NFS refers to NFS client implementations for DOS, Windows, Windows NT, OS/2, and Macintosh operating systems.

NFS protocol must be licensed

The `pcnfsd` daemon can run on the filer only if the NFS protocol is licensed. If the NFS protocol is not licensed, the following warning messages appear:

```
pcnfsd is NOT running
Cannot use pcnfsd without nfs.
(Use the “nfs on” command to enable nfs.)
```

What (PC)NFS clients can do

Using (PC)NFS, clients can mount volumes and directories from the filer after the user provides the UNIX user name and password for authentication.

Clients must be configured to mount the filer.

Alternate access method for (PC)NFS clients

If you want (PC)NFS clients to be able to access filer volumes and directories but you do not want to run the `pcnfsd` daemon on the filer, you must set up another host to provide the (PC)NFS authentication.

(PC)NFS version 2 supported

The filer supports (PC)NFS version 2, except for printer support. With (PC)NFS version 2, clients can use the filer for authentication, including obtaining group permissions (GIDs) and group names through the */etc/group* file or the corresponding `group.byname` NIS map.

Enabling the daemon

Default daemon configuration

The `pcnfsd` daemon is disabled by default. You need to enable the daemon before (PC)NFS clients can access volumes and directories on the filer.

Procedure for enabling the daemon

Complete the steps in the following table to enable the `pcnfsd` daemon.

Step	Action
1	Enter the following command to enable the `pcnfsd` daemon: **`options pcnfsd.enable on`**
2	If you want the `pcnfsd` daemon to start automatically whenever the filer reboots, place the following command in the */etc/rc* file in the filer default volume: `options pcnfsd.enable on`

Setting up (PC)NFS user entries

Procedure for setting up user entries

Use the steps in the following table to set up user entries for (PC)NFS users.

Step	Action	
1	Verify that each user of (PC)NFS has an entry in your UNIX */etc/passwd* file and */etc/shadow* file, if applicable.	
2	**If...**	**Then...**
	The filer has not been configured to use NIS for the password database	Copy */etc/passwd* and */etc/shadow*, if applicable, to the filer. **Results:** The filer is now ready to serve files to NFS clients using (PC)NFS. Each user on the client system can mount the file system from the filer after entering a correct user name and password.
	The filer has been configured to use NIS for the password database	No action is necessary.

Example of how to copy files manually

This example shows how to copy the */etc/passwd* and */etc/shadow* files from a UNIX host to the filer. In this example, the filer root volume is mounted to the UNIX host at the following mount point:

/nfs/toaster

The following commands copy the */etc/passwd* and */etc/shadow* files from the UNIX host to the filer:

```
cp /etc/passwd /nfs/toaster/etc/passwd
cp /etc/shadow /nfs/toaster/etc/shadow
```

Copying files to the filer automatically

You can set up a UNIX host to copy the */etc/passwd* and */etc/shadow* files automatically. The following table shows you which information to read to learn about automatic copying.

If the filer...	Then read...
Uses NIS to authenticate users	"Copying files to the filer automatically from NIS Makefile"
Uses */etc/passwd* and */etc/shadow* files to authenticate users	"Copying files to the filer automatically with a cron job"

Copying files to the filer automatically from NIS Makefile

If you use NIS, you can modify the Makefile for the NIS master to copy the */etc/passwd* and */etc/shadow* files to the filer automatically.

Example: Put the following code at the end of the NIS Makefile section for `passwd.time`, substituting the name of each filer in the "for" list in place of toaster1, toaster2, and so on. Add the line copying the */etc/shadow* file only if you are using a shadow password file.

```
@mntdir=/tmp/nac_etc_mnt_$$$$;\
if [ ! -d $$mntdir ]; then rm -f $$mntdir; mkdir $$mntdir; fi;\
for filer in toaster1 toaster2 toaster3 ; do \
mount $$filer:/vol/vol0/etc $$mntdir;\
mv $$mntdir/hosts $$mntdir/hosts.bak;\
cp /etc/passwd $$mntdir/passwd;\
cp /etc/shadow $$mntdir/shadow;\
umount $$mntdir;\
done;\
rmdir $$mntdir
```

Copying files to the filer automatically with a cron job

If you do not use NIS, you can establish a `cron` job to update the filer.

Setting the default umask

What is umask?

The `umask` option for the `pcnfsd` daemon controls the permissions that the filer assigns to directories and files created by (PC)NFS clients.

How umask works

The maximum permission for a user, group, or everyone else is 6, which enables the appropriate user to read files, write files, or have no access to files. The value specified in the `umask` option is subtracted from 0666 to determine the permissions assigned to a file.

Example: The following example enables a file's owner to read and write a file, but denies access to the file by anyone else:

```
066
```

The umask format

The format of `umask` consists of three integers (*nnn*) that have values of 0, 2, 4, or 6. The integers specify permissions in the following manner:

- The first (left) integer specifies the permissions that the owner of a file has for that file.
- The second (middle) integer specifies the permissions that members of the same group as the file's owner have for the file.
- The third (right) integer specifies the permissions that everyone who is not the owner or a member of the owner's group has for the file.

Values of umask

The following table shows the `umask` integers that you can specify in any of the `umask` integer positions (*nnn*).

umask integer	Meaning
0	Creates files with permissions that enable users, groups, or everyone else to read from and write to files.
2	Creates files with permissions that enable users, groups, or everyone else to write to files.

umask integer	Meaning
4	Creates files with permissions that enable users, groups, or everyone else to read from files.
6	Creates files with permissions that prevent users, groups, or everyone else to read from or write to files.

The default umask

The filer uses a default umask of 022, which

- enables users to read and write their own files
- enables group members and others to read any users' files

Syntax

The syntax for the `umask` option is as follows:

`options pcnfsd.umask` *nnn*

nnn is the umask.

Section C: Configuring a filer for WebNFS

About configuring a filer for WebNFS

The filer can respond to NFS requests from browsers

The filer can use NFS rather than HTTP to respond to file transfer requests made through Web browsers that support the WebNFS protocol.

The filer does not need a license for the HTTP protocol to respond to WebNFS requests; however, the filer must be licensed for the NFS protocol.

Web browser requirements

To access files through the WebNFS protocol, users type URLs that start with “nfs://”. Web browsers must be capable of sending requests using the WebNFS protocol.

Advantages of WebNFS

With WebNFS, the filer can transfer files much faster than with HTTP because the WebNFS protocol can transfer several files, including graphics files, with only one TCP connection. The HTTP protocol, in version 1.1, requires a separate connection for each file that is transferred.

How WebNFS restricts file access

WebNFS access does not use the `mount` command to enable access to a subtree, and does not consider UID/GID mappings.

Requests are restricted as follows:

- Requests are honored only for files in subtrees that have been exported.
- If a subtree has been exported with the `-o access` option, files in that subtree are not available through WebNFS.

Related information

For information about the `exportfs` command, consult the *na_exportfs*(1) man page.

Setting up WebNFS

Procedure for setting up WebNFS

Use the steps in the following table to set up WebNFS.

Step	Action	
1	Enter the following command to turn on WebNFS: `options nfs.webnfs.enable on`	
2	**If you...**	**Then...**
	Want to specify a public directory, known as the root directory, for WebNFS access	Enter the following commands, replacing *directory* with the path to the root directory: `options nfs.webnfs.rootdir` ***`directory`*** `options nfs.webnfs.rootdir.set TRUE` **Results:** All NFS lookups are done relative to the root directory. All WebNFS clients can access files and directories under the root directory.
	Do not want to specify a public directory for WebNFS access	Do nothing.

Example of specifying WebNFS root directory

To use the */vol/vol0/webfiles* directory as the WebNFS root directory, you enter the following commands:

```
options nfs.webnfs.rootdir /vol/vol0/webfiles
options nfs.webnfs.rootdir.set TRUE
```

Managing WebNFS

Tasks you can perform

You can perform the following tasks to manage WebNFS service:

- Change the root directory.
- Disable the root directory.
- Turn off WebNFS service.

Changing the root directory

Complete the steps in the following table to change the WebNFS root directory.

Note

If you use the `vol rename` command to change the name of the volume in which the WebNFS root directory resides, remember to use this procedure to specify the new name of the root directory.

Step	Action
1	Enter the following command, replacing *newdir* with the path of the new root directory: **`options nfs.webnfs.rootdir newdir`** **Example:** To change the root directory to */vol/vol0/corpwebfiles*, you the following command: **`options web.webnfs.rootdir /vol/vol0/corpwebfiles`**
2	Enter the following command to enable the root directory: **`optins nfs.webnfs.rootdir.set TRUE`**

Disabling the root directory

Complete the step in the following table to disable the root directory.

Step	Action
1	Enter the following command: **`options nfs.webnfs.rootdir.set FALSE`**

Turning off WebNFS service

Complete the step in the following table to turn off WebNFS service.

Step	Action
1	Enter the following command: `options nfs.webnfs.enable off`

Section D: Displaying NFS statistics

About displaying NFS statistics

The nfsstat command displays NFS and RPC statistics

The `nfsstat` command displays statistics about NFS and Remote Procedure Calls (RPCs) for the filer. You can use the output of this command to find performance bottlenecks or inefficiencies in your NFS setup.

Note

A full description of the meaning of NFS statistics is outside the scope of this guide. A good source of information about this topic is *Managing NFS and NIS* by Hal Stern, O'Reilly & Associates, Inc.

Syntax

The syntax of the `nfsstat` command is as follows:

```
nfsstat [ interval ] | [ -h [ ip_address | host_name ] ] | [ -l ] |
[ -z ]
nfsstat -h [ ip_address | host_name ]
nfsstat -l
nfsstat -z
```

Options

The following table describes the options for the nfsstat command.

Option	Description
none	When no options are specified, the command displays statistical information since the last time the filer was rebooted.
interval	When an interval is specified, the command displays statistics continually. The interval specifies the number of seconds the command waits between updates.

Option	Description
`-h`	Displays statistics for a single client. You must provide the client's host name or IP address as an argument to the `-h` option. ◆ To use the `-h` option, you must enable the `nfs.per_client_stats.enable` option by entering the following command: **`options nfs.per_client_stats.enable on`** ◆ Enable the per-client statistics collection mode as soon as possible after you start the filer or reset the counters with `nfsstat -z`. Otherwise, `nfsstat -l` reports incorrectly low percentages and displays statistics that include clients that have generated RPC calls but no NFS calls to the client.
`-l`	Displays statistics that have been collected for all NFS clients.
`-z`	Resets the statistics counters.

Example: no options

The following lines show the output of the `nfsstat` command when you specify no options:

`nfsstat`

```
Server rpc:
TCP:
calls      badcalls   nullrecv   badlen     xdrcall
0          0          0          0          0

UDP:
calls      badcalls   nullrecv   badlen     xdrcall
24         0          0          0          0

Server nfs:
calls      badcalls
24         0

Server nfs V2: (24 calls)
null     getattr   setattr  root      lookup  readlink   read
0 0%     5 21%     0 0%     0 0%      16 67%  0 0%       0 0%
wrcache  write    create   remove    rename  link       symlink
```

```
0 0%     0 0%     0 0%     0 0%      0 0%    0 0%       0 0%
mkdir    rmdir    readdir  statfs
0 0%     0 0%     3 13%    0 0%

Server nfs V3: (0 calls)
null     getattr   setattr  root      lookup  readlink   read
0 0%     0 0%      0 0%     0 0%      0 0%    0 0%       0 0%
write    create   mkdir    symlink    mknod  remove     rmdir
0 0%     0 0%     0 0%     0 0%       0 0%   0 0%       0 0%
rename   link     readdir  readdir+  fsstat  fsinfo    pathconf
0 0%     0 0%     0 0%       0 0%      0 0%   0 0%       0 0%
commit
0 0%
```

Example: using the -l option

The following lines show the output of the `nfsstat` command when you specify the `-l` option:

```
nfsstat -l
172.17.25.13  sherlock   NFSOPS =     2943506 (90%)
172.17.25.16  watson     NFSOPS =     3553686 ( 2%)
172.17.25.18  hudson     NFSOPS =     2738083 ( 1%)
172.17.230.7  conan      NFSOPS =     673247l ( 3%)
172.17.230.8  baker      NFSOPS =   202614527 ( 1%)
172.17.230.9  moriarty   NFSOPS =     1006881 ( 0%)
175.17.230.10 doyle      NFSOPS =        1185 ( 0%)
```

Example: using the -h option

The following lines show the output of the `nfsstat` command when you specify the `-h` option with a host name:

```
nfsstat -h eng_host
Client: 172.17.25.3 (eng_host)  -------------------------------
Server rpc:
calls   badcalls nullrecv badlen  xdrcall
33374   0        0        0       0

Server nfs:
calls   badcalls
33345   0

Server nfs V2:
null    getattr  setattr  root     lookup     readlink  read
0 0%    8410 25% 19 0%    0 0%     13687 41%  42 0%     489 22%
wrcache write    create   remove   rename     link      symlink
0 0%    3225 10% 12 0%    7 0%     9 0%       0 0%      0 0%
mkdir   rmdir    readdir  statfs
```

```
0 0%    0 0%     416 1%   29 0%

Server nfs V3:
null    getattr  setattr  lookup   access     readlink  read
0 0%    0 0%     0 0%     0 0%     0 0%       0 0%      0 0%
write   create   mkdir    symlink  mknod      remove    rmdir
0 0%    0 0%     0 0%     0 0%     0 0%       0 0%      0 0%
rename  link     readdir  readdir+ fsstat     fsinfo    pathconf
0 0%    0 0%     0 0%     0 0%     0 0%       0 0%      0 0%
commit
0 0%
```

Example: resetting counters with the -z option

The following command resets the counters:

```
nfsstat -z
```

Related information

The following table lists additional information that you might find helpful.

To learn about...	Read...
How the Network Appliance custom Management Information Base (MIB) provides the same information as the `nfsstat` command through SNMP	"About the Network Appliance custom MIB" on page 112.
About the `nfsstat` command	The *na_nfsstat(1)* man page.

CIFS Administration 8

What is CIFS?

CIFS (Common Internet File System) is a file-sharing protocol based on the Server Message Block (SMB) protocol widely in use by personal computers and workstations running a wide variety of operating systems. CIFS provides an open, cross-platform mechanism for client systems, including Windows systems, to request file and print services from server systems over a network.

This chapter describes how to manage the CIFS file protocol and users.

What you can do only from the filer command line or FilerView

You can do some operations that affect CIFS administration only from the filer command line or FilerView. These are

- Viewing volumes and examining their status. For additional information about volume operations, see Chapter 3, "Disk and File System Management," on page 71.
- Setting or changing volume and qtree security style and oplocks status. For information about administering qtrees from the filer command line, see Chapter 12, "Qtree administration." For information about how to administer qtrees in FilerView, see FilerView on-line help.

Chapter contents

This chapter discusses the topics listed in the following table.

Read...	To learn about...
"CIFS limitations" on page 270	Describes CIFS limitations when using the ACL editor, User Manager, and Server Manager.
"Limits on CIFS open files, sessions, and shares" on page 271	CIFS limits on a filer.
"Limits for F700 and series filers" on page 272	How to change or examine descriptive information about a filer.
"Adding CIFS users to the filer" on page 275	How to add users to a filer from Windows NT and the filer command line.
"Adding local groups to the filer" on page 278	How to use User Manager to add local groups to the filer.

Read...	To learn about...
"Using CIFS commands with a remote shell program" on page 281	How to use a remote shell to automate CIFS management.
"Enabling guest and generic access" on page 282	How to give access to users who do not have an account on the filer.
"Displaying a filer's shares" on page 284	How to display information about the shared directories on a filer.
"Creating and changing a share" on page 287	How to create a share from Windows NT and the filer and change the value of its parameters.
"Displaying information about shares" on page 294	How to get information about shares from Windows NT and the filer.
"Deleting a share" on page 297	How to delete shares from Windows NT and the filer.
"Creating a home share for each user" on page 299	How to create home shares from Windows NT and the filer.
"Assigning and changing access rights" on page 302	How to define user or group access rights to a share.
"Displaying access rights to an NTFS file" on page 307	How to show who has rights to an NTFS file on the filer.
"Changing UNIX permissions and DOS attributes from Windows" on page 308	How to use SecureShare™ Access to view and change UNIX permissions and DOS attributes.
"Sending a message to all users on a filer" on page 311	How to send a message in an alert box to all users on a filer.
"Event auditing" on page 313	How to monitor access to a file by a specified user.
"Using oplocks" on page 328	How to set and remove opportunistic locking from the filer as a whole.
"Displaying CIFS statistics" on page 330	How to display CIFS information.

Read...	To learn about...
"Displaying CIFS session information" on page 331	How to display information from the filer command line about who is logged in to CIFS.
"Stopping and restarting CIFS sessions" on page 333	How to stop and start CIFS service from the filer command line.
"Reconfiguring the filer for CIFS" on page 337	How to run the `cifs setup` program to reconfigure CIFS on the filer

Effects of renaming a volume on shares

If you change the name of a volume that contains at least one share, whether through Windows NT or the `vol rename` command, the filer automatically offers the share to users in the renamed volume at the next reboot. The new volume name is reflected in the `cifs shares` command.

Scope of this chapter

This chapter does not discuss procedures that take place on the clients and it does not describe how a machine joins a workgroup or domain. For information about these topics, refer to the manuals for your client operating systems or books about PC networking.

CIFS limitations

Introduction

This section describes CIFS limitations when operating on files on the filer.

User Manager limitations

The Policy menu items and the New Users menu item are permanently disabled.

Server Manager limitations

The following Server Manager features are not supported:

- stopping and starting services
- specifying the recipients of alerts

Limits on CIFS open files, sessions, and shares

CIFS file access limits

The filer is subject to limits on file access through CIFS, as summarized in the following tables.

Limits for F200 and F300 series filers

The following table shows access limits for F200 and F300 series filers.

CIFS limits by filer memory	128 MB	192 MB	256 MB
Maximum number of users	1,600	3,000	4,000
Maximum number of open files	32,000	60,000	80,000
Maximum number of locked files	36,096	66,144	88,192
Maximum number of shares	3,200	6,000	8,000

Limits for F500 and F600 series filers

The following table shows access limits for F500 and F600 series filers.

CIFS limits by filer memory	128 MB	192 MB	256 MB	512 MB
Maximum number of users	1,200	2,200	3,200	7,200
Maximum number of open files	24,000	44,000	64,000	144,000
Maximum number of locked files	32,000	50,144	72,192	160,384
Maximum number of shares	2,400	4,400	6,400	14,400

Limits for F700 and series filers

The following table shows access limits for F700 series filers with at least 120 MB of memory.

Type of access	Maximum number
Users	14,825
Files	269,500
Locked files	88,960
Shares	29,650

Changing or viewing the filer's description

When to change or view a filer's description

The description of a filer appears next to its name wherever a machine's description comments appear. Initially, the filer has no description. You might want to change the description to something more informative so that you can distinguish filers from each other. You might want to view the description of a filer to find out, for example, what a particular filer does or who is in charge of it.

Changing a filer's description from Server Manager

To change the description of a filer from Server Manager, follow these steps.

Step	Action
1	Open Server Manager by choosing it from the Start menu.
2	From the Server Manager Computer menu, choose Select Domain. **Result:** The Select Domain window opens.
3	In the Select Domain window, select the filer you want by typing its UNC name, for example, \\FILERNAME, in the Domain field, then clicking OK.
4	Double-click a filer name. **Result:** The Filer Properties window appears.
5	Type a new description of the filer in the Description field.

Step	Action
6	Click OK. **Result:** The new description goes into effect.

Viewing a filer's description from the filer command line

To view the description of a filer, enter the following command:

```
cifs comment
```

Changing a filer's description from the filer command line

To change the description of a filer, enter the following command:

```
cifs comment "description"
```

where *description* is its description, which must be enclosed in quotation marks. The description must be no longer than 48 characters.

Adding CIFS users to the filer

When you add CIFS users

By default, Windows NT users who have an account map to a UNIX account of the same name. If they do not have a Windows NT account, the user becomes a generic user and uses the generic account, which is described in "Generic user accounts" on page 282. In the rare case that you must add a user explicitly, you can do so with the methods described in this section.

You can add CIFS users to the filer at any time. The method you use for adding users to the filer depends on whether you are authenticating with a domain controller or the UNIX password database.

When authenticating with a domain controller

To add a user, create an account for the CIFS user within your Windows NT domain environment. If you want the user to also use UNIX files, either create an entry in the */etc/passw*d file on the filer or include the user in the */etc/usermap.cfg* file.

What is the /etc/usermap.cfg file? The */etc/usermap.cfg* file explicitly maps Windows NT users to the correct UNIX account and UNIX users to a Windows NT account. The file can be coded as follows:

- As a text file, with non-ASCII characters encoded in the NFS character set.
- As a UNICODE file created using the Windows NT tools Notepad or Microsoft Word.

The filer automatically detects which of these forms is in use.

Format of the /etc/usermap.cfg file: The format of the */etc/usermap.cfg* file is a list of text records in the following format:

```
[IP-qual:] [NT-domain\]NTUser [direction] [IP-qual:] UnixUser
```

Lines are processed sequentially.

Format variables: The following table describes the variables in the */etc/usermap.cfg* file description.

Variable	Description
IP-qual	An IP qualifier that the filer uses in matching a user. You use an IP qualifier to narrow a match. *IP-qual* can be an IP address in any of the following formats: ◆ A regular IP address specified as numbers separated by periods (dot notation) with an optional subnet address. For example, 192.4.1.0/24 narrows possible matches to the 192.4.1.0 class C subnet. ◆ A host name. ◆ A network name. ◆ A network name with a subnet specified in dot notation. For example, corpnet/255.255.255.0 specifies the corpnet subnet.
NT-domain	Specifies the domain to match or the domain to use for a mapped UNIX account. The default is the domain in which the filer is installed.
NTUser	Any user-type account name. If the name contains a space, put the name in quotation marks.
direction	Restricts the direction of the mapping. By default, mappings are bidirectional. You can use one of three symbols: ◆ => means NT to UNIX mapping only. ◆ <= means UNIX to NT mapping only. ◆ == means bidirectional mapping. Use this to explicitly indicate a bidirectional mapping.
UnixUser	A UNIX account name.

Symbol conventions: The following symbol conventions are in effect:

- An asterisk (*) matches any name.

  For example, to map all unmapped users to the UNIX “nobody” account, add the following line:

  ```
  *        nobody
  ```

- The null string (“”) matches no name and rejects any user.

  For example, to prevent access completely, add the following line:

  ```
  *        ""
  ```

This line prevents the default mapping of Windows NT users who have an account map to a UNIX account of the same name. Any lines after this line are disregarded by the filer.

- You can use either spaces or tabs as separators.

Name requirements: Windows NT and UNIX names have different requirements, as follows:

- Windows NT names are case-insensitive and can contain non-ASCII characters within the character set in the current code page. Windows NT user names can contain spaces, in which case you must enclose the name in quotation marks.
- UNIX user names are case-sensitive and must be in ASCII.

Default file contents: If the filer is domain authenticated, by default the */etc/usermap.cfg* file contains the following line:

```
domain\administrator    root
```

When authenticating with the UNIX password database

To add a user, enter the user's information into the NIS password and group maps.

Note

If you do not use NIS, create entries for the user in the filer's */etc/passwd* and */etc/group* files.

Adding local groups to the filer

How to add a local group

You add a local group to the filer with the New Local Group window in the User Manager for Domains.

Adding a group with the New Local Group window

Use the New Local Group window to create a new local group as follows

Step	Action
1	Open User Manager for Domains by choosing it from the Start menu.
2	From the User menu, choose Select Domain.
3	Select the filer you want by typing its UNC name, for example, \\FILERNAME, in the Domain field, then clicking OK. **Result:** The User Manager window shows information for the specified filer, as shown below. 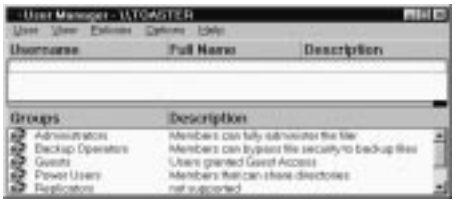
4	From the New User menu, choose New Local Group. **Result:** The New Local Group window appears.

Step	Action
5	Type the name of the new group in the Group Name text box.
6	Type a description for the new group in the Description text box.
7	To add a member, type a user or group name in the Members list box or use the Add Users and Groups window, as described in Step 8.
Adding a name with the Add Users and Groups window	
8	Click Add. **Result:** The Add Users and Groups window appears.
9	Click the arrow next to the List Names From text box to choose a domain that contains names that you want to add. **Result:** A list of names in the selected domain appears in the Names list box.
10	To add a name, type one or more user names in the Add Names list box or select one or more names and click Add.
11	Specify one or more name in either of these ways: ◆ Click one or more names in the Names list box. ◆ Type valid user or group names in the Add Names list box.

Step	Action
12	Click Add in the window where you specified the names. **Result:** The names are added to the Add Names field of the Add Users and Groups window.
13	If you want, display the full name of a user associated with an account name by clicking Show Full Name.
14	Click OK. **Result:** The names appear in the Members field of the New Local Group window.
Final steps in the New Local Group window	
15	To remove one or more names from the list, select a name or names in the Names list, then click Remove.
16	Click OK to put the additions into effect.

Using CIFS commands with a remote shell program

What you can use a remote shell program for

You can use a remote shell program, such as `rsh`, to

- execute CIFS commands
- create scripts containing CIFS commands to automate similar access rights tasks

UNIX example

For example, using `rsh` on the administration host, you can set access rights for a specific user as follows:

```
rsh -l root -n toaster cifs access home jsmith r-x
```

In this example, the filer name is toaster, the user is jsmith, and the share is home. The user has read, execute, and browsing rights to the directory on the filer that has been defined as the home share.

Automating access rights

Because you can use CIFS commands through a remote shell program, you can automate the task of defining access rights for multiple CIFS filers with similar user information. For example, you can create a script containing CIFS commands to enter similar user information for each filer at your site.

Required information in hosts.equiv file

Make sure that the following entries are added to the *hosts.equiv* file in the *etc* directory:

host *user*
host root

where *host* is the host you are using and *user* is your name.

Note

Make sure that you include both lines, or not all remote shell functionality works properly.

Enabling guest and generic access

Two ways to give access to unauthenticated or occasional users

You can enable unauthenticated or occasional users to have access to the filer through CIFS in one of two ways.

For users who are...	Use...
not in a trusted domain	a guest account
authenticated do not have an entry in the */etc/passwd* file on the filer	a generic user account

Guest accounts

If you using Windows NT domain authentication, guests are users who are not in trusted domains.

If you use a UNIX password database for authentication, guests are users who do not have an entry in the database.

Setting up a guest account: To set up a guest account, use the following `options` command:

```
options cifs.guest_account account_name
```

where *account_name* is the name of the guest account, usually *guest* or *nobody*, which is a preconfigured account in the */etc/passwd* file.

If you are using UNIX-style authentication, set the guest account to the name of an account in the UNIX password database, typically *guest*, which is mapped to the UNIX account *nobody* with the same access rights as the user *everyone*.

Disabling guest access: To disable guest access, use the following `option` command:

```
options cifs.guest_account
```

with a “” as *account_name*.

Generic user accounts

- If you are using Windows NT domain authentication, a generic user account is mapped by default with the name “pcuser.” The generic user account enables users who meet the criteria described later in this section to connect

to NTFS or mixed qtrees on the filer. For information about qtrees, see Chapter 12, "Qtree administration," or FilerView on-line help.

Note

For generic user accounts to be active, pcuser must be an account in the */etc/passwd* file.

Who can use the generic user account: To use the generic user account, a user must

- be authenticated
- be in a trusted domain
- not have an entry in the UNIX password database

All users of the generic user account have the same UNIX rights and the Windows NT rights granted by their Windows NT group membership because they appear as one account to the system.

Setting up a generic user account: To set up a generic user account, use one of the following `options` commands with an account name to use as an argument.:

User type	Option	Description
NT	`wafl.default_unix_user`	Specifies the UNIX user account to use when an NT user attempts to log in and that NT user would not otherwise be mapped.
UNIX	`wafl.default_nt_user`	Specifies the NT user account to use when a UNIX user accesses a file with NT security (has an ACL), and that UNIX user would not otherwise be mapped.

Disabling generic user access: To disable generic user access, use the `options` described in the previous table with a `""` as *account_name*.

Displaying a filer's shares

Ways to share folders

If you want to create a folder to be shared by CIFS clients, do one of the following:

- From Windows NT, use Server Manager as described in the next section, "Using Server Manager to display a filer's shares."
- From the filer command line, use the `cifs shares` command to display share information, as described in "Using the cifs shares command to display a filer's shares" on page 285.

Using Server Manager to display a filer's shares

To display a filers's shares with Server Manager and get detailed information, follow these steps.

Step	Action
1	Open Server Manager by choosing it from the Start menu.
2	From the Server Manager Computer menu, choose Select Domain. **Result:** The Select Domain window opens.
3	In the Select Domain window, select the filer you want by typing its UNC name, for example, \\FILERNAME, in the Domain field, then clicking OK. **Result:** A Server Manager window for the filer appears.

Step	Action
4	Choose Computer > Shared Directories. **Result:** The Shared Directories window appears. 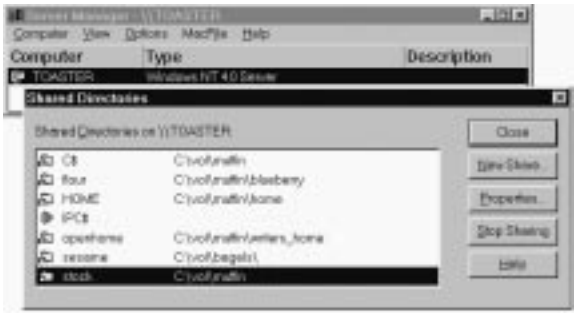
5	To get additional details about a share, double-click a share name. **Result:** The Share Properties window appears. 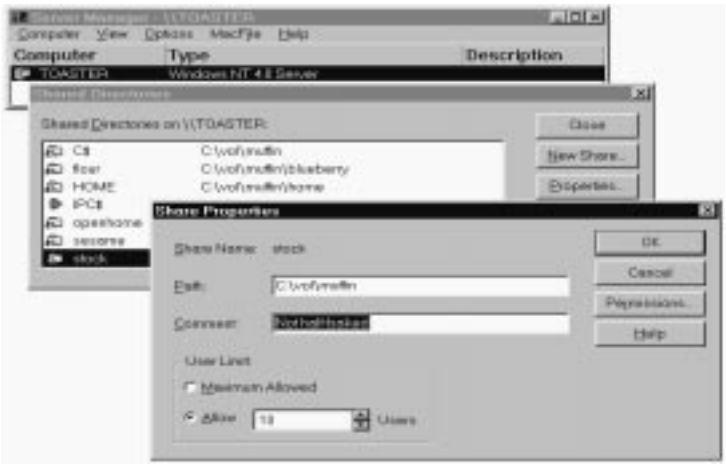

Using the cifs shares command to display a filer's shares

To display the filer's list of shares from the filer's command line, use the `cifs shares` command.

Command syntax: The syntax is as follows:

```
cifs shares
```

Example of displaying a filer's shares: The following example shows a filer's shares:

```
cifs shares
Name         Mount Point                      Description
----         -----------                      -----------
HOME         /vol/muffin/home                 Default Share
                       everyone / Full Control
                       techpubs{g} / Full Control
C$           /vol/muffin                     Remote Administration
                       everyone / Full Control
                       BUILTIN\Administrators / Full Control
openhome     /vol/muffin/writers_home         readable home dirs
             ... user limit=1200
             ... forcegroup=techpubs
                       everyone / --x
                       techpubs{g} / r-x
stock        /vol/muffin                      Not half-baked
             ... user limit=10
                       everyone / Full Control
flour        /vol/muffin/blueberry            Flour power
                       everyone / Full Control
sesame       /vol/bagels/
                       everyone / Full Control
```

Creating and changing a share

Ways to share folders

If you want to create a folder to be shared by CIFS clients, do one of the following:

- From Windows NT, use Server Manager as described in the next section, "Creating a share from Server Manager."
- From the filer command line, use the `cifs shares` command to define a new share, as described in "Creating a share with the cifs shares command" on page 290.

Note

By default, three shares are created during CIFS setup: *C$*, *IPC$*, and *HOME*. In filer console displays, the *C$* share corresponds in UNIX to */vol/vol0* and the HOME share corresponds to */vol/vol0/home*.

Creating a share from Server Manager

To create a share from the Windows NT desktop using Server Manager, create a folder, then share it by following these steps.

Step	Action
1	Create a folder on the filer.
2	Open Server Manager by choosing it from the Start menu.
3	From the Server Manager Computer menu, choose Select Domain. **Result:** The Select Domain window opens.
4	In the Select Domain window, select the filer you want by typing its UNC name, for example, \\FILERNAME, in the Domain field, then clicking OK. **Result:** A Server Manager window for the filer appears.

Step	Action
5	Chose Computer > Shared Directories, then click the New Share button. **Result:** The New Share window appears. 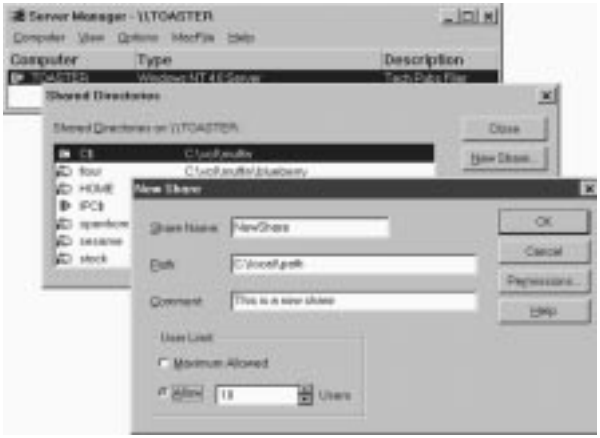
6	Type a name for the new share in the Share Name field. The share name is case-sensitive.
7	Type the local path here of the folder you created in Step 1. This is usually C:*VOL\VOLNAME\FOLDER*.
8	Type a description of the share in the Comments field, if you want.
9	If you do not want to limit the number of users that can connect to the share at the same time, select Maximum Allowed under User Limit.
10	To limit the number of users that can connect to the share at the same time, click the arrows next to Users until the desired number appears.
11	If the share has Windows NT security, click Permissions to set permissions.
12	Click OK. **Result:** The New Share window disappears and the share is created.

Changing the share description and user limit with Server Manager

To change the description and user limit with Server Manager, follow these steps.

Step	Action
1	Open Server Manager by choosing it from the Start menu.
2	From the Server Manager Computer menu, choose Select Domain. **Result:** The Select Domain window opens.
3	In the Select Domain window, select the filer you want by typing its UNC name, for example, \\FILERNAME, in the Domain field, then clicking OK. **Result:** A Server Manager window for the filer appears.
4	Choose Computer > Shared Directories. **Result:** The Shared Directories window appears. 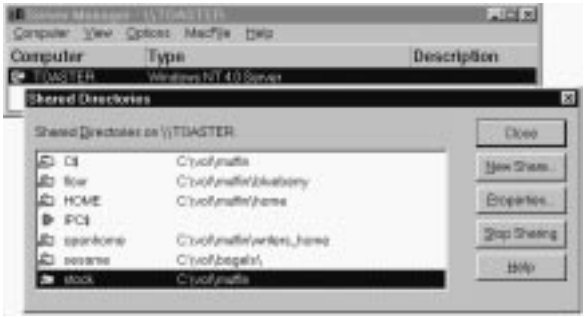

Step	Action
5	Double-click a share name. **Result:** The Share Properties window appears. 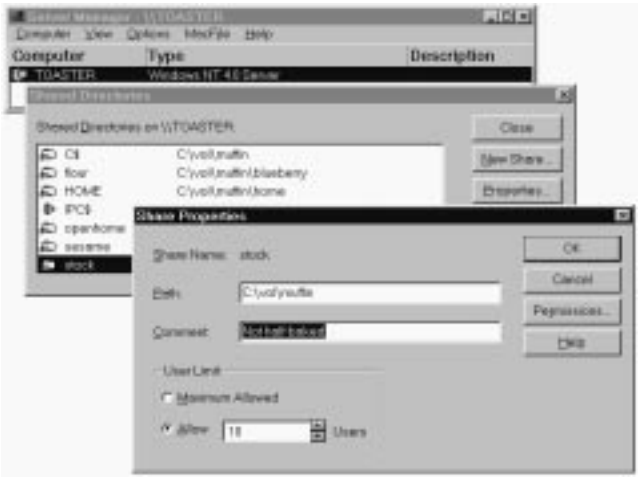
6	Type a new description of the share in the Comments field, if you want.
7	If you do not want to limit the number of users that can connect to the share at the same time, select Maximum Allowed under User Limit.
8	To limit the number of users that can connect to the share at the same time, click the arrows next to Users until the desired number appears.
9	Click Permissions to change share-level permissions.
10	Click OK. **Result:** The window disappears and the new values go into effect.

Creating a share with the cifs shares command

Following is the syntax for the `cifs shares` command for creating a share:

```
cifs shares -add sharename path [ -comment description ]
[ -forcegroup groupname ] [ -maxusers n ]
```

The following table describes the parameters.

Variable	Description
description	A string describing the purpose of the share. It must contain only characters in the current code page. It is required by the CIFS protocol and is displayed in the share list in Network Neighborhood on the client. If the description contains spaces, you must enclose it in single quotation marks.
groupname	The name of the group you want all files in the share to get the group membership of.
n	The maximum number of users that you specify can connect to the share at the same time. The limit on this number is dependent on filer memory, and is shown in "Limits on CIFS open files, sessions, and shares" on page 271.
path	The path name, relative to the root of the filer's file system, of the root directory of the share. **Note** Because the `cifs shares` command is case-sensitive, be sure that you use the appropriate case when entering the path name. Separators used in the path name must be forward slashes (/).
sharename	The name of the share, which is used by CIFS users to obtain access to the directory on the filer. If *sharename* already exists, the `cifs shares -add` command fails. **Caution** Do not create shares whose names end with a dollar sign ($); doing so might cause conflicts with reserved names. In particular, do not create shares called C$ or ADMIN$. C$ is a reserved share name and ADMIN$ is an illegal share name.

Example: The following example creates a new share called *library*:

```
cifs shares -add library /vol/vol0/home/lib 'New file library'
```

This example creates the *library* share and defines it as *HOME\LIB* (*/home/lib* in UNIX notation) in the filer's root volume. With the appropriate access rights, CIFS users can gain access to the *HOME\LIB* directory in the root volume, which is displayed as the *library* share on their computers. For more information about setting access rights, refer to "Assigning and changing access rights" on page 302.

Using the cifs shares command to change the share

To change the description, forced file ownership, and user limits of a share, use the following command:

```
cifs shares -change sharename [ -comment description | -nocomment ]
[ -forcegroup groupname | -noforcegroup ]
[ -maxusers n | -nomaxusers ]
```

The following table describes the parameters.

Parameter or variable	Description
description	A string describing the purpose of the share. It must contain only characters in the current code page. It is displayed in the share list in Network Neighborhood on the client. If the description contains spaces, you must enclose it in single quotation marks.
groupname	The name of the group you want all files in the share to get the group membership of.
n	The maximum number of users that you specify can connect to the share at the same time. The limit on this number is dependent on filer memory, and is shown in "Limits on CIFS open files, sessions, and shares" on page 271.
-nocomment	Specifies no description.
-noforcegroup	Specifies no particular group to own the files that are created in the share.
-nomaxusers	Specifies no maximum number of users that can connect to the share at the same time.

Parameter or variable	Description
sharename	The name of the share, which is used by CIFS users to obtain access to the directory on the filer. It must contain only characters in the current code page.

Displaying information about shares

Methods of displaying information about shares

You can display information about a specific share or all shares on the filer.

To display information about shares, do either of the following:

- From the Windows NT desktop, use Server Manager as described in the next section, "Using Server Manager to view information about shares."
- From the filer, use the `cifs shares` command as described in "Using the cifs shares command to view information about shares" on page 295.

Using Server Manager to view information about shares

To display information about shares from Server Manager, follow these steps.

Step	Action
1	Open Server Manager by choosing it from the Start menu.
2	From the Server Manager Computer menu, choose Select Domain. **Result:** The Select Domain window opens.
3	In the Select Domain window, select the filer you want by typing its UNC name, for example, \\FILERNAME, in the Domain field, then clicking OK. **Result:** A Server Manager window for the filer appears.
4	Double-click the name of a filer. **Result:** The Filer Properties window appears.

Step	Action
5	Click the Shares button. **Result:** Information about each share on the filer appears, as shown below. 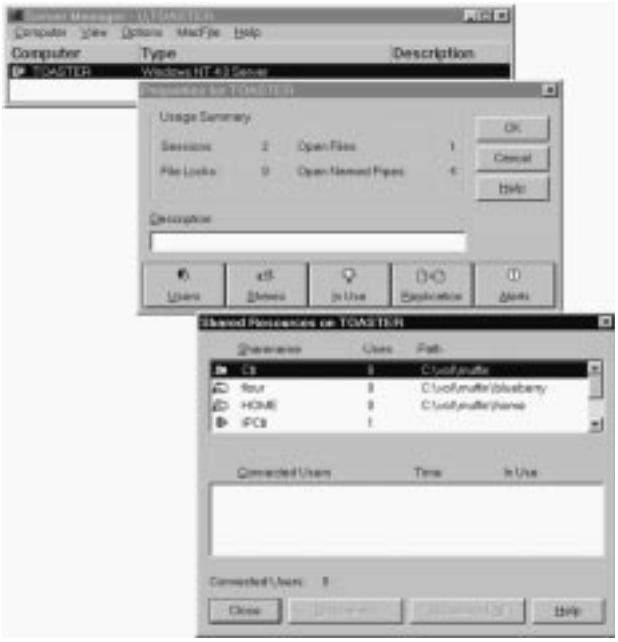

Using the cifs shares command to view information about shares

To display information about a share or shares that exist, enter the following command:

```
cifs shares sharename
```

where *sharename* is the specific name of the share about which you want information. If you want information about all shares, leave *sharename* blank.

Examples of displaying share information: The following command example displays the share information only for the share *library* created by the `cifs shares -add` command described in "Creating a share with the cifs shares command" on page 290.

```
cifs shares library
```

The following command displays information about all shares:

```
cifs shares
```

The sample output is from `cifs shares`:

```
Name         Mount Point                      Description
----         -----------                      -----------
HOME         /vol/muffin/home                 Default Share
                       everyone / Full Control
                       techpubs{g} / Full Control
C$           /vol/muffin                     Remote Administration
                       everyone / Full Control
                       BUILTIN\Administrators / Full Control
openhome     /vol/muffin/writers_home         readable home dirs
             ... user limit=1200
             ... forcegroup=techpubs
                       everyone / --x
                       techpubs{g} / r-x
stock        /vol/muffin                      Not half-baked
             ... user limit=10
                       everyone / Full Control
flour        /vol/muffin/blueberry            Flour power
                       everyone / Full Control
sesame       /vol/bagels/
                       everyone / Full Control
```

Deleting a share

How to delete a share

To delete shares, do either of the following:

- From the Windows NT desktop, use Server Manager as described in the next section, "Using Server Manager to delete a share."
- From the filer, use the `cifs shares` command, as described in "Using the cifs shares command to delete shares" on page 298.

Using Server Manager to delete a share

To delete a share from Server Manager, follow these steps.

Step	Action
1	Open Server Manager by choosing it from the Start menu.
2	From the Server Manager Computer menu, choose Select Domain. **Result:** The Select Domain window opens.
3	In the Select Domain window, select the filer you want by typing its UNC name, for example, \\FILERNAME, in the Domain field, then clicking OK. **Result:** A Server Manager window for the filer appears.

Step	Action
4	In the Server Manager window for the filer, choose Computer > Shared Directories, then click the Properties button. **Result:** The Share Properties window appears, as shown below. 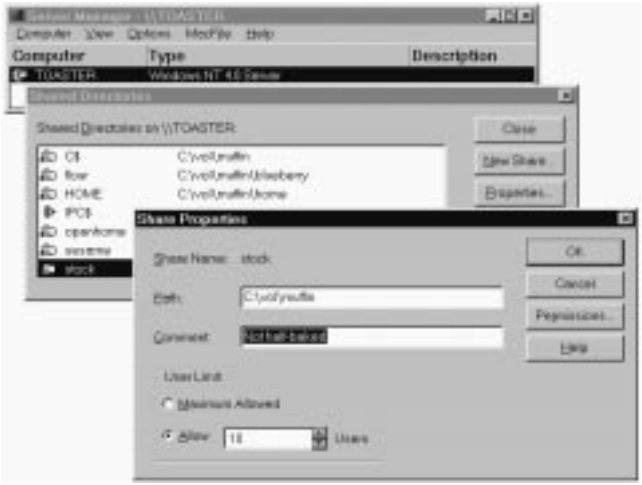
5	Click Stop Sharing, then click OK. **Result:** The folder is no longer shared.

Using the cifs shares command to delete shares

To delete a share, use the `cifs shares -delete` command.

Command syntax: The `cifs shares -delete` command has the following syntax:

```
cifs shares -delete sharename
```

where *sharename* is the specific name of the share that you want to delete.

Example: The following command deletes the share created by the `cifs shares -add` command described in "Creating a share with the cifs shares command" on page 290.

```
cifs shares -delete library
```

Creating a home share for each user

When to create a home directory

You can create a share that contains home directories of registered CIFS users.

For example, if there are users called *user1* and *user2*, the share contains directories for user1 and user2. When user1 connects to the filer and asks for its list of shares, the display shows a share called user1, but not user2 or any other individual user.

Note

If Domain1\user and Domain2\user are the same, they do not have different home directories. To prevent access by the wrong user, set Windows permissions and UNIX permissions at the root of the user's home directory.

Accessing a home directory

Users access their home directories in the same way as any other share. That is, users can open the share with Network Neighborhood, by mapping a drive, or by using a UNC name. The UNC name is

`\\filer\`*`username`*

Share name length limitations

Because share names are truncated to 12 characters, the home directory name might show a truncated version of the user's account name.

For example, consider the 13-letter name *administrator*. From the filer point of view, the home directory for administrator has a name that exactly matches the account name *administrator*, but use of that directory is offered under the truncated share name *administrato*.

If there were users *administrator* and *administrator1*, they would see an offer for their own home directory, but both would see it as an offer to share the name *administrato*. Therefore, each user gets the correct directory, even though the share names appear the same.

Creating a share containing user home directories

The following procedure creates a share that automatically contains home directories for CIFS users:

1. Enable the home directory option by entering the following command.

   `options cifs.home_dir homedirpath`

   Where *homedirpath* is the UNIX path name that will be mapped to the share CIFS.HOMEDIR. This share automatically contains CIFS user home directories.

   If the directory containing the home directories is */vol/vol1/homes*, enter the following command:

   `options cifs.home_dir /vol/vol1/homes`

2. Within the directory specified by the `cifs.home_dir` option, create a directory for each user. Make sure that the following conditions are met:
   - Each directory name matches the user's login name exactly.
   - Each user is the owner of the directory.

Creating share home directories

Using the `cifs.home_dir` option to create home directories is useful only if users do not need to read or write other users' home directories. If they must access other users' home directories, follow this procedure:

1. Create an additional share, from either Windows or the filer, that maps to the same path name as the CIFS.HOMEDIR share.
2. From either Windows or the filer, assign each user the appropriate access permissions to other users' home directories.

Example from the filer: The following example shows how to create an additional share, assign user access rights to the share, and display the share information from the filer command line. If authentication is through the */etc/passwd* file, UNIX permissions are shown; otherwise, Windows NT permissions appear.

```
cifs shares -add enghomes /vol/vol1/homes \
   -comment "Readable home directories"
```

```
“Readable home directories”
```

```
cifs access enghomes -g engineering r-x
```

```
cifs shares
Name         Mount Point                       Description
----         -----------                       -----------
enghome      /vol/vol1/homes                   Readable home
directories
                        engineering{g} / r-x
C$           /vol/vol0                        Remote Administration
                        BUILTIN\Administrators / Full Control
HOME         /vol/vol0/home                    Default Share
                        everyone / Full Control
CIFS.HOMEDIR /vol/vol1/homes                   Home Directories
                        everyone / Full Control
```

Result: Users in the *engineering* group can read all home directories in HOMES on the root volume, which corresponds to the share defined by the `options cifs.home_dir` command. However, they can only write to their own home directories, which reside in the CIFS.HOMEDIR share.

Assigning and changing access rights

When to assign or change access rights

After you create a share, you define the user or group access rights to the share. If a group or a user no longer exists, you can remove the corresponding entry from an ACL.

Methods of assigning or changing access rights to a share

To assign access rights, use either of the following methods:

- From Windows, follow the instructions in the next section, "Assigning or changing access rights with Server Manager."
- From the filer command line, do one of the following:
  - Add access with the `cifs access` command, as described in "Giving access with the cifs access command" on page 305.
  - Remove a user or group with the `cifs access -delete` command, as described in "Removing a user or group with the cifs access -delete command" on page 306.

Assigning or changing access rights with Server Manager

To assign or change access rights with Server Manager, use the Access Through Share Permissions window by following these steps.

Step	Action
1	Open Server Manager by choosing it from the Start menu.
2	From the Server Manager Computer menu, choose Select Domain. **Result:** The Select Domain window opens.
3	In the Select Domain window, select the filer you want by typing its UNC name, for example, \\FILERNAME, in the Domain field, then clicking OK. **Result:** A Server Manager window for the filer appears.

Step	Action
4	In the Server Manager window, click Computer > Shared Directories, then click the Properties button. **Result:** The Share Properties window appears. 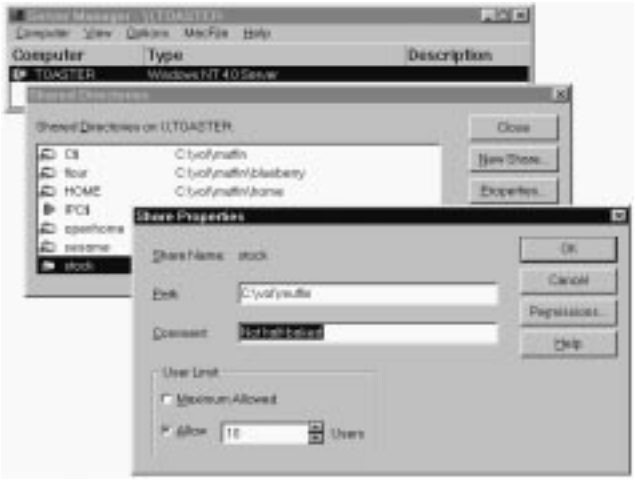
5	Click Permissions. **Result:** The Access Through Share Permissions window appears.

Step	Action
6	To add a new user or group, click Add. **Result:** The Add Users and Groups window appears. 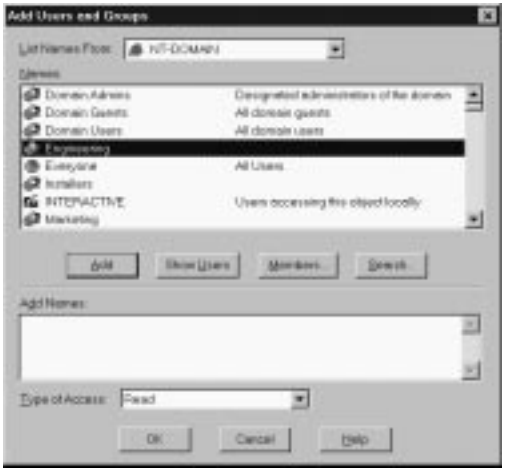
Adding a name with the Add Users and Groups window	
7	Click the arrow next to the List Names From field to choose a domain that contains names that you want to add. **Result:** A list of names appears in the Names list box.
8	Specify one or more names in either of the following two ways: ◆ Click one or more names in the Names list box. ◆ Type valid user or group names in the Add Names text box.
9	Click Add in the window where you selected the names. **Result:** The names are added to the Add Names field of the Add Users and Groups window.
10	If you want, add or modify the access type by selecting one or more names and choosing an access type from the Type of Access list.
Steps in the Access Through Share Permissions window	

Step	Action
11	To assign or change an access type, select a name or names in the Names list, then click the arrow next to Type of Access and select an access type.
12	To remove one or more names from the list, select a name or names in the Names list, then click Remove.
13	Click OK.
Final step in the Properties window	
14	Click OK to put the changes into effect.

Giving access with the cifs access command

You use the `cifs access` command to assign access to a share from the filer command line. To change the permissions, run the command with the new permissions.

Command syntax: The `cifs access` command has the following syntax:

`cifs access` *`share`* `[ -g ]` *`user|group rights`*

The `-g` flag specifies that the access rights are defined for a group.

rights can be UNIX-style permissions or Windows NT-style rights:

- UNIX-style permissions are defined as `r`, `w`, and `x`, which mean read, write and delete, and execute and browse, respectively. To deny a right to a user, use a hyphen (-).
- Windows NT-style rights are No Access, Read, Change, and Full Control.

Examples: Here are some examples of assigning access rights from the filer:

- `cifs access library -g engineering rwx`
- `cifs access library domain\joed Change`

Note

The group *everyone* is reserved. When you use it in an ACL, the group *everyone* means every CIFS user. For example, to give every CIFS user read, write, and execute rights to the *library* share, you enter the following command:
`cifs access library everyone rwx`

Removing a user or group with the cifs access -delete command

If a user or group no longer exists, you can remove the corresponding entry from an ACL. Use the following command syntax to remove an entry in an ACL:

```
cifs access -delete share user | group
```

Following are some examples of removing entries from the *library* share.

The following command removes the *engineering* group from the *library* share.

```
cifs access -delete library engineering
```

The following command removes the user *joed* from the *library* share.

```
cifs access -delete library domain\joed
```

Displaying access rights to an NTFS file

Access rights display methods

You can display access rights to an NTFS file from Windows.

Displaying access rights from the Windows desktop

To display access rights to a file from the Windows desktop, follow these steps.

Step	Action
1	Right-click a file and choose Properties from the pop-up menu. **Result:** The Properties sheet appears.
2	Click the Security tab. **Result:** If the file is a Windows file, the Security sheet appears.
3	Click Permissions. **Result:** Permissions are displayed, as shown below. 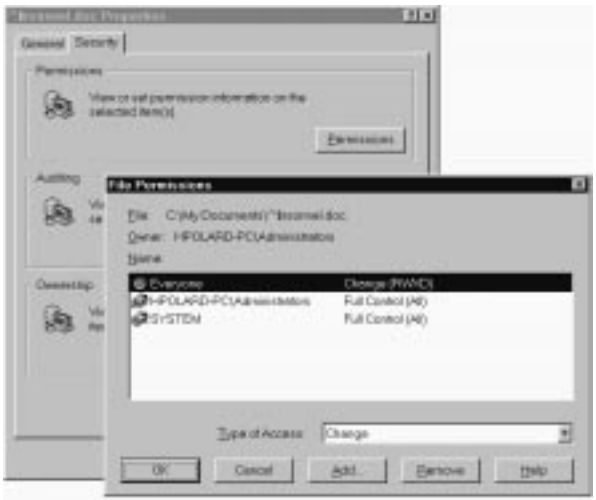

Changing UNIX permissions and DOS attributes from Windows

How to change UNIX permissions

Some files and directories on the filer have both Windows and UNIX-style permissions. You can change Windows-style permissions from Windows by editing Windows permissions. To change UNIX permissions in a UNIX file system from the Windows desktop or to change DOS attributes, you use the SecureShare Access tool.

To use the SecureShare Access tool from a client, you must install the tool on the client. For instructions about downloading the SecureShare Access tool, see the *Software and Firmware Upgrade Guide* or *Start Here*.

Note

To change UNIX permissions, you must understand them. Explaining UNIX permissions is outside the scope of this guide. Consult literature about UNIX for an explanation of UNIX permissions.

Displaying SecureShare Access

To display SecureShare Access, follow these steps.

Step	Action	Result
1	Select the files and directories whose permissions you want to change.	The items are highlighted.
2	Right-click one of the items you selected.	A pop-up menu appears.
3	Choose Properties from the pop-up menu.	The Properties dialog box appears.
4	Click the SecureShare tab.	SecureShare Access appears.

Changing the permissions of a single item

If you select only one item, SecureShare Access appears as follows:

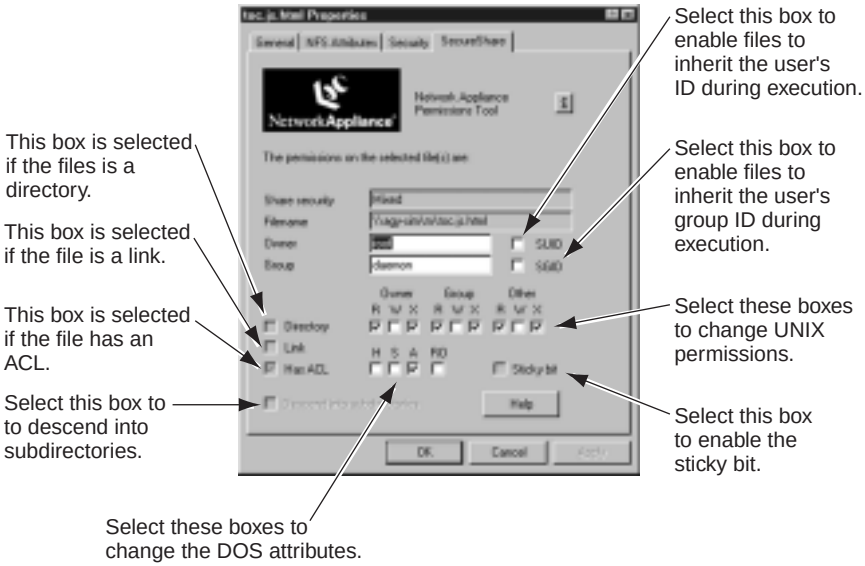

Caution

SecureShare Access has no undo feature. Use it very carefully.

Qtree security style effects: The security style of the qtree from which you select an item has the following effects:

- In NTFS-style qtrees, all the edit fields and check boxes are disabled. Use standard Windows NT tools to manipulate NTFS files.
- In NFS-style qtrees, the Has ACL check box is meaningless and is disabled.
- In mixed-style qtrees, if a file has an ACL, the ACL is removed when you click the OK or Apply button. A dialog box prompts you to confirm that this is what you want.

Note

If you click OK or Apply on any page or tab of a Properties sheet, it puts into effect all changes that you made on all pages or tabs of the Properties sheet.

Recursive application of changes: If one or more of the files is a directory, the Descend Into Subdirectories check box is enabled. Select this check box to apply the changes you make recursively to the contents of any directory you selected. Links are not changed, but their targets are.

Changing the permissions of multiple items

You can change the permissions of several items at once by selecting them, then displaying SecureShare Access by following Steps 2 through 4 in "Displaying SecureShare Access" on page 308.

In the multiple-item display, the following conditions are in effect that are different than the single-item display.

- Permissions; DOS attributes; or characteristics; for example, whether an item is a directory, that are common to all the items you selected appear with a white background.
- Permissions; DOS attributes; or characteristics; for example, whether an item is a directory, where one item differs from the rest of the items you selected appear with a light gray background.
- Permissions or DOS attributes that you set apply to all items you selected.
- Turning one of the check boxes gray preserves the permissions on any of the selected items.

Sending a message to all users on a filer

When to send a message

You might want to send a message to all users on a filer to tell them of important events. The message appears in an alert box. For example, you might need all users to close any open files on the filer, but not want to stop CIFS services and, therefore, not take advantage of the messaging function of the `cifs terminate` command.

How to send the message

To send a message to all users on a filer, follow these steps:

Step	Action
1	Open Server Manager by choosing it from the Start menu.
2	From the Server Manager Computer menu, choose Select Domain. **Result:** The Select Domain window opens.
3	In the Select Domain window, select the filer you want by typing its UNC name, for example, \\FILERNAME, in the Domain field, then clicking OK. **Result:** A Server Manager window for the filer appears.
4	In the Server Manager window, from the Computer menu, choose Send Message. The Send Message window appears. 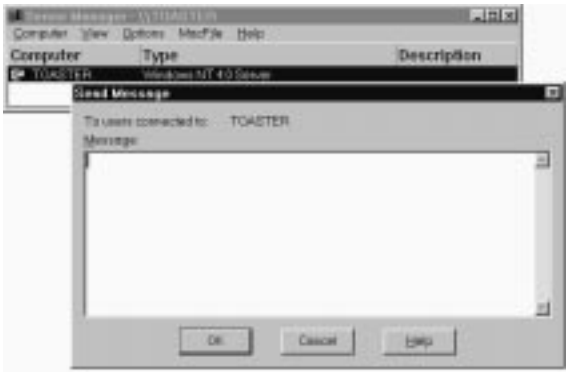
5	Enter a message in the text box of the Send Message window.

Step	Action
6	Click OK to send the message.

Event auditing

You can audit file access events

Data ONTAP enables you to monitor reads and writes of a specified file on the filer by a specified user. The procedure for doing so is the same as for Windows NT.

Note

The file on the filer must be in a mixed or ntfs volume or qtree. You cannot audit events on a file in a UNIX volume or qtree.

You can specify the logging of successes, failures, or both for each type of event as described in "Setting a system ACL on a file" on page 321. You can view the events on the filer using the Windows Event Viewer, as described in "Viewing events in a security log" on page 325.

Why use event auditing

You use event auditing to troubleshoot access problems, to check for suspicious activity on a system, or to investigate a security breach.

- If someone who should have access to a file cannot get access, examining the event can provide a clue to resolving the problem.
- If a file has been changed in a way that indicates a security breach, event auditing might provide clues about the nature of the breach if subsequent changes occur.

Active event log naming

The active event log is the file to which system writes access logging information. The default active event log is */etc/log/adtlog.evt*. You can specify another active event log with the `cifs.access_logging.filename` option. You can change active event logs at any time. The other logs remain available for reading until you delete them.

Log access

Anyone can read an event log. You cannot write to an event log or clear one, but you can delete one.

Event log detail displays

How to examine an event in detail

You examine an event in detail by double-clicking it and looking at the resulting event details display. There are the following kinds of event details displays:

- Windows file access detail displays
- UNIX file access detail displays
- Unsuccessful file access detail display
- Lost record event detail display

Windows file access detail displays

Windows file access detail displays show the following information displayed in the detail screens.

Field	Description
Object Server	The name of the subsystem server process calling the audit check function. This is always SECURITY, because this is a security log.
Object Type	The type of object being accessed.
Object name	The name (such as a file name) of the object being accessed.
New Handle ID	The new handle identifier of the open object.
Operation ID	A unique identifier associating multiple events resulting from a single operation.
Process ID	The identifier of the client process accessing the object.
Primary User Name	The user name of the user requesting the object access. When impersonation is taking place, this is the user name with which the server process is logged on.

Field	Description
Primary Domain	The name of the computer, or SYSTEM if the user identified by Primary User Name is SYSTEM. If the computer is a member of a Windows NT Server domain, this can also be the name of the domain containing the primary user's account.
Primary Logon ID	A unique identifier assigned when the primary user logged on.
Object User Name	Your login.
Client Domain	The name of your computer or the domain containing the client user's account.
Client Logon ID	A unique identifier assigned when the client user logged on.
Accesses	The types of accesses to the object that were attempted.
Privileges	Your privileges.

UNIX file access detail displays

UNIX file access detail displays show the same kind of information as the Windows file access event detail screens, but instead of an object name, NFS access appears because the file is being accessed through NFS. In addition, UNIX file access detail displays show the following information about the file that you are monitoring:

- The ID of the volume in which the file is located
- The ID of the latest snapshot in which the file is located
- The inode of the file

This information enables you to find the file using the `ls -i` command.

Unsuccessful file access detail display

An unsuccessful file access detail display appears when a user could not access a file. For example, an unsuccessful file access occurs when a user tries to access a file but does not have permission to access it.

The display shows the ID of the user who tried to access the file and an indication that the access attempt was unsuccessful.

Lost record event detail display

If the system could not create an audit record, the Lost record event detail display gives a reason, such as the following reason:

```
Internal resources allocated for the queueing of audit messages
have been exhausted, leading to the loss of some audits.
      Number of audit records discarded:           1
```

Event auditing overview

Description

This overview of how to audit events contains steps that are described in detail in later sections.

Steps

To audit events, complete the following steps.

Step	Action
1	Enable CIFS access logging, as described in "Enabling CIFS access logging" on page 318.
2	Set a system ACL on the files you want to audit, as described in "Setting a system ACL on a file" on page 321. This involves specifying the users or groups whose access to the file you want to audit.
3	Use the Event Viewer, as described in "Viewing events in a security log" on page 325 to view the security log.

Enabling CIFS access logging

Description

Use this procedure to enable CIFS access logging. You follow this procedure when you want to start logging file access events by a specified user on a particular file. After you complete this procedure, the filer is ready to process access logging information.

Step

To enable CIFS access logging, complete this step.

Step	Action
1	Enter the following command: `options cifs.access_logging.enable on`

Disabling CIFS access logging

Description

Use this procedure to disable CIFS access logging. You follow this procedure when you want to stop logging file access events and free the resources that are used by access logging. After you complete this procedure, access logging stops.

Step

To disable CIFS access logging, complete this step.

Step	Action
1	Enter the following command: `options cifs.access_logging.enable off`

Specifying the active event log

Description

Use this procedure to specify the active event log. You follow this procedure when you want to specify the name of a file to which the system writes access logging information. You can do this at any time. After you complete this procedure, if access logging is enabled, the system writes access logging information to the file you specified.

Prerequisites

You must have the path name of the file to use. This file must be in an existing writable directory.

Step

To specify the active log, complete this step.

Step	Action
1	Enter the following command at the command line: **`options cifs.access_logging.filename`** ***`path_name`*** *path_name* is the path name in UNIX format of the file to which you want the system to write access logging information.

Setting a system ACL on a file

Description

Use this procedure to sets a system ACL (SACL) on a file. You must complete this procedure to monitor access activity on a file. You complete this procedure when you have decided to monitor access to a file by specified users or groups. The procedure is the standard Windows NT procedure for setting a SACL on a file. After you complete this procedure, if access logging is enabled, the filer logs accesses to the file by the users or groups you specified.

Prerequisites

You must have the following items to complete the procedure:

- The name of the file you want to monitor
- The name of the users or groups whose access to the file you want to monitor

Steps

To set a SACL on a file, complete these steps.

Step	Action
1	Right-click the file you want to monitor. **Result:** The Properties tab for that file appears.
2	Click the Security tab. **Result:** The Security properties window appears.

Step	Action
3	Click the Auditing button. **Result:** A blank File Auditing window appears for the file you specified.
4	Click the Add button. **Result:** The Add Users and Groups window appears.
5	Click the arrow next to the List Names From text box to choose a domain that contains names that you want to add. **Result:** A list of names in the selected domain appears in the Names list box.

Step	Action
6	To add a name, type one or more user names in the Add Names list box or select one or more names and click Add.
7	Specify one or more names in one of these ways: ◆ Click one or more names in the Names list box. ◆ Type valid user or group names in the Add Names list box.
8	Click Add in the window where you selected the names. **Result:** The names are added to the Add Names field of the Add Users and Groups window.
9	If you want, display the full name of a user associated with an account name by clicking Show Full Name.

Step	Action
10	Click OK. **Result:** The names appear in the Name field of the File Auditing window and the Events to Audit check boxes are enabled, as shown below. In the following figure, checks indicate the boxes that you can select. 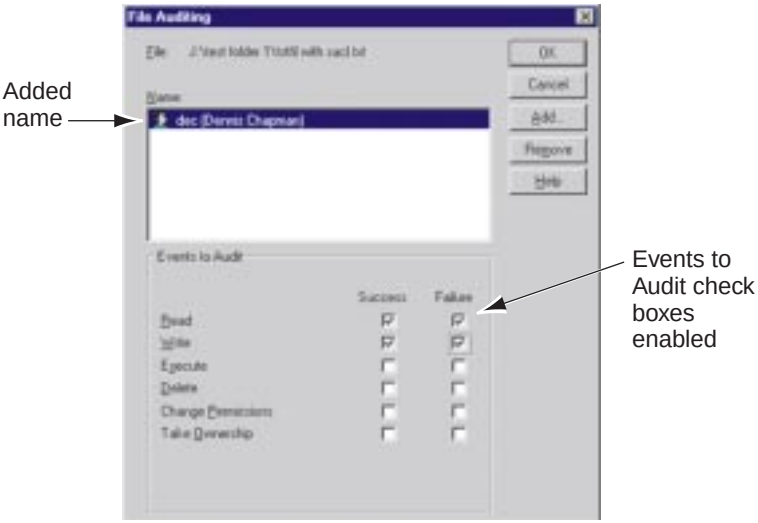 **Note** Execute, Delete, Change Permissions, and Take Ownership events are not currently supported.

Viewing events in a security log

Description

Use this procedure to view events in a security log and check the events on the file and users you specified in the procedure "Setting a system ACL on a file" on page 321. After you complete this procedure, access information is displayed.

Prerequisites

You should have the name of the security log that you want to view.

Steps

To view events in a security log, complete these steps.

Step	Action
1	Open the Event Viewer by choosing the menus shown below. 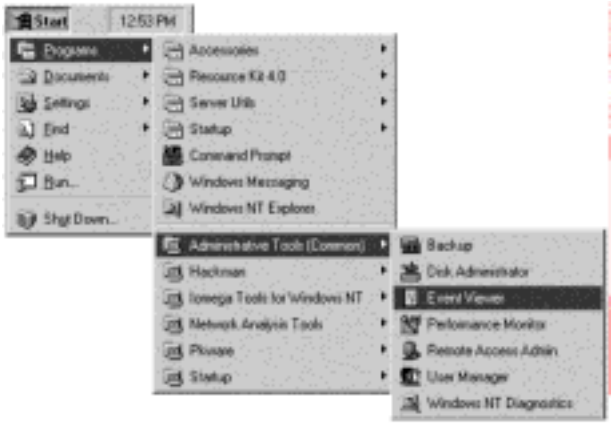 **Result:** The event viewer appears.

Step	Action
2	From the Log menu of the Event Viewer, choose Open. **Result:** An Open window appears.
3	Choose the share that contains the log you want to look at. **Result:** A list of files appears.
4	Click the log file you want to examine. An Open File Type window appears. 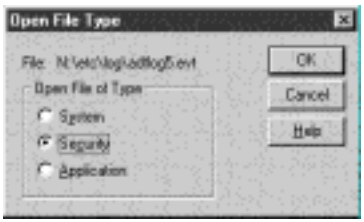
5	If Security is not selected, select it.

Step	Action
6	Click OK. **Result:** The Event viewer appears and displays a list of Security events as shown below. 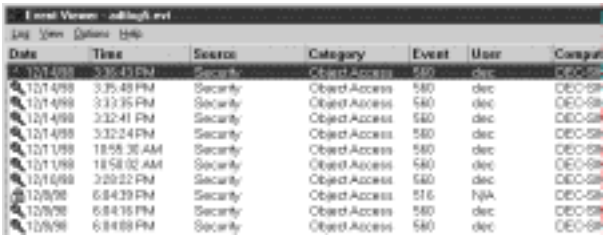 In the display, symbols preceding event entries have the following meanings: ◆ Key - Successful access attempts ◆ Lock -Unsuccessful access attempts
7	To view an event detail, double-click the event. **Result:** An event detail screen appears.

Using oplocks

What oplocks do

Oplocks (opportunistic locks) enable the redirector on a CIFS client in certain file-sharing scenarios to perform client-side caching of read-ahead, write-behind, and lock information. A client can then work with a file (read or write it) without regularly reminding the server that it needs access to the file in question. This improves performance by reducing network traffic.

When to use oplocks

Under some circumstances, if a process has an exclusive oplock on a file and a second process attempts to open the file, the first process must relinquish the oplock and access to the file. The redirector must then invalidate cached data and flush writes and locks, resulting in possible loss of data that was to be written.

Data loss possibilities: Any application that has write-cached data can lose that data under the following set of circumstances:

- It has an exclusive oplock on the file.
- It is told to either break that oplock or close the file.
- During the process of flushing the write cache, the network or target system generates an error.

Error handling and write completion: The cache itself does not have any error handling—the applications do. When the application makes a write to cache, the write is always completed. If the cache, in turn, makes a write to the target system over a network, it must assume that the write is completed because if it does not, the data is lost.

When to turn oplocks Off

CIFS oplocks on the filer are On by default.

You might turn CIFS oplocks Off under either of the following circumstances:

- You are using a database application whose documentation recommends that oplocks be turned Off.
- You are handling critical data; that is, you have a good network but you cannot afford even the slightest data loss.

Otherwise, you can leave CIFS oplocks On.

Turning oplocks On and Off globally

You can turn CIFS oplocks On or Off globally for the entire filer or for individual qtrees, which are special directories that are described in detail in Chapter 12, "Qtree administration."

Turning oplocks Off: You turn all CIFS oplocks Off with the following `options` command:

```
options cifs.oplocks.enable off
```

Turning oplocks On: You turn CIFS oplocks On with the following `options` command:

```
options cifs.oplocks.enable on
```

Turning oplocks On or Off at individual clients

You can turn CIFS oplocks On or Off at individual clients. Turning CIFS oplocks On at the filer does not override any client-specific settings. Turning CIFS oplocks Off at the filer disables all oplocks to or from the filer.

For additional information: For additional information about oplocks, consult the Microsoft Knowledge Base at *http://www.microsoft.com/kb*.

Displaying CIFS statistics

How and why to display CIFS statistics

You use the `cifs stat` command to display statistics about CIFS operations that take place on your filer. You use the `cifs stat` display for diagnostic purposes.

Statistics displays with the cifs stat command

You can use the `cifs stat` command in two forms:

- If you specify a time interval, the command displays statistics at the specified intervals.
- If you do not specify a time interval, the command displays CIFS statistics that have accumulated since the last reboot.

Example of cifs stat output: The following command displays statistics every second:

```
cifs stat 1
    GetAttr   Read   Write   Lock   Open/Cl   Direct   Other
      175    142      54     70       254      643      50
      232     76     123     44       321      154      17
      152    120      34    111        12      435      76
```

The following table describes the fields.

Field	Description
`GetAttr`	Attribute operations.
`Read`	Read data operations.
`Write`	Write data operations.
`Lock`	Lock operations.
`Open/Cl`	Open and close operations.
`Direct`	Directory operations.
`Other`	Other operations, such as deletes and logoffs.

Displaying CIFS session information

CIFS session information you can display

You can display information about connected users and the number of shares and open files for each user. You can also display information about a specific connected user.

Displaying information with the cifs sessions command

The `cifs sessions` command syntax is as follows:

`cifs sessions <username>`

Displaying information about all connected users: To display information about all connected users, use the following command syntax:

```
cifs sessions
```

Sample output is

```
Server Registers as 'TOASTER' in group WNT-DOMAIN
WINS Server: 272.320.0.4
PC style Access Control is being used
Using domain controller WNT-DOMAIN-PDC for authentication
========================================
PC (user)           #shares       #files
SMITHPC (qsmith)      1             1
PETERSPC (zpeters)    2             3
```

Displaying information about one user: To display the information for a particular user, specify the user name in the command; for example:

```
cifs sessions ghopper

users
    shares/files opened

  HALEY-HOME1   (ghopper)
    ENG-USERS
             Read-denyW   -
\GHOPPER\SRC\PROD\COMMON\HTTPD\httpd_fast.c


  HALEY-PC      (ghopper)
    ENG-USERS
```

Displaying connected user security information: To display security information for each connected user, use `cifs sessions` with the `-s` option. After the first two lines, detailed information for each connected user is displayed. The following example lists only one user.

```
cifs sessions -s

users
  Security Information

HOLARD-PC     (root)
        ***************
        UNIX uid = 0
        user is a member of group daemon (1)
        user is a member of group www (204)
        user is a member of group well (0)
        user is a member of group http (500)

        NT membership
                WNT-DOMAIN\root
                WNT-DOMAIN\Domain Users
                WNT-DOMAIN\Domain Admins
                WNT-DOMAIN\SU Users
                WNT-DOMAIN\Installers
                BUILTIN\Users
                BUILTIN\Administrators
        User is also a member of Everyone, Network Users,
        Authenticated Users
        ***************
```

Stopping and restarting CIFS sessions

Ways to stop CIFS sessions

If you want to stop CIFS sessions for all clients or for a single client, do one of the following:

- From Windows NT, use Server Manager as described in the next section, "Disconnecting users with Server Manager."
- From the filer command line, use the `cifs shares` command to display share information, as described in "Using the cifs terminate command" on page 334.

Disconnecting users with Server Manager

To stop CIFS sessions with Server Manager, follow these steps.

Step	Action
1	Open Server Manager by choosing it from the Start menu.
2	From the Server Manager Computer menu, choose Select Domain. **Result:** The Select Domain window opens.
3	In the Select Domain window, select the filer you want by typing its UNC name, for example, \\FILERNAME, in the Domain field, then clicking OK. **Result:** A Server Manager window for the filer appears.
4	In the Server Manager window, double-click the name of a filer. **Result:** The Properties window for the filer appears.

Step	Action
5	Click Users. **Result:** The User Sessions window appears. 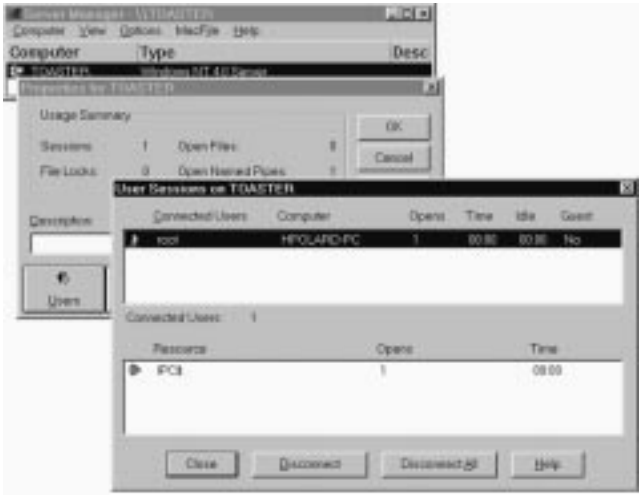
6	To disconnect one or more users, do one of the following: ◆ To disconnect a single user or selected users, select them, then click Disconnect. ◆ To disconnect all users, click Disconnect All. **Result:** The selected users are disconnected. **Note** If at least one of the selected users has open resources, an alert box appears for you to confirm or cancel your command.

Using the cifs terminate command

You can stop CIFS service for a specific client or shut down CIFS service from the filer by using the `cifs terminate` command. Always terminate all CIFS sessions before you reboot or turn Off the filer.

You can specify a single client or all clients, and the time delay, in minutes, before the CIFS sessions are terminated, as shown in the following command syntax:

`cifs terminate [client] [[-t] time]`

Variable	Description
client	Name of the client for which you are ending a CIFS session.
time	Number of minutes before the termination happens.

The cifs terminate command not persistent

The `cifs terminate` command disables CIFS sessions only between the time you enter the command and the next reboot. After each reboot, if your filer is licensed and configured to run CIFS, the filer automatically starts CIFS service. If you want to prevent CIFS from restarting after a reboot, remove the */etc/cifsconfig.cfg* file from the filer or rename the file.

Time delay

You can delay the termination of CIFS service after you enter the `cifs terminate` command.

Default time delay: The default time delay is five minutes.

Changing the time delay: When you use the `-t` option, the command counts down from the time specified.

Canceling the cifs terminate command

If you want to cancel the `cifs terminate` command, press Ctrl-C before the end of the countdown.

Note

The `halt` command automatically invokes the `cifs terminate` command.

Caution

The `reboot` command also stops CIFS service; however, it does not provide a time delay during which users can save their open files before the disconnect. Changes that have not been saved to disk are lost if the CIFS client has an open file when it is disconnected from the filer.

Examples of the cifs terminate command

Here are some examples of the `cifs terminate` command.

Terminating CIFS service for all users on the filer: To terminate CIFS service for all users on the filer after 10 minutes, enter the following command:

```
cifs terminate -t 10
```

The `cifs terminate` command displays an alert message on CIFS clients that warns them of the pending shutdown of CIFS service.

Note

Windows 95 and Windows for Workgroup clients must have the WinPopup program configured before they can display the alert message.

Console display: Here is the display for `cifs terminate` when you do not specify a time:

```
cifs terminate

There are currently 35 CIFS users that have 37 open files
Disconnecting while files are open may cause data loss!!
How many minutes should I wait? [5]:
minutes left 4
minutes left 3
minutes left 2
minutes left 1
CIFS shutting down
```

Terminating a CIFS session for a specific client: To terminate a CIFS session for a particular client, specify the name of the computer in the command. For example, the following command terminates a CIFS session for a computer named PETERSPC after 10 minutes:

```
cifs terminate PETERSPC -t 10
```

Using the cifs restart command to restart CIFS service

To restart CIFS service, use the `cifs restart` command, as follows:

```
cifs restart

CIFS server is registering...
CIFS server is running.
```

Reconfiguring the filer for CIFS

When to reconfigure a filer for CIFS

You can reconfigure the filer for CIFS service at any time, for example, if you want to change the authentication method from PDC (Primary Domain Controller) to UNIX password database.

How to reconfigure a filer for CIFS

To reconfigure a filer for CIFS, follow these steps:

1. Enter the `cifs terminate` command to stop CIFS service.
2. Enter the `cifs setup` command to reconfigure CIFS service.
   - If you enter only valid information, when you exit the program, the filer automatically restarts CIFS using the new configuration information.
   - If you enter some invalid information, for example, you mistype a domain name, when you exit the program, the filer restarts CIFS with the previous configuration.

HTTP Administration 9

Chapter contents

This chapter describes how to perform the following tasks.

Read...	To learn about...
"Starting HTTP service" on page 340	Starting HTTP service.
"Protecting Web pages with passwords" on page 342	Protecting Web pages with passwords.
"Using the HTTP virtual firewall" on page 345	Using the HTTP virtual firewall.
"Using virtual hosting" on page 346	Enabling virtual hosting.
"Specifying MIME Content-Type values" on page 348	Specifying MIME Content-Type values.
"Translating URLs" on page 350	Translating URLs.
"Displaying HTTP connection information" on page 353	Displaying HTTP connection information.
"Displaying HTTP statistics" on page 354	DISPLAYING HTTP statistics.

Note

You can use the filer as an HTTP server *only* if you purchased the license for HTTP. Without the license, you can use an HTTP client (Web browser) only to display the filer's man pages and to use FilerView.

Starting HTTP service

Procedure for starting HTTP service

To start HTTP service on your filer, follow these steps:

Step	Action
1	Enable the `httpd` daemon by entering the following command: **`options httpd.enable on`**
2	Use the following command syntax to specify the root directory that contains the files and directories to be read by HTTP clients: `options httpd.rootdir` *`directory`* For example, if the root directory is */vol/vol0/home/users/pages*, enter the following command: **`options httpd.rootdir /vol/vol0/home/users/pages`**
3	If you want to limit the size of the */etc/log/httpd.log* log file to other than the default of 2,147,483,647 bytes (2 GB minus 1 byte), use the following command: `options httpd.log.max_file_size` *`bytes`*
4	Make a copy of */etc/httpd.mimetypes.sample* and name the copy */etc/httpd.mimetypes*. If the */etc/httpd.mimetypes* file is missing, the HTTP client uses the information in */etc/httpd.mimetypes.sample*.

Note

If you want these options to remain active after rebooting, you must add them to the */etc/rc* file.

The procedure for starting HTTP service is now complete, and clients can display text files under the root directory by using a Web browser. If the filer will transfer files other than text files, for example, image files and audio files, follow the instructions in "Specifying MIME Content-Type values" on page 348 to configure your filer so that the appropriate MIME Content-Type header is sent with each file transferred.

Procedure for testing HTTP service

To test the filer's HTTP service, follow these steps.

Step	Action
1	Create an HTML file in the root directory for HTTP. For example, create a file named *myfile.html* in the HTTP root directory, which is */vol/vol0/home/users/pages*, assuming that the HTML root directory is */vol/vol0/home/users/pages*.
2	Start a Web browser on a client and specify the URL of the HTML file in the browser. For example, if your filer is *toaster* and the root directory for HTTP is */vol/vol0/home/users/pages*, enter this URL: `http://toaster/myfile.html` The path component of the URL is a path name relative to the HTTP root. Do not specify the complete path name to the file in the URL.

Note

If the URL names a directory, for example, *http://toaster/home/pages*, the filer automatically tries to transfer the *index.html* file within the directory. If *index.html* does not exist, the filer returns "Error 404. No such file or directory."

Protecting Web pages with passwords

Configuration files for password protection

You can restrict access to a specified directory so that only specified users or groups have access to it.

Password protection involves three configuration files:

- */etc/httpd.access*
- */etc/httpd.passwd*
- */etc/httpd.group*

The /etc/httpd.access file

The */etc/httpd.access* file contains directives that govern authentication for each directory. The filer supports the following directives:

- `directory`
- `AuthName`
- `require user`
- `require group`

These directives are compatible with the Apache Web server directives, but the filer ignores all other directives.

The directory directive: Specifies a directory tree to be protected and encloses all other directives. The syntax of the `directory` directive is as follows:

```
<Directory directory_name>
directive ...
</Directory>
```

The AuthName directive: Specifies a "realm," that is, an alias for the directory that appears instead of the directory name in the browser's password dialog box when a user tries to access the directory. Whatever follows `AuthName` is the name of the realm. The name of the realm can contain spaces. The syntax of the `AuthName` directive is as follows:

```
AuthName realm name
```

The require user directive: Specifies the users who can access the directory. The syntax of the `require user` directive is as follows:

```
require user user_id [, user_id, ... ]
```

The require group directive: Specifies the groups that can access the directory. The syntax of the `require group` directive is as follows:

```
require group group_id [, group_id, ... ]
```

The /etc/httpd.passwd file

The */etc/httpd.passwd* file contains the *user_id* and encrypted-password pairs. The pairs have the following format:

```
user_id:encrypted_passwd
```

The pairs are copied in from a machine on which the user has a password.

The /etc/httpd.group file

The */etc/httpd.group* file contains a *group_id* and a list of *user_ids* in that group in the following format:

```
group_id: user_id [user_id  ....]
```

The lists are copied in from a machine that has a similar list.

Web page protection examples

The following */etc/httpd.access* file restricts access to */vol/vol0/home/htdocs/private/spec* to only the user bob:

```
<Directory /vol/vol0/home/htdocs/private/spec>
AuthName polard Private Stuff
<Limit GET>
require user bob
</Limit GET>
</Directory>
```

The <Limit GET> and </Limit GET> directives, which might have been imported from an Apache or NCA Web server, are not supported. To be used on a filer, the file does not need to be edited to remove the Limit GET directive; the filer ignores the directive and all other directives not mentioned in this chapter.

The following sample procedure restricts user access to a particular directory:

1. Enter the following lines in the */etc/httpd.access* file:

```
<Directory /vol/vol0/home/htdocs/private/specs>
```

```
AuthName Social commentary
require group engineering
</Directory>
```

2. Enter the following line in */etc/httpd.group*:

```
engineering: bob larry nancy rose
```

The */vol/vol0/home/htdocs/private/specs* directory is now accessible only to the group engineering, which consists of the following user IDs:

- bob
- larry
- nancy
- rose

Using the HTTP virtual firewall

About the HTTP virtual firewall

The HTTP virtual firewall feature enables you to maintain security on your filer.

You can restrict HTTP requests by marking the subnetwork interface over which they arrive as “untrusted.” An untrusted interface provides only HTTP access to your filer on a read-only basis.

Mark an interface untrusted if it meets all the following conditions:

- You know you are going to service HTTP requests over that interface.
- You don’t want to allow requests through protocols other than HTTP.
- You want to restrict access to the filer through that interface to read-only access.

By default, a subnetwork interface is trusted.

Syntax

Mark an interface as untrusted or trusted by setting an option to the `ifconfig` command. Following are examples of the command:

- To mark the *f0* interface as untrusted, enter the following command:

  `ifconfig f0 untrusted`

- To mark the *f0* interface as trusted, enter the following command:

  `ifconfig f0 trusted`

Using virtual hosting

About virtual hosting

Virtual hosting enables a filer to respond to requests directed to more than one IP address through a single physical interface. This means that a filer with only one physical interface can host several IP addresses.

Virtual hosting enables, for example, an Internet provider to host several Web sites but have only one physical interface. An HTTP server can use the destination IP address of an incoming HTTP request to find the directory that contains the HTTP pages belonging to the virtual host.

To set up and enable virtual hosting

To enable virtual hosting, you

- direct HTTP requests by putting subdirectory and host or address entries in the */etc/httpd.hostprefixes* file
- map virtual host addresses to the virtual host interface with the `ifconfig` command

Directing HTTP requests

To direct HTTP requests, use the following format in the */etc/httpd.hostprefixes* file:

`prefix` `[`*`host-name-or-address ...`*` ]`

Variable	Description
prefix	Specifies a subdirectory in the HTTP root directory, which is defined by the `options httpd.rootdir` command.
host-name-or-address	Specifies an HTTP host name or an IP address. You can have more than one of each.

For example, the line

```
/customer www.customer.com 192.225.37.102
```

means that an HTTP request that comes for the interface with address 192.225.37.102, or with an HTTP 1.1 Host: header specifying *www.customer.com*, is directed to */customer*, and the requestor cannot get a file outside the */customer* directory.

If the HTTP server receives an HTTP request that is destined for one of its virtual host IP addresses, in this example 192.225.37.102, the destination IP address is used to select the virtual host root directory from the */etc/httpd.hostprefixes* file.

Mapping virtual host addresses

To map virtual host addresses to the virtual host interface, use the `ifconfig` command, as follows:

- Add a new IP virtual host address mapping with the following command:

  ```
  ifconfig vh alias address
  ```

  where *`address`* is an IP address.

  The use of the `vh` interface indicates to the system that you are adding a virtual host address rather than adding an IP alias address to a network interface.

- Delete virtual host addresses with the following command:

  ```
  ifconfig vh -alias address
  ```

Note

If you need to create a virtual subnet with many contiguous addresses, the IP address can be a subnet address.

Specifying MIME Content-Type values

About MIME Content-Type values

You can configure the filer to send the appropriate MIME (Multipurpose Internet Mail Extensions) Content-Type header in each response to a get request. The header shows the MIME Content-Type value of the file, which tells the browser on the client how to interpret the file.

For example, if the MIME Content-Type value shows that the file being transferred is an image file and the client is configured properly, the browser can render the image by using a graphics program.

The filer determines the MIME Content-Type value of a file by mapping the file name suffix, or example, *.gif*, *.html*, or *.mpg*, according to information in the */etc/httpd.mimetypes* file.

Note

On a Windows 95 or Windows NT 4.0 client, the */etc/httpd.mimetypes.sample* file name is not displayed in its entirety. By default, the Explorer displays the file name as */etc/httpd.mimetypes*. If you are using Windows, from the Explorer View menu, select Options, then the View tab and, in the dialog box, make sure that there is no check mark in the check box next to “Hide MS-DOS file extensions for file types that are registered.”

Modifying MIME Content-Type mappings

To modify MIME Content-Type mappings or to add MIME Content-Types, edit the entries in */etc/httpd.mimetypes*. Entries are in the following format:

```
# An optional comment.
suffix   Content-Type
```

Lines preceded by the # sign are comments. The file name suffix is case-insensitive. Following are sample entries:

```
# My clients’ browsers can now use
# PICT graphics files.
pct        image/pict
pict       image/pict
```

In the sample entries, files whose names ended with *.pact* or *.pact* are mapped to the MIME Content-Type value of image/pict. The first field in the Content-Type value describes the general type of data contained in the file; the second field is

the data subtype, which shows the specific format in which the data is stored. If the browser on the client is configured to start a graphics program as a helper application, the user can view a file named *file.pict* as a graphics file on the client.

For more information about MIME, refer to RFC 1521.

Translating URLs

How the filer responds to URLs

You can specify that the filer's response to an HTTP request be dependent on the URL. For example, you can configure the filer to redirect a particular request to a specific directory, or to prevent access to a particular directory that is specified in the URL.

How the filer maps its responses to URLs is defined in a configuration file named */etc/httpd.translations*. Each entry in the configuration file contains up to three fields in the following format:

```
rule template result
```

Field	Description
`rule`	Defines the response of the filer to a request.
`template`	Specifies a component of a URL.
`result`	Depends on the rule, as described in the following section.

Translation rules supported by the filer

This section explains the meanings of the rules. It also describes the format required for each type of entry in */etc/httpd.translations*.

The map rule: The map rule specifies that if a component of a URL matches the template, the request is mapped to another directory on the same host as defined in the `result` field.

For example, the following */etc/httpd.translations* entry causes any requests to a URL containing the */image-bin* directory to be mapped to the */usr/local/http/images* directory:

```
map /image-bin/* /usr/local/http/images/*
```

The redirect rule: The redirect rule specifies that if a component of a URL matches the template, the request is redirected to the URL defined in the `result` field. The `result` field for the redirect rule must be specified as a complete URL beginning with *http://* and the host name.

For example, if */etc/httpd.translations* contains the following entry

```
redirect /cgi-bin/* http://cgi-host/*
```

the filer redirects CGI requests to another HTTP server named cgi-host. This is essential for calls to *cgi-bin* because the filer does not execute them.

The pass rule: The pass rule specifies that if a component of a URL matches the template, the filer accepts the request, processes the request as is, and disregards other rules.

For example, if */etc/httpd.translations* contains the following entry

```
pass /image-bin/*
```

the filer processes the request for any URL containing */image-bin* as is, even though there is another rule specified as follows:

```
map /image-bin/* /usr/local/http/images/*
```

If the pass rule includes the `result` field, the filer accepts the request, processes the request by using the URL defined in the `result` field, and disregards other rules.

The fail rule: The fail rule specifies that if a component of a URL matches the template, the filer denies access to that component and disregards other rules.

For example, if */etc/httpd.translations* contains the following entry

```
fail /usr/forbidden/*
```

the filer does not provide access to the */usr/forbidden* directory.

How the filer processes rules

The filer processes the rules defined in */etc/httpd.translations* in the order they are listed, and applies the rule if the URL matches the template. However, the filer stops processing other rules after it applies a pass or fail rule.

In the `template` or `result` field of an */etc/httpd.translations* entry, you can use asterisks (*) as wildcard characters, as follows:

- In the `template` field, the wildcard character matches zero or more characters, including the slash (/) character.
- In the `result` field, the wildcard character represents the text expanded from the match in the `template` field. Include the wildcard character in the `result` field only if you used a wildcard character in the `template` field.

- If you use multiple wildcard characters, the first one in the `result` field corresponds to the first one in the `template` field, the second one in the `result` field corresponds to the second one in the `template` field, and so on.

Following is an example showing how a wildcard character is used:

```
# Redirect all cgi requests to my cgi server
redirect /cgi-bin/*    http://cgi-host/cgi-bin/*
```

This redirect rule specifies that all CGI requests are redirected to another host named cgi-host.

For example, if the filer receives the following requests

```
http://toaster/cgi-bin/displayfares
http://toaster/cgi-bin/displaydates
```

the filer expands the wildcard character to displayfares and displaydates and redirects the requests to the host named cgi-host. To the client, the results of these requests are the same as the results of the following requests:

```
http://cgi-host/cgi-bin/displayfares
http://cgi-host/cgi-bin/displaydates
```

Displaying HTTP connection information

Information in the /etc/log/httpd.log file

You can read the */etc/log/httpd.log* file if you are interested in the following types of information for each HTTP connection:

- IP address of HTTP client.
- Name of authorized users, if the requested page is protected, making requests. The names are in the */etc/httpd.passwd* file. If the page is not protected, dashes appear instead of a name.
- Time of connection in `dd/mm/yy:hh:mm:ss` format; gmt is used.
- Request line from connecting host, for example, `get /netapp.html`.
- STATUS code returned by the server, as defined in the HTTP 1.0 specifications.
- TOTAL bytes sent in response by the filer, not including the MIME header.

Following is an example of the */etc/log/httpd.log* file:

```
192.9.77.2  - - [26/Aug/1996:16:45:50] "GET /top.html" 200 1189
192.9.77.2  - - [26/Aug/1996:16:45:50] "GET /header.html" 200 531
192.7.15.6  - - [26/Aug/1996:16:45:51] "GET /logo.gif" 200 1763
198.9.200.2 - - [26/Aug/1996:16:45:57] "GET /task/top.html" 200 334
192.9.20.5  authuser [26/Aug/1996:16:45:57] "GET /task/head.html"
200 519
```

Displaying HTTP statistics

httpstat statistic types

The `httpstat` command displays four types of statistics about HTTP operations on the filer:

Column	Description
`gets`	Successful requests for files.
`badcalls`	Requests for nonexistent files.
`open conn.`	Number of HTTP connections currently open.
`peak conn.`	Largest number of simultaneous HTTP connections since the filer was booted or since the `-z` option was used.

Syntax

The syntax for the `httpstat` command is as follows:

```
httpstat [ -t|-z ][ interval ]
```

If you use no arguments, `httpstat` displays HTTP statistics accumulated since the last reboot or since the last time the `-z` argument was used.

The `-z` argument resets both the gets and badcalls counters.

The `-t` argument displays statistics since the last filer reboot.

You can specify the interval, in seconds, at which the filer displays the statistics. Following is an example of `httpstat`:

```
httpstat
HTTPD statistics:
gets          badcalls          open conn.    peak conn.
451           11                 5             17
```

Cluster Administration 10

Introducing clusters

About clusters

A cluster consists of a pair of NetApp filers that are connected and configured in a special way to provide fault tolerance. When one filer in a cluster undergoes a system failure and can't reboot, the partner filer in the cluster can take over the failed filer's functions and serve the data on the failed filer's disks. This is known as a takeover. In addition to an automatic takeover, an operator may issue a `cf takeover` command at any time to force a takeover.

The filers in a cluster have the following characteristics:

- They are connected to each other through a Cluster Interconnect adapter through which they:
  - constantly check whether the other filer is functioning
  - mirror log data each other's NVRAM
  - synchronize each other's time
- They use two Fiber Channel Arbitrated Loops (FC-AL)
  - Each filer uses the A loop to manage its own disks
  - Each filer uses the B loop in takeover mode to manage its partner's disks
- They own their spare disks and do not share them with the other filer.

If one filer in a cluster stops functioning physically, its partner filer:

- creates a virtual filer using the failed filer's identity
- accesses the failed filer's disks and serves its data to clients
- maintains its own identity and continues to serve its own data to clients

When the failed filer is functioning again:

- An operator initiates a giveback that terminates the virtual filer on the partner
- The failed filer resumes normal operation, serving its own data
- The cluster resumes normal operation, with each filer ready to take over for its partner if the partner fails

Chapter contents

This chapter discusses the topics described in the following table.

Read...	To learn about...
"Understanding cluster concepts" on page 357	What clusters are and how they work.
"Configuring interfaces for takeover" on page 361	How to configure interfaces to maintain connectivity if a takeover occurs.
"Understanding takeover" on page 369	How the takeover process works and its result.
"Understanding giveback" on page 372	How the giveback process works.
"Managing a cluster" on page 374	How to manage a cluster in general.
"Managing the cluster in normal mode" on page 375	How to manage a cluster when no takeover has taken place.
"Managing takeover" on page 377	How to manage a cluster after a takeover has taken place.
"Managing filers in takeover mode" on page 379	How to manage the filer that has taken over its partner.
"Managing a virtual filer" on page 381	How to manage the virtual filer that represents the filer that has been taken over.
"Performing a giveback" on page 386	How to restore a cluster to normal operation.

Cluster diagram

The following figure diagrams how the cluster hardware components are connected.

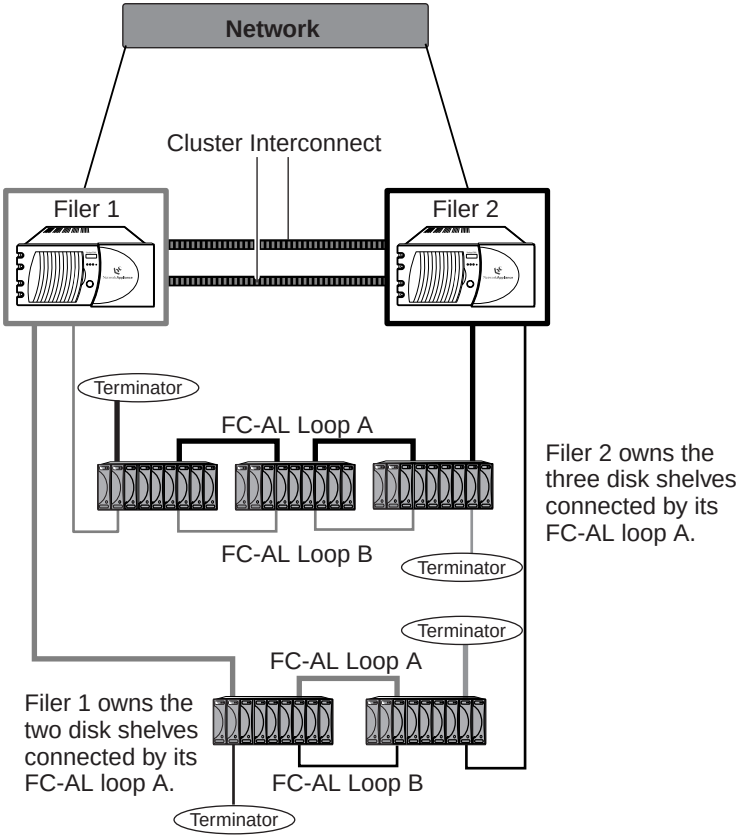

Cluster components and terminology

The following table describes components and explains terms that apply to clusters.

Term	Explanation
cluster	A pair of F630 filers connected so that one filer can detect when the other is not working and, if so, can serve the failed filer's data.
cluster interconnect	Cables and adapters with which the two filers in a cluster are connected and over which heartbeat and WAFL™ log information are transmitted when both filers are running.
cluster monitor	Software that administers the relationship of filers in the cluster. You access the cluster monitor with the `cf` command.
failed filer	The physical filer that has ceased operating. It remains the failed filer until giveback succeeds.
FC-AL disk shelves	In normal operation, each filer in a cluster owns the set of disk shelves that are connected to it through the FC-AL A loop.
giveback	The return of identity from the virtual filer to the failed filer, resulting in a return to normal operation; the reverse of takeover.
heartbeat	A repeating signal transmitted from one filer to the other that indicates that the filer is in operation. Heartbeat information is also stored on disk.
takeover filer	A filer that remains in operation after the other filer stops working and that hosts a virtual filer that manages access to the failed filer's disk shelves and network connections. The takeover filer maintains its own identity and the virtual filer maintains the failed filer's identity.
local filer	The filer you are logged in to.

Term	Explanation
mailbox disk	One of a set of disks owned by each filer that is used to store the cluster state information of the filer. If that filer stops operating, the takeover filer uses the information in the mailbox disks in constructing a virtual filer. Mailbox disks are also used as filesystem disks.
normal mode	The state of a filer when there is no takeover in the cluster.
NVRAM mirror	A synchronously updated copy of the contents of a filer's NVRAM (Non-Volatile Random Access Memory) contents kept on the partner filer.
partner	From the point of view of the local filer, the other filer in the cluster.
partner mode	The method you use to communicate through the command line interface with the virtual filer during a takeover. Partner mode is accessed using the `partner` command.
takeover	The emulation of the failed filer's identity by the takeover filer in a cluster.
takeover mode	The method you use to interact with a filer while it has taken over its partner. The console prompt indicates when the filer is in a takeover state.
virtual filer	A software copy of the failed filer that is hosted by the takeover filer. The virtual filer appears to users and administrators like a functional version of the failed filer. For example, it has the same name as the failed filer.

Normal cluster operation

During normal operation, in addition to serving data, the cluster performs the activities described in the following table.

Action	Description
Heartbeat transmission	The heartbeat signals transmitted over the cluster interconnect cable tell each filer the state of the other filer. This information is written to disks on each filer that the cluster monitor designates as mailbox disks. The disks provide redundancy.
Monitoring	The cluster monitor software looks at the heartbeat signal, commands from filers, and disk connections, and initiates a takeover under the following conditions: ◆ The operator initiates a takeover using the cf takeover command ◆ The partner is determined to have stopped functioning due to a lack of heartbeat activity ◆ The partner's power-on-self-test (POST) fails **Note** If the filer that is not transmitting a heartbeat is rebooting, no takeover occurs. Takeover may be disabled if NVRAM mirroring is not functioning.
NVRAM mirroring	Each filer dedicates half of its NVRAM to a synchronously updated copy of its partner's NVRAM. If a takeover occurs, the takeover filer uses the cluster data in the part of the NVRAM dedicated to the failed filer to ensure that no information is lost.
Time synchronization	Filers' times can deviate from each other. A daemon updates the times on filers periodically to keep them synchronized.

Configuring interfaces for takeover

Interface takeover

When one filer takes over another, it also takes over the interfaces of the failed filer so that it can maintain the same network connections as before the takeover. You can configure interfaces to enable the takeover of interfaces to happen properly. When you configure a new filer, you use the `setup` command to configure interfaces, but if you are adding interfaces, you use the `ifconfig` command, as described in the next section.

Network connections

Each node must be able to assume the identity of its partner on the networks to which the filers are connected.

- Each node in a cluster must have interface attachments to the same collection of subnetworks.
- Each filer must have interfaces that can take over each IP address and MAC (Media Access Control) address on the other filer.

In a noncluster configuration, each interface acts in one role: supporting the identity of the local node. In a cluster configuration, each interface can act in one of two roles, as described in the following table.

- Normal mode means operation when there is no takeover.
- Takeover mode means that the node has taken over the failed filer.

Interface type	Normal mode	Takeover mode
Dedicated	Supports the identity of the local node	Idle.
Hot standby	Idle	Supports the identity of the virtual filer.
Shared	Supports the identity of the local node.	Supports the identity of both the local node and the failed filer.

Takeover configurations

You can have two kinds of interface configurations for takeover.

- A configuration using separate mirrored NICs. You use this configuration with cards that support only single MAC addresses. For each interface, you configure a dedicated interface on one filer and a hot standby interface on the other.
- A configuration using NICs with dual MAC capability. A pair of dual-mac interfaces, one interface to a filer, mirror each other.

FORE/IP and LANE interfaces

You can configure FORE/IP and LANE interfaces as follows:

- You must configure FORE/IP interfaces as dedicated and hot standby interfaces.
- For LANE, you can configure ATM adapter interfaces as mirrored interface pairs, provided that after takeover, a given ELAN is configured only once on an adapter.

You can specify FORE/IP and LANE interfaces using interface names.

Takeover configuration with dedicated and hot standby interfaces

In this configuration, you use a pair of NICs for each interface. Each NIC can use only one MAC address. One NIC acts as a dedicated interface and the other acts as a hot standby.

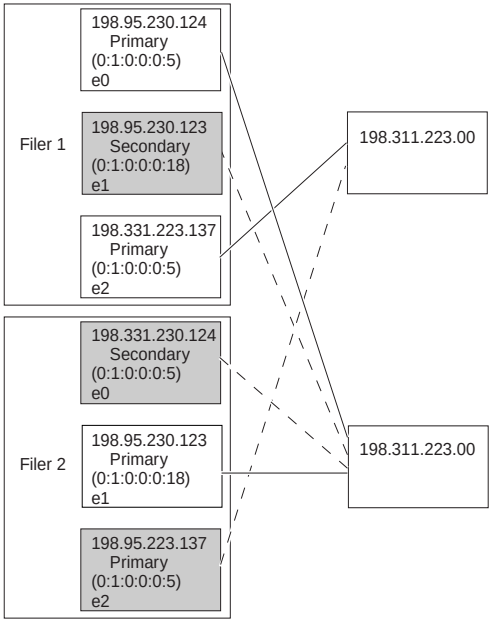

If Filer 1 fails, the following occurs on takeover:

- *e0* stops functioning, but *e0* on Filer 2 substitutes for the failed interface and handles Filer 1's network connection with the 230 network.
- *e2* stops functioning, but *e2* on Filer 2 substitutes for the failed interface and handles Filer 1's network connection with the 223 network.

If Filer 2 fails, *e1* stops functioning, but *e1* on Filer 2 substitutes for the failed interface and handles Filer 2's network connection with the 230 network.

The procedure for setting up a configuration with individual mirrored NICs is in "Configuring dedicated and hot standby interfaces" on page 365.

Takeover configuration with shared interfaces

In this configuration, you use a pair of NICs that support two MAC addresses to mirror two interfaces, as shown below with the interfaces that connects to the 230 net.

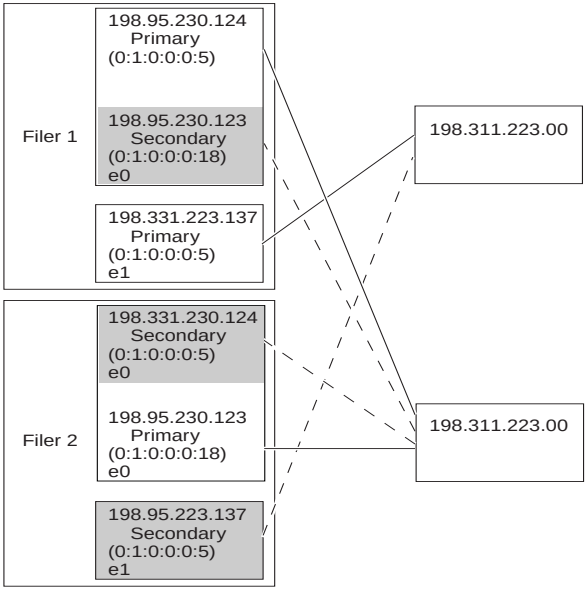

With the configuration shown in the preceding figure, if Filer 1 fails, the following occurs on takeover:

- *e0* stops functioning, but the secondary address on *e0* on Filer 2 handles Filer 1's network connection with the 230 network.
- *e1* stops functioning, but *e1* on Filer 2 substitutes for the failed interface and handles Filer 1's network connection with the 223 network.

If Filer 2 fails, *e1* stops functioning, but *e1* on Filer 2 substitutes for the failed interface and handles Filer 2's network connection with the 223 network.

Configuring dedicated and hot standby interfaces

Description

This procedure describes how to configure a pair of interfaces so that the hot standby interface can replace a dedicated interface on the other filer in the cluster if one filer fails and is taken over. You do this when converting or installing a filer or when you are adding network connections to a cluster. When you complete this procedure, on both filers, you will have a hot standby NIC that can take over the network connection of a dedicated NIC on the other filer if there is a takeover.

Prerequisites

You must have available the following items:

- Four single MAC-address interfaces or ATM adapter interfaces, two on each filer in the cluster
- Two IP addresse or ATM interface names, the addresses of one interface or ATM interface name on each filer in the cluster

Steps

To configure a pair of interfaces for takeover, complete these steps.

Step	Action
1	Enter the following command at the command line or in the */etc/rc* file of the filer with the dedicated local interface: **`ifconfig`** ***`lint laddr`*** *lint* is the i name of the dedicated local interface *laddr* is the IP address or ATM interface name of the dedicated local interface.
2	Enter the following command at the command line or in the */etc/rc* file of the other filer in the cluster: **`ifconfig`** ***`sint`*** **`partner`** ***`saddr`*** *sint* is the name of the hot standby interface *laddr* is the IP address or ATM interface name of the hot standby interface.

Step	Action
3	Enter the following command at the command line or in the */etc/rc* file of the other filer than the filer you used in Step 1: **`ifconfig lint laddr`** *lint* is the interface name of the dedicated local interface *laddr* is the IP address or ATM interface name of the dedicated local interface. This should be the same address or name you used for *saddr* in Step 2.
4	Enter the following command at the command line or in the */etc/rc* file of the filer you used in Step 1 of this procedure: **`ifconfig sint partner saddr`** *sint* is the interface name of the hot standby interface *saddr* is the IP address of the hot standby interface. This should be the same address or name you used for *laddr* in Step 1.

LANE example

Configuring interfaces with mirrored interfaces

Description

This procedure describes how to configure a pair of mirrored, dual-MAC interfaces, one on each filer in the cluster, so that one can replace the other if one filer fails and is taken over. You do this when converting or installing a filer or when you are adding network connections to a cluster. When you complete this procedure, you will have a NIC in each filer that can take over the connection of the other if there is a takeover.

Prerequisites

You must have available the following items:

- Two dual-MAC NICs, one installed on each filer in the cluster
- Two IP addresses, the addresses of one interface on each filer in the cluster.

Caution

For LANE, when an adapter is shared, an ELAN cannot be configured more than once on the adapter after takeover. This means that if both filers in a cluster pair are on the same ELANs, each filer needs a minimum of two ATM cards to support full takeover.

Steps

To configure a pair of dual-MAC NICs for takeover, complete these steps.

Step	Action
1	Enter the following command at the command line or in the */etc/rc* file of one filer in the cluster: **`ifconfig`** ***`lint laddr`*** **`partner`** ***`paddr`*** *lint* is the interface number of the local interface. *laddr* is the IP address or ATM interface name of the local interface. paddr is the IP address ATM interface name that *lint* services when the local filer has taken over the other filer in the cluster.

Step	Action
2	Enter the following command at the command line or in the */etc/rc* file of one filer in the cluster: **`ifconfig lint laddr partner paddr`** *lint* is the interface number of the local interface. *laddr* is the IP address ATM interface name of the local interface. This is usually the *paddr* that you used in the previous step. *paddr* is the IP address ATM interface name that *lint* services when the local filer has taken over the other filer in the cluster. This is usually the *laddr* that you used in the previous step.

ELAN example

The following command specifies that adapter *a4* will takeover all the ELANs on the partner *a5* interface.

ifconfig a4 partner a5

Note

If *a4* and partner *a5* have been configured with common ELANs, then *a4* will not take over those ELANs.

A cluster pair might be configured with ELANs as follows:

Filer and adtapter	Interfaces	Filer and adtapter	Interfaces
filer1-a4	*el0* default *el1* nineKMTU	filer1-a5	*el2* eighteenKMTU
filer2-a	*el0* defaul *el1* nineKMT	filer2-a5	*el2* eighteenKMTU

To achieve full takeover, use the following command on one filer:

`ifconfig a4 partner a5`

and the following command on the other filer:

`ifconfig a5 partner a4`

Understanding takeover

Takeover process

The takeover process is shown in the following table.

Stage	Description
1	One of these three things happen: ◆ One filer in a cluster undergoes a hardware or power failure and can't reboot. That filer is the failed filer, and the surviving filer is the takeover filer. ◆ An operator issues the `cf takeover` command from a filer. The filer from which the command was issued becomes the takeover filer; the other filer becomes the failed filer. ◆ An operator issues the `halt` command without the `-f` option.
2	The takeover filer reserves the disk drives on its FC-AL B loop (which is the failed filer's FC-AL A loop) so that they do not receive commands from the failed filer.
3	The takeover filer creates a virtual filer that takes over the resources and identity of the failed filer and executes the equivalent of a reboot of the virtual filer. For example, it uses network interface information from the failed filer's configuration files to duplicate the failed filer's network services.
4	The virtual filer goes through a process similar to the reboot process. ◆ All filer information preserved across filer reboot is preserved in a takeover. ◆ Filer information that is lost in a reboot, such as CIFS connection state, is lost in a takeover.
5	Users of the failed filer experience the effects of a reboot of the failed filer, and then client service continues.

Messages indicating a successful takeover

During a typical takeover, the cluster monitor issues messages like the following on the console of the takeover filer. In this case, the partner of filer1 was halted.

```
filer1> Tue Jul 21 23:02:02 GMT [filer1: cf_main]: Cluster monitor:
partner halted
Tue Jul 21 23:02:11 GMT [filer1: cf_main]: Cluster monitor: partner
not responding
```

Timeout messages such as the following appear as filer1 waits for signs of life from its partner:

```
Tue Jul 21 23:02:34 GMT GMT [filer1: cf_main]: Cluster monitor:
takeover scheduled in 30 seconds
```

Then the following messages appear:

```
Tue Jul 21 23:03:06 GMT [filer1: cf_takeover]: Cluster monitor:
takeover started
```

Reboot messages follow as the virtual filer boots up, then the following message appears and indicates a successful takeover:

```
Tue Jul 21 23:03:22 GMT [filer1: cf_takeover]: Cluster monitor:
takeover completed
filer1(takeover)>
```

Takeover result diagram

The following figure shows the state of a cluster after a takeover of Filer 1 by Filer 2.

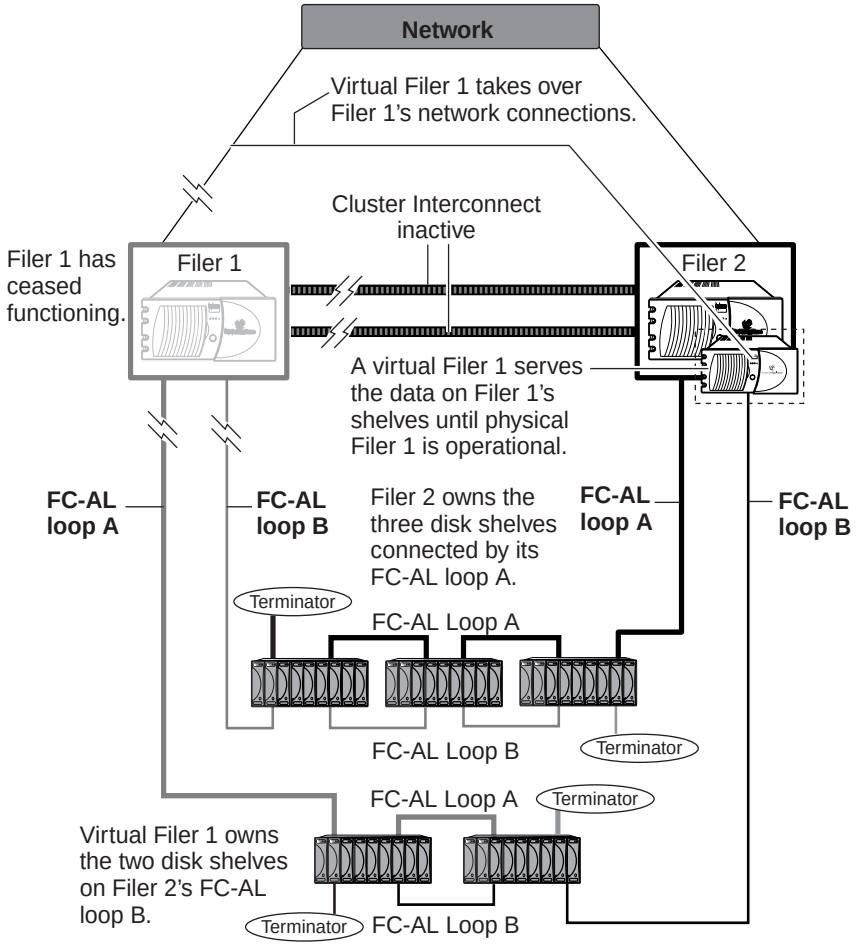

Understanding giveback

Giveback process

The giveback process is described in the following table.

Stage	Description
1	An operator issues a `cf giveback` command.
2	The takeover filer updates the failed filer's disks.
3	The takeover filer destroys the virtual filer so that no more writes can occur to the failed filer's disks.
4	The takeover filer releases the disk reservations that it held on behalf of the failed filer. This returns control to the failed filer of its FC-AL A loop disks.

When the failed filer comes up, it returns to normal cluster operation, and the users of that filer experience the effects of a reboot.

Note

Giveback occurs on the takeover filer and does not depend on the state of the failed filer.

Messages indicating a successful giveback

During a typical giveback, the cluster monitor issues messages like the following on the console of the takeover filer:

```
filer1(takeover)> Tue Jul 21 22:37:34 GMT [filer1: rc]: Cluster
monitor: giveback initiated by operator
Tue Jul 21 22:37:34 GMT [filer1: cf_giveback]: Cluster monitor:
giveback started
Tue Jul 21 22:37:35 GMT [filer1: cf_giveback]: Cluster monitor:
giveback completed
```

The preceding message indicates that the giveback process was successful. When the failed filer reboots, the following messages appear:

```
Tue Jul 21 22:38:35 GMT [filer1: cf_main]: Cluster monitor: partner
ok
```

```
Tue Jul 21 22:38:35 GMT [filer1: cf_main]: Cluster monitor:
takeover of partner enabled
```

The last message indicates that the filers are operating normally and therefore either filer can take over for its partner if necessary.

Note

After a giveback, the takeover filer's ability to automatically take over its partner is not reenabled until the partner reboots successfully. If the partner fails to reboot, you can enter the `cf takeover` command to manually initiate a takeover of the partner.

Forcing a giveback

There are two variations of the cf giveback command that force a giveback to occur:

- `cf giveback -f`
- `cf forcegiveback`

cf giveback -f: The `cf giveback` command will not perform giveback if the following long running or stateful operations are underway on the virtual filer:

- parity check
- disk scrub
- active cifs sessions
- dump or restore
- kernel download
- snapshot creation or deletion
- quota initialization

If the `cf giveback` command fails because any of the above operations are running, you can force a giveback by entering the `cf giveback -f` command. The cf giveback -f command forces the long running and stateful operations to terminate.

cf forcegiveback: In the event that the filer panics during the giveback process, you may use the `cf forcegiveback` command to release the partner's identity and resources.

Caution

Using the cf forcegiveback command may result in the loss of recently written data. The command should be used only under extreme circumstances.

Managing a cluster

What managing a cluster entails

Managing a cluster entails tasks described in the following topics:

- "Managing the cluster in normal mode" on page 375. You can enable or disable takeover, and monitor the cluster in normal mode to verify that it is functioning properly.
- "Managing takeover" on page 377. You can initiate a takeover.
- "Managing filers in takeover mode" on page 379. After a takeover, you can administer the takeover filer, perform maintenance on the failed filer, and initiate a giveback.
- "Managing a virtual filer" on page 381. You manage the virtual filer with the `partner` command.
- "Performing a giveback" on page 386. You can initiate or force a giveback.

Managing the cluster in normal mode

Cluster management overview

You manage a cluster in normal mode by performing the following actions:

- Monitoring the status of the cluster as a whole with the `cf status` and `cf partner` commands, as described in "Monitoring cluster status" and "Displaying the partner's name" on this page.
- Viewing information about the cluster, as described in "Displaying cluster configuration information" on this page.
- Controlling whether time in a cluster is synchronized, as described in "Controlling time synchronization" on page 376.
- Performing license operations on the cluster feature, as described in "Licensing and unlicensing the cluster feature" on page 376.

Monitoring cluster status

To find out whether the cluster feature is enabled and whether the other filer in the cluster is up, enter **`cf status`**.

A display like the following appears:

```
filer1> cf status
Cluster enabled, filer2 is up.
```

This tells you that both filer1 and filer2 are functioning and available for takeover.

If you see a `Cluster disabled` message, it might mean that

- you disabled the cluster feature. To reenable the cluster feature, use the `cf enable` command.
- the system disabled clustering because the interconnect failed.

Displaying the partner's name

To find out the name of the other filer in the cluster, enter **`cf partner`**.

Displaying cluster configuration information

To find out about cluster configuration, use the `sysconfig` or `vol status` command. In a cluster, the `sysconfig` and `vol status` displays are special in that they display information about both filers.

For each filer, the `sysconfig` command display includes disks on both FC-AL loop A and FC-AL loop B.

- The information about disks on FC-AL loop A is the same as for filers not in a cluster.
- The information about disks on FC-AL loop B is for hardware only; the `sysconfig` command displays information about the adapters supporting the disks. The command does not show whether a disk on FC-AL loop B is a file system disk, spare disk, or parity disk.

Controlling time synchronization

You can disable time synchronization on a filer or in a cluster. Normally, both filers should have the same value for this option.

- To disable time synchronization on a filer, use the following command:
  `options cf.timed.enable off`
- To enable time synchronization on a filer, use the following command:
  `options cf.timed.enable on`

Licensing and unlicensing the cluster feature

You can license or unlicense the cluster feature only in normal mode with takeover disabled. Complete these steps to unlicense the cluster feature.

Step	Action
1	If the filer has taken over its partner, enter the following command to perform a giveback: **`cf giveback`**
2	Enter the following command to disable clustering: **`cf disable`**
3	Enter the following command to unlicense clustering: **`license cluster=disable`**
4	Enter the following command to reboot the filer and make the change take effect: **`reboot`**

Managing takeover

Takeover management tasks

You manage takeover by performing either of the following tasks:

- Enabling and disabling takeover for the cluster, as described in the next section, "Enabling and disabling takeover."
- Initiating a takeover of the partner by the local filer, as described in "Initiating a takeover" on page 377.

Enabling and disabling takeover

You might want to disable takeover if you are doing maintenance work that normally causes a takeover, then reenable takeover when you have finished.

- You disable takeover on a cluster with the `cf disable` command.
- You enable takeover on a cluster with the `cf enable` command.

You can issue this command from either filer in the cluster.

Halting a filer without a takeover

To halt a filer and prevent its partner from taking over, enter the following command:

```
halt -f
```

This is useful if you need to perform maintenance on both the filer and its disks.

Initiating a takeover

You might want to initiate a takeover to do maintenance work on a filer in the cluster while still serving the data on the disks of that filer.

You can initiate a takeover in the two ways described in the following table.

Method	Description
`cf takeover`	Initiates a takeover of the partner of the local filer. Takeover is aborted if time-consuming operations, such as a dump, are in progress.
`cf takeover -f`	Initiates a takeover of the partner of the local filer regardless of whether time-consuming operations, such as a dump, are in progress.

Method	Description
`cf forcetakeover`	Tells the cluster monitor to ignore some configuration problems that otherwise prevent a takeover, such as unsynchronized NVRAM due to a faulty interconnect connection. It then initiates a takeover of the partner of the local filer. **Caution** Use `cf forcetakeover` only in certain controlled circumstances, when you cannot get `cf takeover` to succeed. You can lose uncommitted data with this command.

Managing filers in takeover mode

Takeover mode tasks

You can do the following tasks in takeover mode:

- Administer the local filer as you would any normal filer. Take into account the minor differences between a filer in normal mode and a filer in takeover mode described in the next section,"Takeover mode differences from normal mode."
- Enter partner mode and administer the virtual filer, as described in "Managing a virtual filer" on page 381.
- Initiate or force a giveback, as described in "Performing a giveback" on page 386.

Takeover mode differences from normal mode

The following subsections describe differences between managing a filer in normal mode and doing so in takeover mode.

Takeover mode prompt: In takeover mode, the prompt looks like this:
`filer1(takeover)>`

This indicates that filer1 is in takeover mode.

Statistics in takeover mode: Statistics are treated differently in the following commands:

`httpstat`

`ifstat`

`netstat`

`nfsstat`

`sysstat`

In takeover mode, statistics for the preceding commands differ from the statistics in normal mode in the following ways:

- Each display reflects the sum of operations that take place on the takeover filer plus the operations on the virtual filer. Therefore, the display does not differentiate between the operations on the takeover filer and the operations on the virtual filer.
- The statistics displayed by each of these commands are cumulative.

- After giving back the failed partner's resources, the takeover filer does not subtract the statistics it performed on behalf of the failed filer in takeover mode.
- The giveback does not reset (zero out) the statistics.

  If you want accurate statistics from a command after a giveback, reset (zero out) the statistics as described in the man page for the command you are using.

SNMP in takeover mode: You can have different settings on each filer for SNMP options, but any statistics gathered while a filer has been taken over does not distinguish between filers.

Managing a virtual filer

Virtual filer management overview

You manage a virtual filer as you do any filer, including managing disks, with the following exceptions, which are described in greater detail later in this section:

- A virtual filer can access only failed filer disks.
- Some commands are unavailable.
- Some displays differ from normal displays.

Accessing the virtual filer from the takeover filer

You access the virtual filer from the takeover filer in takeover mode with the `partner` command. You can issue the `partner` command in two forms:

- The `partner` command without an argument toggles between partner mode, in which you manage the virtual filer, and takeover mode, in which you manage the takeover filer. The prompt changes to the partner mode prompt, which has the form:

  virtual_filer/*takeover_filer*>,

  The following example shows the change from takeover mode to partner mode and back:

  ```
  filer1(takeover)> partner
  Login from console: filer2
  filer2/filer1> Thu Aug 20 16:44:39 GMT [filer1:
  rc]: Login from console: filer2
  partner
  Logoff from console: filer2
  filer1(takeover)> Thu Aug 20 16:44:54 GMT [filer1:
  rc]: Logoff from console: filer2
  filer1(takeover)>
  ```

  The `partner` command with a filer command as an argument executes the filer command on the virtual filer in partner mode, then returns to takeover mode, as shown in the following example:

  ```
  filer1(takeover)> partner cf status
  filer2 has been taken over by filer1.
  filer1(takeover)>
  Accessing the virtual filer remotely
  ```

You can also access the virtual filer using a remote shell program such as `rsh` without using the `partner` command.

Using Telnet: With Telnet, you can access directly only the takeover filer. To access the virtual filer, log in to the takeover filer, then use the `partner` command.

Using a remote shell: To execute a command on the virtual filer using a remote shell, issue the command to that filer, which has the same name as the failed filer.

Virtual filer tape drives inaccessible

You cannot access tape devices on the virtual filer. You must use the tape drives on the takeover filer instead.

Virtual filer commands

Almost all the commands that are available to a normal filer are available on the virtual filer in partner mode. The commands in the following subsections are exceptions.

Unavailable commands: The following commands are not available on a virtual filer:

`cf disable`

`cf enable`

`cf forcegiveback`

`cf forcetakeover`

`cf giveback`

`cf takeover`

`date` *`date`*

`halt`

`ifconfig partner`

`ifconfig -partner`

`ifconfig mtusize`

`license cluster`

`rdate`

`reboot`

`timezone` *`name`*

Commands with different behaviors: The commands in the following table behave differently in partner mode than they do on a takeover filer.

Command	Difference
`ifconfig [interface]`	Displays the following: ◆ Virtual filer interface mappings based on the failed filer's */etc/rc* file. ◆ Virtual filer's interface names rather than the takeover filer's interface names that have performed the takeover. ◆ Only interfaces that have been configured.
`mt`	Uses the tape devices on the takeover filer. This is because the failed filer has no access to its tape devices.
`netstat -i`	Appends a plus sign (+) to shared interfaces. A shared interface is one that has two IP addresses assigned to it: an IP address for the filer in which it physically resides, and an IP address for it's partner filer in the cluster.
`shelfchk`	Does not distinguish between the takeover filer's disk shelves and the virtual filer's disk shelves. The `shelfchk` command has the same result in partner mode as it does in takeover mode.
`sysconfig`	When it displays hardware information, the `sysconfig` command displays information only about hardware that is attached to the takeover filer. It does not display information about any hardware that is attached only to the failed filer. For example, the disk adapter information that the `partner sysconfig -r` command displays is about the disk adapters on the takeover filer.
`uptime`	Displays how long the failed filer has been down and the host name of the takeover filer.

Command	Difference
`vol status`	When it displays hardware information, the `vol status` command displays information only about hardware that is attached to the takeover filer. It does not display information about any hardware that is attached only to the failed filer. For example, the disk adapter information that the `partner vol status -r` command displays is about the disk adapters on the takeover filer.

Backing up and restoring from the virtual filer

You can perform dumps and restores on the virtual filer. However, neither the takeover filer nor the virtual filer can access the tape drives or other peripheral devices on the failed filer.

Performing dumps on a virtual filer: Any dump commands directed to the failed filer's tape drives are executed on the takeover filer's tape drives. Therefore, any `dump` commands that you execute through the `cron` command succeed only under the following conditions:

- The device names are the same on both filers in the cluster.
- The dump commands for the filer and the virtual filer are not scheduled to occur during the same time period—the filer and the virtual filer cannot access the tape drives simultaneously.

Caution

Make sure there that is a tape on which you want the backup to be written in the tape device for each dump. Any tape in the tape device can get overwritten by a dump.

Performing restores on a virtual filer: To perform restores to the virtual filer, follow these steps:

Step	Action
1	Put the restore tape in a drive accessible by the takeover filer.
2	Use the `restore` command in partner mode or as an argument to the `partner` command.

Verifying the link between clustered filers

You can verify the status of the link between clustered filers by entering the following command:

```
cf status
```

The following table shows the messages that may be displayed by the `cf status` command.

Message	Meaning
`Cluster enabled, partner_name is up.`	The cluster is operating normally.
partner_name_1 `has taken over` *partner_name_2*.	One filer in the cluster has taken over the other filer in the cluster.
`Interconnect not present.`	The system does not recognize the existence of an interconnect device.
`Interconnect is down.`	The interconnect is unable to access the partner. This may be due to cabling problems or the partner may be down.
`Interconnect is up.`	The interconnect is active and is able to transmit data to the partner.
`X link is down; Y link is up.`	The X link (upper cable) is currently down. This condition may occur briefly when a cluster node is booting. If this condition persists, check the cable seating.
`X link is up; Y link is down.`	The Y link (lower cable) is currently down. This condition may occur briefly when a cluster node is booting. If this condition persists, check the cable seating.

Note

Even if one of the interconnect links (X or Y) is down, cluster communication continues over the link that is up.

Performing a giveback

Initiating a giveback

To restore a cluster to normal operation, from the takeover filer, enter one of the commands in the following table.

Command	Description
`cf giveback`	Initiates a giveback to the failed filer. **Note** The cf giveback command will not succeed if any of the following events are taking place: ◆ Dump processes ◆ Outstanding CIFS sessions ◆ Parity checks ◆ RAID reconstructs ◆ RAID scrubs ◆ Volume creation
`cf giveback -f`	Terminates the following events on the virtual filer, then initiates a giveback to the failed filer: ◆ Dump processes ◆ Outstanding CIFS sessions ◆ Parity checks ◆ RAID reconstructs ◆ RAID scrubs ◆ Volume creation

Command	Description
`cf forcegiveback`	Forces the takeover filer to give back the resources of the virtual filer even if the takeover filer detects an error that would normally prevent a complete giveback, such as data not being flushed from NVRAM to the failed filer's disks. When the failed filer reboots as a result of a forced giveback, it displays the following message if data from the virtual filer did not fully synchronize with the virtual filer's disks: `partner giveback incomplete, some data may be lost` **Caution** Use `cf forcegiveback` only in desperate circumstances, when you cannot get `cf giveback` to succeed. You can lose uncommitted data with this command.

What happens when giveback is interrupted

If a failure or power outage happens to the takeover filer during the giveback process, when the failure is repaired or power is restored, the giveback process stops and the takeover filer returns to takeover mode. Some operations, such as quota initialization, will always cause a giveback to fail.

Snapshots 11

Chapter contents

This chapter explains the NetApp snapshot system. The following table lists the topics available.

Read...	To learn about...
"Understanding snapshots" on page 390	Snapshots in general.
"Snapshot commands and options" on page 394	Snapshots commands.
"Understanding snapshot disk consumption" on page 399	How snapshots work with disk consumption.
"Managing snapshot disk consumption" on page 403	How to manage snapshot disk usage.
"Accessing snapshots from clients" on page 408	How to access snapshots from NFS and CIFS clients.

Understanding snapshots

What is a snapshot?

A *snapshot* is a read-only copy of the entire file system—it reflects the state of the file system at the time the snapshot was created.

Accessing snapshots

Any client of a filer can access snapshots to recover old versions of files; for example, files that were accidentally changed or deleted. The snapshot feature enables users to restore their own files without help, because files in snapshots can be viewed and copied by those who have permission to do so with the original files.

Simplifying tape backup

Snapshots also simplify tape backup. The filer `dump` command automatically creates a snapshot of the active file system, if necessary, before backing up the data to tape. However, it is not necessary if you are backing up an existing snapshot. Because a snapshot is a read-only copy of the file system, it does not change even when files in the active file system are changing. As a result, `dump` can make a safe and consistent backup without requiring you to take the filer off-line.

Snapshots use little disk space

The filer uses a copy-on-write technique to create snapshots quickly without consuming any disk space. Snapshots begin to consume extra space only as blocks in the active file system are modified and written to new locations on disk. For more information about the copy-on-write technique used by snapshots, refer to "How snapshots work" on page 391.

Creating snapshots for your needs

The filer creates and deletes snapshots automatically at preset intervals. You can also create and delete snapshots manually. Each volume on the filer can have up to 20 different snapshots at one time.

Snapshots maintain original file permissions

Snapshot files carry the same permissions and inode numbers as the original files, keeping the integrity of the security system intact. Inodes are data structures that hold information about files on the filer. There is an inode for each file and a file is uniquely identified by the file system on which it resides and its inode number on that system.

Note

The inode number for a file in a snapshot is the same as the inode number for the corresponding file in the active file system. As a result, some programs on UNIX clients consider the two files to be the same. For example, if you use the GNU `diff` program to compare the two files, it does not find any differences between them. To make sure that the two files have different inode numbers before the comparison, copy one of the files to another name.

How snapshots work

When the filer creates a snapshot, it doesn't copy disk blocks; instead, it identifies all blocks in the file system as belonging to the snapshot as well as to the active file system.

Example: Consider a particular file named *foo* in a newly created snapshot. Because the snapshot was just created, the snapshot version of *foo* has the same contents as the version in the active file system. The same blocks on disk store both versions, so the snapshot version of *foo* consumes no disk space.

Later, if *foo* is deleted, the blocks holding the data for *foo* are no longer part of the active file system, but they are still part of the snapshot. Therefore, deleting *foo* from the active file system does not free any disk space.

Diagram of a snapshot

The following figure illustrates how disk space is used before and after *foo* is removed.

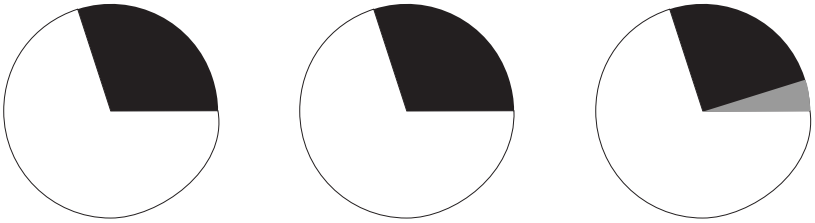

Before any snapshot is taken, disk space is consumed by the active file system only.

After a snapshot is taken, the active file system and snapshot point to the same disk blocks. The snapshot doesn´t use extra disk space.

After the *foo* file is deleted, the space previously used by *foo* is referenced by the snapshot. That´s why deleting active file system data doesn´t free disk space.

Space used by the active file system
Space used by the snapshot only
Unused disk space

Changing the contents of *foo* creates a similar situation. New data written to *foo* cannot be stored in the same disk blocks as the current contents because the snapshot is using those blocks to store the old version of *foo*. Instead, the new data is written to new disk blocks, so there are two separate copies of *foo* on disk—a new copy in the active file system and an old one in the snapshot. This technique of duplicating disk blocks only as they are modified is called *copy-on-write*.

In some directories, most data remains unchanged from day to day. For example, a user with a 10-MB home directory might change only 100 KB to 500 KB on a typical day. When files change slowly, snapshots can be kept on-line for days or even weeks before they begin to consume unacceptable amounts of disk space. In other directories, data changes quickly. If a large percentage of data changes every day, there might not be room to keep snapshots for even a few hours. To accommodate the needs of different users, create multiple volumes on the filer. In this way, you can apply different snapshot schedules to different volumes.

In summary, when the filer creates a snapshot, it doesn’t use any disk space, but as files in the active file system are changed or deleted, the snapshot uses more and more disk space. How often files are changed and deleted determines the number of snapshots the filer can create and the length of time the snapshots can be kept.

Snapshot commands and options

Snapshot commands

The commands related to snapshots are listed in the following table. If the volume name is omitted in any of these commands, the command applies to the root volume.

Command	Meaning
`snap list`[a] *volume_name*	Lists all available snapshots.
`snap create` *volume_name snapshot_name*	Creates a snapshot with a specified name.
`snap delete` *volume_name snapshot_name*	Deletes a specified snapshot.
`snap rename` *volume_name from to*	Renames a snapshot.
`snap reserve` *volume_name*	Reserves a percentage of the disk space for snapshots.
`snap sched` *volume_name*	Schedules automatic snapshots.

a. Some of the information generated by this command is available through SNMP using the Network Appliance custom MIB. The custom MIB is described in Chapter 4, "Using SNMP," on page 111."

Snapshot options

The following options for the `vol options` command affect snapshots in the specified volume. The options remain in effect after the filer reboots.

Options	Descriptions
`nosnap`	Disables automatic snapshots. By default, this option is disabled.
`nosnapdir`	Makes the *.snapshot* directory that is present at client mount points or the root of the CIFS share invisible. It also turns off access to the *.snapshot* directory and all *.snapshot* directories under the mount point or the root of the CIFS share. By default, this option is disabled.

Note

The `dump` command does not work if the `nosnapdir` or `nosnap` option is on.

Automatic snapshot creation

The filer uses the `snap sched` command to create snapshots automatically and to keep them on-line for a predetermined amount of time.

Types of automatic snapshots: The following table describes the three types of automatic snapshots.

Type	Description
Weekly	The filer creates these every Sunday at midnight. Weekly snapshots are called *weekly.n*, where *n* is an integer. *weekly.0* is the most recent weekly snapshot, and *weekly.1* is the next most recent weekly snapshot. When the filer creates a weekly snapshot, the value of *n* is adjusted for all weekly snapshots. The higher the value of *n*, the older the snapshot.
Nightly	The filer creates these every midnight except when a weekly snapshot is scheduled to occur at the same time. If the number of weekly snapshots is nonzero and it's the day of the week that weekly snapshots occur, no nightly snapshot is created. Nightly snapshots are called *nightly.n*, where *n* is an integer. *nightly.0* is the most recent nightly snapshot, and *nightly.1* is the next most recent nightly snapshot. When the filer creates a nightly snapshot, the value of *n* is adjusted for all nightly snapshots. The higher the value of *n*, the older the snapshot.

Type	Description
Hourly	The filer creates these on the hour at specified hours, except at midnight, if a nightly or weekly snapshot is scheduled to occur at the same time. This occurs either if the number of nightly snapshots in the schedule is nonzero, or if the number of weekly snapshots in the schedule is nonzero and it's the day of the week that weekly snapshots occur. Hourly snapshots are called *hourly.n*, where *n* is an integer. *hourly.0* is the most recent hourly snapshot, and *hourly.1* is the next most recent hourly snapshot. When the filer creates an hourly snapshot, the value of *n* is adjusted for all hourly snapshots. The higher the value of *n*, the older the snapshot.

Example 1 of snap sched command

The following example shows a sample `snap sched` command:

```
snap sched vol1 2  6  8@8,12,16,20
```

Volume name

Create a snapshot every week and keep a maximum of two.

Create a snapshot every day and keep a maximum of six.

Create a snapshot every hour, or at the times listed in the optional time list and keep a maximum of eight.

Optional list of times, in 24-hour format at which an hourly snapshot is created.

In the `snap sched` command, the first argument after the volume name in the example indicates how many weekly snapshots to keep (2), the second argument indicates how many nightly snapshots to keep (6), and the third argument indicates how many hourly snapshots to keep (8). A zero in any of the three positions disables snapshots for that interval.

The argument for hourly snapshots can include an optional list of numbers indicating the hours at which the filer creates the snapshots in 24-hour time (8, 12, 17, 20). If the argument is omitted, the filer creates an hourly snapshot.

The default snapshot schedule is

`snap sched` *`volume_name`* `0 2 6@8,12,16,20`

Example 2 of snap sched command

Following is an example of the `snap sched` command:

`snap sched` *`volume_name`* `2 6 8@8,12,16,20`

Snapshots created by this schedule: The following list describes the snapshots created by the example:

- weekly snapshots, and keeps the two most recent
- daily snapshots, and keeps the six most recent
- hourly snapshots at 8:00 a.m, noon, 4:00 p.m., and 8:00 p.m., and keeps the eight most recent

The following list shows the snapshots that are created by this snapshot schedule in 1998 (when January 11 is a Sunday):

```
% ls -lu .snapshot

total 64
drwxrwsrwx  2 root 4096 Jan 14 12:00 hourly.0
drwxrwsrwx  2 root 4096 Jan 14 08:00 hourly.1
drwxrwsrwx  2 root 4096 Jan 13 20:00 hourly.2
drwxrwsrwx  2 root 4096 Jan 13 16:00 hourly.3
drwxrwsrwx  2 root 4096 Jan 13 12:00 hourly.4
drwxrwsrwx  2 root 4096 Jan 13 08:00 hourly.5
drwxrwsrwx  2 root 4096 Jan 12 20:00 hourly.6
drwxrwsrwx  2 root 4096 Jan 12 16:00 hourly.7
drwxrwsrwx  2 root 4096 Jan 14 00:00 nightly.0
drwxrwsrwx  2 root 4096 Jan 13 00:00 nightly.1
drwxrwsrwx  2 root 4096 Jan 12 00:00 nightly.2
drwxrwsrwx  2 root 4096 Jan 10 00:00 nightly.3
drwxrwsrwx  2 root 4096 Jan 09 00:00 nightly.4
drwxrwsrwx  2 root 4096 Jan 08 00:00 nightly.5
drwxrwsrwx  2 root 4096 Jan 11 00:00 weekly.0
drwxrwsrwx  2 root 4096 Jan 04 00:00 weekly.1
```

Result: This schedule keeps the eight most recent hourly snapshots, created at 8 a.m., noon, 4 p.m., and 8 p.m. every day, the six most recent daily snapshots, and the two most recent weekly snapshots. Whenever the filer creates a new snapshot

of a particular type, it deletes the oldest one and renames the existing ones. On the hour, for example, the filer deletes *hourly.7*, renames *hourly.0* to *hourly.1*, and so on. The nightly snapshot schedule jumps from January 12 to January 10 because there is a weekly snapshot on January 11.

Note

On a UNIX client, if you use `ls -l` instead of `ls -lu` to list the snapshot creation times, the times are not necessarily all different. The times listed by `ls -l` reflect the modification times of the directory at the time of each snapshot, and are not related to the times at which the snapshots are created.

The `snap sched` command is persistent across reboots. There is no need to put the command in the */etc/rc* file.

User-defined automatic snapshots

You can create snapshots at predefined times instead of using the hourly, daily, and weekly schedules.

Example: For example, if you want to create two snapshots for the volume named vol1 each week, you can set up a `cron` job on the administration host to run twice each week at an appropriate time to execute the following `snap create` command:

```
rsh toaster snap create vol1 filename
```

Note

The `snap create` command does not accept a snapshot name containing a slash (/).

Understanding snapshot disk consumption

About snapshot disk consumption

It is important to understand the amount of disk space snapshots consume and the amount of disk space they are likely to consume. The following sections explain how to determine the amount of disk space used by snapshots.

Disk consumption by multiple identical snapshots

Suppose a snapshot contains a 1-MB file that hasn't changed since the filer created the snapshot. If that file is removed from the active file system, the snapshot then consumes 1 MB of disk space.

The same version of that 1-MB file might be referenced by several snapshots: *hourly.0*, *hourly.1*, and *hourly.2*. If these snapshots all contain the 1-MB file that hasn't changed since the filer created those snapshots, only 1 MB of disk space is consumed by the snapshots even though all three snapshots contain the file.

Using the df command to display snapshot use

To provide information about snapshot disk utilization, the `df` command on the filer treats snapshots as a partition different from the active file system.

Sample df command output: Following is a partial sample `df` command output:

`df`

```
Filesystem            kbytes    used    avail   capacity
/vol/vol0             3000000  2000000  1000000  65%
/vol/vol0/.snapshot   1000000  500000   500000   50%
```

Note

The numbers in this example were rounded off to make the example easier to understand. Also, to make the output easier to read, the "Mounted on" column is not included in the sample `df` output in the following sections.

In this example, the *vol0* volume contains 4 GB of disk space. It has 1 GB (or 25%) reserved for snapshots (the idea of reserving space for snapshots is described in more detail later). That leaves 3 GB for the active file system, and 2 GB of the file system is in use.

It is important to understand that the `/vol/vol0/.snapshot` line counts data that exists only in a snapshot. Because data that also exists in the active file system needs to be stored on disk anyway, it is misleading if the filer charged the space to snapshots. In the example, half of the 1 GB reserved for the snapshot is used.

How the snapshot reserve works

By default, the snapshot reserve is 20% of disk space. For information about how to adjust the amount of the snapshot reserve, refer to "Changing the snapshot reserve" on page 406.

This section uses several examples to explain the advantages of reserving disk space for snapshots.

Snapshots use deleted active file disk space: If the filer created a snapshot when the disks were full, removing files from the active file system wouldn't create any free space because everything in the active file system would also be referenced by the newly created snapshot. The filer would have to delete the snapshot before it could create any new files.

The following example shows how disk space being freed by deleting files in the active file system ends up in the snapshot:

If the filer creates a snapshot when the active file system is full and there is still space remaining in the snapshot reserve, the `df` command output is as follows:

```
Filesystem              kbytes    used   avail   capacity
/vol/vol0/              3000000 3000000       0     100%
/vol/vol0/.snapshot     1000000  500000  500000      50%
```

If you remove 100 MB of files, the disk space used by these files is no longer part of the active file system, so the space is reassigned to the snapshots instead. If you enter the `df` command, the output is as follows:

```
Filesystem              kbytes    used   avail   capacity
/vol/vol0/              3000000 2900000  100000      97%
/vol/vol0/.snapshot     1000000  600000  400000      60%
```

The filer reassigns 100 MB of space from the active file system to the snapshot reserve. Because there was reserve space for snapshots, removing files from the active file system freed space for new files.

Administering snapshot disk space: Even with the snapshot reserve, the job of administering snapshot disk space consumption is important. There is no way to prevent snapshots from consuming disk space greater than the amount

reserved for them. Consider what would happen in the example if all files in the active file system were deleted. Before the deletion, the `df` output was as follows:

```
Filesystem              kbytes    used   avail   capacity
/vol/vol0/              3000000 3000000       0     100%
/vol/vol0/.snapshot     1000000  500000  500000      50%
```

After removing all the data in the file system, the `df` command generates the following output:

```
Filesystem              kbytes    used   avail   capacity
/vol/vol0/              3000000 2500000  500000      83%
/vol/vol0/.snapshot     1000000 3500000       0     350%
```

Explanation: The entire 3 GB in the active file system moved into snapshots, along with the 500 MB that were in snapshots before, making a total of 3.5 GB of snapshot data. This is 2.5 GB more than the space reserved for snapshots. Because the active file system cannot use that space, the space shows up as used by the active file system even though no files are there.

Recovering disk space for file system use: Whenever snapshots consume more than 100% of the snapshot reserve, the system is in danger of becoming full. In this case, you can create files only after you remove enough snapshots.

Example: For example, if 500 MB of data were added to the active file system in the preceding example, a `df` command generates the following information:

```
Filesystem              kbytes    used   avail capacity
/vol/vol0               3000000 3000000      0     100%
/vol/vol0/.snapshot     1000000 3500000      0     350%
```

As soon as the filer creates a new snapshot, every block in the system is referenced by some snapshot. Therefore, no matter how many files you remove from the active file system, there is still not room to add any more. The only way to recover from this situation is to remove enough snapshots to free more disk space. Refer to "Displaying snapshot statistics" on page 404 for information about how to use the `snap list` command to determine which snapshot to delete.

Effects of snapshots on quotas

Quotas do not count disk space consumed by snapshots. If snapshots were included in the quota calculations, users could end up in a state where they could not create any new files until all snapshots containing their old files expired.

Managing snapshot disk consumption

About snapshot management

This section describes

- how to schedule snapshots to suit your environment
- how to determine a reasonable snapshot reserve
- how to adjust the amount of space snapshots use

The examples in this section are for a volume named *home*.

Scheduling snapshots

The best way to manage the amount of space consumed by snapshots in each volume is to use the `snap sched` command to adjust the schedule of snapshot creation.

Following are some suggested strategies for scheduling and retaining snapshots:

- If users rarely lose files or typically notice lost files right away, use the default snapshot schedule. For example, this is the schedule that creates a snapshot every day and keeps two:

```
snap sched home 0 2 6@8,12,16,20
```

- If users commonly lose files and need to restore them, Network Appliance recommends that you delete the snapshots less often than you would in the preceding example.

  On many systems only 5% or 10% of the data changes each week, so the snapshot schedule of six nightly and two weekly snapshots consumes 10% to 20% of disk space. Considering the benefits of snapshots, it is worthwhile to reserve this amount of disk space for snapshots. Following is the recommended snapshot schedule, which keeps six daily snapshots and two weekly snapshots:

```
snap sched home 2 6 8@8,12,16,20
```

- If the data changes very quickly, reduce the number of snapshots scheduled. For example, if a volume is filled and emptied each day, for example, a volume storing large temporary files for a CAD application, it might not make sense to use daily or weekly snapshots at all.

  On a very active volume, schedule snapshots every hour and keep them for just a few hours, or turn off snapshots. For example, the following schedule creates a snapshot every hour and keeps three:

```
snap sched home 0 0 3
```

This schedule doesn't consume much disk space, and it lets users recover files in recent snapshots as long as they notice their mistake within a couple of hours.

- When you create a new volume on a filer, the new volume inherits the snapshot schedule from the root volume. After you use the volume for a while, check how much disk space the snapshots consume in the volume.

Displaying snapshot statistics

The `snap list` command shows the amount of disk space used by snapshots in a specified volume. This command enables you to see how much disk space each snapshot uses, and helps you determine an appropriate snapshot reserve.

Command output

Following is an example of the command output. If you don't specify a volume name in the command, the output contains statistics about each volume.

```
snap list vol0

Volume vol0
%/used       %/total    date          name
----------  ----------  ------------  --------
0% ( 0%)    0% ( 0%)    Jan 19 08:01  hourly.0
1% ( 1%)    1% ( 1%)    Jan 19 00:01  nightly.0
2% ( 2%)    2% ( 2%)    Jan 18 20:01  hourly.1
3% ( 2%)    2% ( 2%)    Jan 18 16:01  hourly.2
3% ( 2%)    3% ( 2%)    Jan 18 12:01  hourly.3
5% ( 3%)    4% ( 3%)    Jan 18 00:01  nightly.1
7% ( 4%)    6% ( 4%)    Jan 17 00:00  nightly.2
8% ( 4%)    7% ( 4%)    Jan 16 00:01  nightly.3
10%( 5%)    9% ( 4%)    Jan 15 00:01  nightly.4
```

The %/used column : The %/used column shows space consumed by snapshots as a percentage of disk space being used in the volume. The first number is cumulative for all snapshots listed so far, and the second number is for the specified snapshot alone.

- The first number is equal to

$$\frac{\text{cumulative snapshot space x 100\%}}{\text{cumulative snapshot space + file system space}}$$

- The second number is equal to

$$\frac{\text{this snapshot x 100\%}}{\text{this snapshot + file system space}}$$

The %/total column : The %/total column shows space consumed as a percentage of total disk space in the volume.

- The first number is equal to

$$\frac{\text{cumulative snapshot space x 100\%}}{\text{total disk space in this volume}}$$

- The second number is equal to

$$\frac{\text{this snapshot x 100\%}}{\text{total disk space in this volume}}$$

  "Cumulative snapshot space" is the total space used by this snapshot and all other more recent snapshots (the ones preceding this snapshot in the `snap list` output).

Output summary: The %/used number is more useful for planning the snapshot reserve because it is more likely to remain constant as the file system fills.

The example shows a volume that keeps five nightly snapshots and four hourly snapshots. That is, the volume uses the following command for creating snapshots regularly:

```
snap sched vol0 0 5 4@8,12,16,20
```

The `snap list` output shows that the overhead for snapshots is only 10%, so the default snapshot reserve of 20% seems to be a waste of disk space. Assuming that this pattern of change holds up, a reserve of 12% to 15% provides a safe margin to ensure that removing files frees disk space when the active file system is full.

The values in parentheses, that is, the values that show the space used by an individual snapshot, are useful in identifying a particular snapshot to remove when the file system is full. However, deleting a particular snapshot doesn't necessarily release the total amount of disk space indicated, because other snapshots might be referring to the same blocks. Refer to "Adjusting disk space used by snapshots" on page 406 for further information about how to select a snapshot file for deletion to reclaim disk space.

If you do not want the total amount of disk space consumed by all snapshots to exceed a certain percentage of the used disk space, use the cumulative values in the `snap list` output to determine which snapshots to delete. In the preceding example, if you don't want more than 5% of used disk space to be spent by snapshots, delete all snapshots listed below *nightly.1* in the `snap list` output; that is, *nightly.2*, *nightly.3*, and *nightly.4*. After deleting the snapshots, *nightly.1* and all the other more recent snapshots consume 5% of the used disk space.

Changing the snapshot reserve

The snapshot reserve can be used only by snapshots, not by the active file system.

The default snapshot reserve is 20% of the available disk space. To change the reserve, enter the following command:

`snap reserve volume_name percent`

For example:

`snap reserve vol0 25`

With no arguments, the `snap reserve` command displays the percentage of disk space reserved for snapshots in each volume.

Note

Snapshots can exceed the snapshot reserve space.

Adjusting disk space used by snapshots

This section describes how to use the `snap list` output to determine which snapshot file to delete to free the most disk space.

In the sample `snap list` output in "Displaying snapshot statistics" on page 404, the cumulative disk space used by snapshots gradually increases from top to bottom.

For example, in the %/used column, the cumulative space used by *hourly.1* is 2% and the cumulative space used by *hourly.2* is 3%. This is not always the case.

Example: Consider a filer with a 100-MB file system that has not changed since the first snapshot was taken. The `snap list` command on this filer displays the following output:

```
%/used       %/total      date          name
--------     ---------    ------------  ---------
0% (0%)      0% (0%)      May 05 16:00  hourly.0
0% (0%)      0% (0%)      May 05 12:00  hourly.1
0% (0%)      0% (0%)      May 05 08:00  hourly.2
```

The cumulative disk space used by snapshots does not increase because no changes were made to the file system. However, if you had deleted 20 MB from the file system before the filer took the *hourly.0* snapshot, the `snap list` command would have displayed the following output:

```
%/used       %/total      date          name
--------     ---------    ------------  ---------
0%  (0%)      0% (0%)     May 05 16:00  hourly.0
20% (20%)     1% (1%)     May 05 12:00  hourly.1
20% (20%)     1% (1%)     May 05 08:00  hourly.2
```

In the %/used column, the cumulative values for *hourly.1* and *hourly.2* are both 20%, but the cumulative value for *hourly.2* is not 40%. This is because both snapshots point to the same 20 MB of data, the data that you just deleted.

The cumulative values for *hourly.1* and *hourly.2* are different if you delete and create data between snapshots in the following way:

1. Delete 20 MB of data and create 20 MB of new data after *hourly.2*.
2. Delete the 20 MB of data created in Step 1 after *hourly.1*.

After the data deletions and additions, the `snap list` command displays the following output:

```
%/used       %/total      date          name
--------     ---------    ------------  ---------
0%  (0%)      0% (0%)     May 05 16:00  hourly.0
20% (20%)     1% (1%)     May 05 12:00  hourly.1
33% (20%)     2% (1%)     May 05 08:00  hourly.2
```

In this scenario, *hourly.1* and *hourly.2* each consume 20% of the used disk space: 20 MB out of 100 MB. However, this time they reference different data blocks. Cumulatively, they consume 40 MB, which is about 33% of the disk space used: 120 MB, which is 40 MB used by snapshots plus 80 MB in the file system.

Before trying to conserve space by deleting a large snapshot file, examine the cumulative values in the `snap list` output. If two adjacent snapshot files show little difference in the cumulative values, most of the data referenced by the snapshots is the same. In this case, removing one of the snapshots doesn't free much disk space.

If you find snapshots confusing and hard to manage, use the default snapshot schedule and the default snapshot reserve because these settings are appropriate for most environments. When you create a new volume on a filer, remember that the new volume inherits the snapshot schedule from the root volume. After you use the volume for a while, check how much disk space the snapshots consume in the volume. If the disk space seems high, decrease the amount of time that snapshots are kept or increase the snapshot reserve.

As you use snapshots, continue to watch the statistics change over time. The statistics help you gain a better understanding of how snapshots work.

Accessing snapshots from clients

About client access to snapshots

Snapshots can be accessed by any user with the appropriate permissions. Every directory in the filer's active file system contains a directory named *.snapshot*, through which users can access old versions of files in that directory. How users gain access to snapshots depends on the file-sharing protocol used: NFS or CIFS.

NFS client access to snapshots

The following illustration shows the directory structure on a client with the *vol0* volume of a filer named toaster mounted on */n/toaster*.

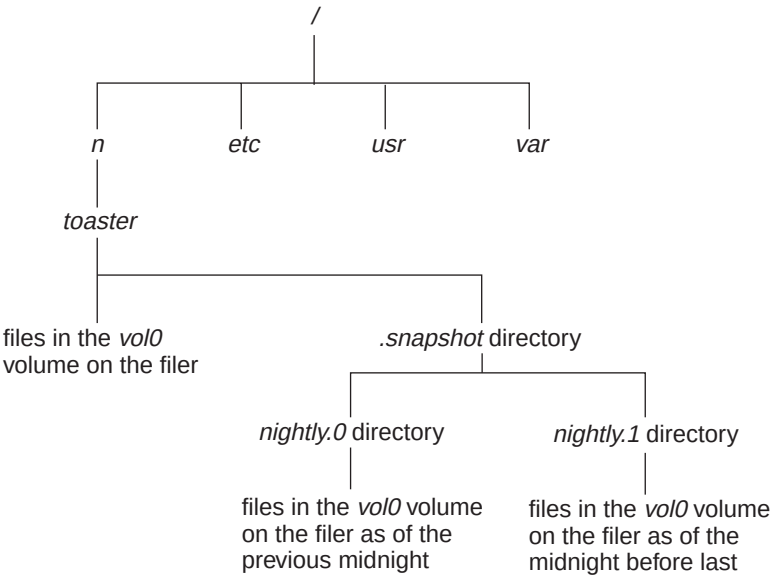

Explanation: In this example, the client can obtain access to snapshots by way of */n/toaster/.snapshot*. This might seem to contradict the explanation of snapshot access in the preceding section, because it shows a *.snapshot* directory only at the mount point instead of in every directory in the tree.

Actually, the *.snapshot* directory in the mount point is "real" to make the `pwd` command work, whereas the *.snapshot* directories in all other directories are "magic"; that is, can be accessed when they are referenced by name but do not show up in a directory listing.

For example, at the mount point of a filer file system, a directory listing looks like this:

```
ls -a

.     ..     .snapshot     dir1     dir2
```

The same command entered in a directory below the mount point does not show the *.snapshot* directory; for example:

```
cd dir1
ls -a

.     ..    file1     file2
```

If you enter the `ls` command with the directory name *.snapshot*, you can see a listing of the snapshots for the *dir1* directory:

```
ls .snapshot

hourly.0     hourly.4     nightly.0     nightly.4
hourly.1     hourly.5     nightly.1     nightly.5
hourly.2     hourly.6     nightly.2     weekly.0
hourly.3     hourly.7     nightly.3     weekly.1
```

If *.snapshot* were to show up in every directory, it would cause many commands to work improperly. For instance, all recursive commands for removing files would fail because everything below *.snapshot* is read-only. Recursive copies would copy everything in the snapshots as well as files in the active file system, and a `find` command would generate a list much longer than expected.

CIFS client access to snapshots

To CIFS clients, the snapshot directory appears only at the root of a share. For example, if a user's home directory is a share named bill that corresponds to the */vol/vol0/home/bill* directory, only the */vol/vol0/home/bill/.snapshot* directory is visible. When this user displays the contents of the home directory, the snapshot directory is displayed as *~snapshot* if the operating system supports long file names and as *~snapsht* if the operating system supports only short file names.

Note

The snapshot directory is visible in that it is displayed in a directory listing or File Manager display if the client operating system is configured to show hidden files.

In each directory within the share, a snapshot directory exists but is not visible to clients. For example, if the client operating system supports long file names, the applications on that operating system can use the snapshot at each level of the share by using *.snapshot*, *~snapshot*, or *~SNAPSHT* as the directory name. You cannot, however, display the directory name in any listing.

Determining snapshot versions

From an NFS client: the best way to find all versions of a particular file preserved in snapshots is to use the `ls` command. The following example shows how to find all versions of *foo*:

```
ls -l foo .snapshot/*/foo
-rw-r--r--  1 smith 0 Jan 14 09:40  foo
-rw-r--r--  1 smith 0 Jan 13 18:39  .snapshot/nightly.0/foo
-rw-r--r--  1 smith 0 Jan 12 19:17  .snapshot/nightly.1/foo
```

The version of *foo* in the active file system was last modified on January 14, but the old versions available in the snapshots were modified on January 13 and January 12. Although users can use standard UNIX commands to examine the saved versions of *foo*, they cannot modify or remove these older versions because everything beneath *.snapshot* is read-only.

From a CIFS client: use the `Find` command to search for the file in the *~snapshot* directory. For example, if a user maps the home share to drive F: and wants to find all versions of *foo* in snapshots, the user can use the `Find` command to search for *foo* in the *f:\~snapshot* folder.

Determining access times

When the filer creates a snapshot, the access time of each file in the snapshot is updated to the snapshot creation time.

From an NFS client: you can use the `ls -lu` command, which shows the access times of files, to determine when snapshots were created. Following is an example of the `ls -lu` command:

```
ls -lu foo .snapshot/*/foo
-rw-r--r--  1 smith 0 Jan 14 09:40  foo
-rw-r--r--  1 smith 0 Jan 14 00:00  .snapshot/nightly.0/foo
-rw-r--r--  1 smith 0 Jan 13 00:00  .snapshot/nightly.1/foo
```

From a CIFS client: you can determine the access time of a file by checking its properties.

Qtree administration 12

About qtrees

A qtree is a special subdirectory of the root directory of a volume.

Qtree parameters: You can set the following parameters on a qtree:

- security style: NTFS (Windows NT file system), UNIX, or mixed
- oplocks setting: On or Off
- disk space and file limits, as described in Chapter 13, "Quotas and Maximum Number of Files."

Note

Qtrees used to be called quota trees. Any quota trees created in previous releases of Data ONTAP are now qtrees.

Volumes and qtrees: A volume has all the properties of a qtree except

- It can have qtrees under it.
- It has different defaults than a qtree.

Unless expressly mentioned otherwise, whatever applies to qtrees also applies to volumes.

Note

You cannot create a qtree inside another qtree.

Uses of qtrees: You use qtrees to group files that have similar characteristics, much in the way that you use volumes. However, qtrees are much more flexible than volumes. For details, see "Using qtrees" on page 413.

Chapter contents

This chapter discusses how to administer qtrees from the filer command line. You can also administer qtrees from FilerView.

This chapter discusses the topics listed in the following table.

Read...	To learn about...
"Using qtrees" on page 413	What qtrees are and what you can do with them.
"Qtree security styles" on page 415	Security styles and how to set and change the security style of a qtree.

Read...	To learn about...
“Qtree file access models” on page 418	How to look at qtrees using file access models.
“Creating a qtree” on page 420	How to create qtrees.
“Modifying the security style of a qtree” on page 421	How to set and change the security style of a qtree.
“Modifying qtree oplocks settings” on page 422	How to set and change the oplocks setting of a qtree.
“Displaying qtree information” on page 424	How to display information about all qtrees on a filer so you can see how many there are and what their properties are. This enables you to make informed decisions about qtree management.

Using qtrees

What you can do with qtrees

You can use qtrees in the following two ways:

- Group files that have the same security style and oplocks setting, such as files related to a particular activity, for example, a project, without having to create a separate volume for them.
- Perform quick and easy backups.

Using a qtree for a project: One way to group files is to set up a qtree for a project, such as one maintaining a database. Setting up a qtree for a project enables you to do the following actions:

- Set the security style of the project without affecting the security style of other projects.

  For example, you use NTFS-style security if the members of the project use Windows files and applications. Another project in another qtree can use UNIX files and applications, while yet another project can use both Windows and UNIX files.
- Set oplocks (if the project uses Windows) as appropriate to the project without affecting other projects.

  For example, if one project uses a database that requires no oplocks, you can turn oplocks Off on that project's qtree. If another project uses oplocks, it can be in another qtree that has oplocks set to On.
- Limit the disk space and number of files available to a project so that it does not use up resources that other projects and users need. As the needs of the projects and available resources change, you can easily change the limits on the qtree. For instructions about managing space using qtrees, see Chapter 13, "Quotas and Maximum Number of Files."

Using a qtree for backups: You can back up individual qtrees. You would do so to

- add flexibility to backup schedules
- modularize backups
- keep the size of each backup to one tape

For details, see Chapter 14, "Data Backup."

Qtree and volume defaults

Volumes and qtrees have the default values shown in the following table.

Parameter	Qtree default	Volume default
`oplocks`	On	On
`security`	The style of the volume's root directory	UNIX

Note

Quota trees that were created in previous releases of Data ONTAP have, as defaults, UNIX security style and oplocks On.

Moving files between qtrees

In UNIX, you cannot move a file into or out of a qtree with a rename operation. As a result, the `mv` command on some UNIX systems fails if you try to move a file into or out of a qtree. You can always move the file by copying it, then deleting the original.

In Windows, you can move a file into or out of a qtree.

Qtree security styles

Types of security styles

There are three kinds of security styles, described briefly in the following table.

Style	Behavior
NTFS	Exactly like Windows NT NTFS: Files and directories have Windows NT file-level permission settings. **Note** To use NTFS security, make sure that the filer is licensed for CIFS.
UNIX	Exactly like UNIX: Files and directories have UNIX permissions.
mixed	Both NTFS and UNIX security are allowed: a file or directory can have either Windows NT permissions or UNIX permissions.

Qtree security styles in detail

The following table describes the security styles in detail and the effects of changing to each style.

Security style	Description	Effect of changing to the style
NTFS	User access is determined as follows: **CIFS requests:** Windows NT permissions determine user access if Windows NT permissions have been set on a file. **NFS requests:** Windows NT permissions and a mapped CIFS identity determine access. UNIX groups are not used in the mapping from a UNIX identity to a CIFS identity\. **Note** You cannot change permissions in an NTFS qtree from a UNIX client.	Windows NT permissions determine file access for a file that had them if the change is from a mixed qtree. Otherwise, UNIX-style permission bits determine file access for files created before the change.
UNIX	User access depends on the protocol, as follows: **CIFS requests:** Windows users are mapped to a UNIX UID and UNIX permissions determine access. In a UNIX qtree, a user cannot set Windows NT permissions. A Windows user can change UNIX permissions from Windows using SecureShare Access, as described in "Sending a message to all users on a filer" on page 311. **NFS requests:** Only the user's UID, GID, and UNIX-style permission bits of the file or directory determine user access.	The filer disregards any Windows NT permissions established previously.

Security style	Description	Effect of changing to the style
mixed	Both NTFS and UNIX style permissions are permitted. The security style of a file is the style most recently used to set permissions on that file. See the information in "NTFS" on page 415. **Caution** Changing NTFS permissions on a file recomputes UNIX permissions on that file. Changing UNIX permissions or ownership on a file deletes any NTFS permissions on that file.	None.

Qtree file access models

Kinds of file access models

You can use the following four file access models in working with qtrees:

- CIFS user accessing a file with Windows NT security
- CIFS user accessing a file with UNIX security
- NFS user accessing a file with Windows NT security
- NFS user accessing a file with UNIX security

CIFS access to Windows files

CIFS accesses to Windows files obey Windows security rules.

CIFS access to UNIX files

The following principles apply to accessing UNIX files from CIFS:

- All CIFS users are mapped to UNIX UIDs and GIDs.
- File accesses use UNIX security or PC security, as chosen during the `cifs setup` program.
  - UNIX-style permissions are determined by the rights associated with the UNIX UID and GID.
  - PC-style permissions are determined by the rights assigned in a share's Access Control List (ACL) and are limited by the UNIX permissions assigned to a file.
- PC security is like FAT (File Allocation Table) file system security with per-file permissions:
  - If the owner of a file or directory accesses an item, the owner permissions are checked to see whether they allow access.
  - If someone other than the owner of a file or directory accesses an item, the group permissions are checked to see whether they allow access.

NFS access to Windows files

The following principles apply to accessing Windows files from NFS:

- Windows NT permissions are mapped to UNIX permissions.
- Each Windows NT user who sets Windows NT permissions is mapped to a UNIX user and UNIX group, except that if the owner is a generic user, the owner is mapped to root with restrictions.

- Windows NT permissions for Owner are mapped to UNIX owner permissions.
- Windows NT permissions for Everyone are mapped to UNIX Group and UNIX Other permissions.

NFS access to UNIX files

NFS accesses to UNIX files obey UNIX security rules.

Creating a qtree

How to create a qtree

To create a qtree, use the following command:

```
qtree create pathname
```

Result: The qtree *pathname* is created, with the following properties:

- Volume: the root volume, unless you specify another volume
- Name: *pathname*
- Security style: that of the root directory of the volume
- Oplocks settings: On

Creating a qtree in the root volume

If *pathname* does not begin with a slash (/), the qtree is created in the root volume. For example:

```
qtree create news
```

creates the qtree */vol/vol0/news*, where */vol/vol0/* is the default name for the root volume. For information about volumes, see.“Volume concepts” on page 80

Creating a qtree in a volume other than the root volume

If you want to create a qtree (for example, *news*) in a particular volume (for example, *users*), use the following command:

```
qtree create /vol/users/news
```

Modifying the security style of a qtree

When to change the security style of a qtree

There are many circumstance under which you might want to change qtree style. Two examples are

- Because the default security style of a qtree is that of its root volume, you might want to change the security style of a qtree after creating it to the style you want.
- You might also want to change the security style to accommodate other users or files; for example, if you start with an NTFS qtree and subsequently want to include UNIX files and users, you might want to change the qtree to a mixed qtree.

How to change the security style of a qtree

To change the security style of a qtree, use the following command:

```
qtree security [pathname [mixed | ntfs | unix]]
```

Example with a qtree: To change the security model of */vol/users/docs* to be the same as Windows NT, use

```
qtree security /vol/users/docs ntfs
```

Example with a volume: To change the security model of the root directory of the *users* volume to mixed so that, outside of a qtree in the volume, one file can have NTFS security and another UNIX security, use

```
qtree security /vol/users/ mixed
```

Note

When you create an NTFS qtree or change a qtree to NTFS, by default, every Windows user is given full access. You must change the permissions if you want to restrict access to the qtree for some users. If you do not set NTFS file security on a file, UNIX permissions are enforced.

Modifying qtree oplocks settings

When to change oplocks settings

You might want to change qtree oplocks settings when you add or remove software, or when the kind of data you are using changes. For detailed information about oplocks, see "Using oplocks" on page 328.

Changing oplocks settings

To change the oplocks setting of a qtree, follow these steps:

1. Make sure that the `cifs.oplocks.enable` option is set the way you want.
2. Use the qtree oplocks command, as follows:

```
qtree oplocks [name [enable | disable]]
```

The command takes effect immediately.

Note

If you disable the oplocks feature on a qtree, existing oplocks in the qtree are not broken.

Example with a qtree: To enable oplocks in the */vol/users/docs* qtree if oplocks are disabled and the `cifs.oplocks.enable` option is set to On, enter the following command:

```
qtree oplocks /vol/users/docs enable
```

Example with a volume: To disable oplocks in the entire *users* volume if oplocks are enabled and the `cifs.oplocks.enable` option is set to On, enter the following command:

```
qtree oplocks /vol/users/ disable
```

This disables only files and directories that were not in a qtree when oplocks were enabled.

Effect of the cifs.oplocks.enable option

Setting the `cifs.oplocks.enable` option has the following effects:

- If the `cifs.oplocks.enable` option is set to Off, all oplocks on the filer are turned off.

- If the `cifs.oplocks.enable` option is set back to On, the setting for each qtree comes into effect and oplocks are turned on for those qtrees where oplocks are enabled.

Displaying qtree information

How to display qtree information

To display all attributes of all qtrees on a filer, including quota trees created in previous versions of Data ONTAP by the `quota qtree` command, use the `qtree` command with no arguments.

The qtree command display

The `qtree` command lists for a filer the items described in the following table.

Field	Contents
`Volume`	The volumes on a filer. Bear in mind that a volume is itself a qtree.
`Qtree`	Qtrees that are not volumes; each is listed next to its volume.
`Style`	The security style of each qtree.
`Oplocks`	The oplocks setting of each qtree.

Example qtree display: For example:

`qtree`

```
Volume   Qtree    Style Oplocks
-------- -------- ----- --------
bagels            unix  enabled
bagels   sesame   unix  enabled
muffin            ntfs  enabled
muffin   bran     unix  enabled
```

Explanation of example qtree display: In the example

- Because *bagels* and *muffin* are volumes, each has a security style and oplocks setting.
- *Sesame* is a qtree in the *Bagels* volume. Its security style and oplocks setting happen to be the same as that of its parent volume, bagels.
- *Bran* is a qtree in the *muffin* volume. Its security style is different from that of its parent volume, *muffin*. Files in *bran* have UNIX-style permissions; files in *muffin* but not in *bran* have NTFS-style permissions.

Quotas and Maximum Number of Files 13

Chapter contents

This chapter describes how to manage disk space on the filer. The following table lists the types of information available.

Read...	To learn about...
"Restricting or tracking disk usage by using disk quotas" on page 426	Restricting disk usage by using quotas.
"Increasing the maximum number of files" on page 437	Increasing the maximum number of files in a volume.
"The df command" on page 438	Using the `df` command to monitor the current disk space usage.

Restricting or tracking disk usage by using disk quotas

About disk quotas

Filer disk quotas restrict disk space and the number of files used by a user, a group, or a qtree. For information about how to create a qtree, refer to Chapter 12, "Qtree administration." This chapter discusses how to manage disk quotas by editing the */etc/quotas* file.

Format of the quotas file

To set up disk quotas, create a *quotas* file in the */etc* directory.

Note

Keep a record of your *quotas* file in a safe place and update it as you change it, in case you must do a restore without having access to the root volume.

Following is a sample *quotas* file:

```
#Quota Target       type                 disk     files
/vol/home/user/joe  user                 500M      10K
21                  group                750M      75K
/vol/eng/proj1      tree                 750M      75K
writers             group@/vol/eng/proj1 300M      50K
*                   user                 50M       10K
```

Note

If the quota is a tree quota, the field in the `type` column of the *quotas* file displays `tree`, not `qtree`.

Keep a record of your *quotas* file in a safe place and update it as you change it, in case you must do a restore without having access to the root volume.

Quota Target field: Specifies the user, group, or qtree on which you want to impose restrictions. You can assign more than one quota to a user or group, but only one quota to a qtree. The entries can be in any order.

Quota target for a user quota: You specify a user with one of the following targets:

- a file or subdirectory whose UID matches the user
- the user's name, as defined in the */etc/passwd* file or the NIS password map
- the user's UID

The methods are equivalent, and inform the filer of the UID of the target. A file or directory is used only as the source of a UID; there are no quota implications for that file or directory. The UID of the user must not be 0.

Any file or subdirectory you use in the Quota Target field is referenced repeatedly throughout the life of the system, so if you use a path name, choose a path name that will last for as long as the user account remains on the system. For example, use a user's home directory for a user quota.

Quota target for a group quota: You specify a group with one of the following targets:

- a file or subdirectory whose GID matches the group
- the group's name
- the group's GID

The methods are equivalent. A file or directory is used only as the source of a GID; there are no quota implications for that file or directory. The GID of the group must not be 0.

Quota target for a tree quota: To create a tree quota, use the `quota qtree` command to create a directory in the root directory of a volume. The quota target in the quotas file for a tree quota is the complete path name of this directory.

Quota target for default quotas: Use an asterisk (*) in the Quota Target field to specify a default for the user or group quotas. Defaults do not apply to tree quotas. The default value applies to the following users or groups:

- New users or groups that are created after the default entry has taken effect. For example, if the maximum disk space for default user quotas is 500 MB, any new user can use up to 500 MB of disk space.
- Users or groups that are not explicitly mentioned in the *quotas* file. For example, if the maximum disk space for default user quotas is 500 MB, users for whom you have not specified a user quota in the *quotas* file can use up to 500 MB of disk space.

To override a default for a specific user or group, specify a quota for that user or group.

Type field: You can enter one of the following values in the Type field to define the quota type:

- **user:** If a user quota applies just to a tree and not to the entire volume, specify `user@`*`tree`*, where *tree* is the name of a qtree. If a user quota applies to a volume other than the root volume, append `@/vol/`*`volume`* to

the quota type. For example, `user@/vol/marketing` means that the user quota applies to the *marketing* volume.

- **group:** If a group quota applies just to a tree and not to the entire volume, specify `group@tree` where *tree* is the name of a qtree. If a group quota applies to a volume other than the root volume, append `@volume` to the quota type. For example, `group@/vol/marketing` means that the group quota applies to the *marketing* volume.
- **tree:** A tree quota is similar to a disk partition, but you can increase or decrease the size of a tree quota at any time.

Disk field: Specifies the maximum amount of disk space that the quota target can use. In this field, K is equivalent to 1,024 bytes, M means 2^20 bytes, and G means 2^30 bytes. If you omit the K, M, or G, the default is K.

If you want to track the disk usage but do not want to impose a disk usage limit on the quota target, enter a dash (-) in the disk field.

Note

Do not put a blank in the Disk field; it acts as white space. The filer regards the following entries as equivalent:

```
#Quota Target      type      disk      files
/export            tree                75K
/export            tree      75K
```

Files field: Specifies the maximum number of files that the quota target can use. Use K to indicate 1,024 files. For example, 75K means 76,800 files. Use M to mean 2^20 and G to mean 2^30. You can omit the K, M, or G. For example, if you enter 100, it means the maximum number of files is 100. A blank in this field means there is no restriction on the number of files that the quota target can use.

If you want to track the number of files but do not want to impose a limit on the number of files that can be used by the quota target, enter a dash (-) in the files field.

Sample quotas file

Following is a sample quotas file that includes different kinds of quotas:

```
#Quota Target         type                  disk      files
/vol/home/user/jdoe   user                  500M      10K
108                   user                  500M      10K
jsmith                user@/vol/rls         500M      10K
publications          group                 750M
/vol/home/eng         group@/vol/cad        750M      75K
```

```
/vol/cad/proj1        tree                  750M      75K
writers               group@/vol/cad/proj1  150M
*                     user                   50M      15K
*                     user@/vol/cad/proj1    50M      10K
*                     group                 750M      85K
*                     group@/vol/cad/proj1  100M      75K
mhoward               user                  150M     100K
mhoward               user@/vol/cad/proj1    75M      75K
mfisher               user                    -        -
```

Any operation that creates files or writes to them must satisfy all applicable quotas. The following list describes the effects of the sample quotas file:

- The owner of */vol/home/user/jdoe* and the user whose UID is 108 can each use 500 MB of disk space and 10,240 files in the root volume.
- The user whose user name is jsmith can use 500 MB of disk space and 10,240 files in the *rls* volume.
- The group publications can use 750 MB of disk space with no restrictions on the number of files in the root volume.
- The group that owns */vol/home/eng* can use 750 MB of disk space and 76,800 files in the *cad* volume.
- The qtree proj1 in the *cad* volume can use 750 MB of disk space and 76,800 files.
- The writers group can use 150 MB of disk space and an unlimited number of files in the proj1 qtree provided that the quotas on the proj1 qtree are not exceeded.
- Any user not otherwise mentioned in this file can use 50 MB of disk space and 15,360 files in the root volume.
- In the proj1 qtree, any user not otherwise mentioned in this file can use 50 MB of disk space and 10,240 files.
- Any group not otherwise mentioned in this file can use 750 MB of disk space and 87,040 files in the root volume.
- In the proj1 qtree, any group not otherwise mentioned in this file can use 100 MB of disk space and 76,800 files.
- The user mhoward can use 150 MB of disk space and 102,400 files in the root volume. In the proj1 qtree, mhoward can use 75 MB of disk space and 76,800 files.
- There is no limit as to how much disk space or how many files the user mfisher can use. You can, however, use the `quota report` command to display the amount of disk space and the number of files used by this user.

Both user and group quotas apply to the entire specified volume (or the root volume if no volume name is specified). This is true even if the quota target in the *quotas* file is specified in the form of a path name. For example, if the quota target for the user named jdoe is */vol/home/user/jdoe*, the filer imposes the quotas on all files written by jdoe, not just the ones written to */vol/home/user/jdoe*.

The quota command

The quotas file specifies what the restrictions are on users, groups, and trees. Whether these restrictions take effect depends on the `quota` command. The `quota` command enables you to do the following tasks:

- enable and disable quotas on a per-volume basis
- resize quotas on a per-volume basis
- display information about all active quotas or about quotas that apply to a specified path

Some of the information from `quota` commands is available through SNMP using the Network Appliance custom MIB. For more information about the MIB, refer to "About the Network Appliance custom MIB" on page 112.

Enabling or disabling quotas: The `quota on|off` command enables or disables quotas for all volumes or a specific volume. Use the following syntax when using the command:

```
quotas [on|off] volume
```

Use the following command to enable quotas for a volume:

```
quota on volume
```

This command computes the disk usage of each quota target. The computation can take a few minutes to complete for a large number of quotas. To find out how much quota initialization the filer has completed, use the `quota` command without any arguments. For example:

```
quota

vol0: quotas are on.
vol1: quotas are initializing (24% done).
vol3: quotas are off.
```

Use the following command to disable quotas for a volume:

```
quota off volume
```

Because the filer remembers whether quotas are on or off even after it reboots, there is no need to add a `quota` command to */etc/rc*.

Resizing quotas: The `quota resize` command updates active quotas without recalculating disk usage, and is faster than `quota off` followed by `quota on`. An active quota is one that appears in the output of the `quota report` command, discussed in "Displaying information about quotas" on page 433.

You use the `quota resize` command after changing limits for a group or user. For example, jdoe has a disk quota of 500 MB, as shown in the following example.

```
#Quota Target         type         disk      files
/vol/home/user/jdoe   user         500M       10K
*                     user          10M       10K
```

Use the `quota resize` command to update the quota if jdoe's disk quota were increased, as shown in the following example:

```
#Quota Target         type         disk      files
/vol/home/user/jdoe   user         600M       10K
*                     user          10M       10K
```

How quota resize affects newly added quota targets: The `quota resize` command usually ignores newly added quota targets. For each entry that it skips, the `quota resize` command prints the following message:

`quota resize: new entry on line` *n* `in /etc/quotas`

where *n* is a line number.

The `quota resize` command does not ignore newly added quota targets in the following situation:

If a default quota applies to the creator of a file, an active quota record is created for the owner of the file. If a newly added quota target is the user or group that has an active quota record, the `quota resize` command does not ignore the newly added quota target. That is, if a user or group has written to a file that is under the control of a default quota, a newly created entry in the quotas file for this user or group takes effect after a `quota resize` command.

After you edit the quotas file, if you want to make sure that all entries take effect, enter the `quota off` command followed by the `quota on` command. In this way, all quotas listed in the file become active.

Creating an active quota for a quota target: You can use the `quota resize` command so that quotas can take effect on targets that have not created any files.

Because a quota becomes active when a quota target has written a file, you can make an entry in the quotas file an active quota by following these steps:

1. Create a file.
2. Change ownership of the file to the quota target.

Because the quota target now has an active quota record, you can enter the `quota resize` command, including the volume name, to make the quota entry go into effect. This procedure takes less time than executing the `quota off` command followed by the `quota on` command because only the newly active quota is recalculated.

Displaying information about quotas: The `quota report` command displays the current consumption of files and space for each quota target and for each user and group that is under quota restrictions.

For example, if a user quota is specified in the quotas file for a user named jdoe, the quota report shows how many files and how much disk space have been used by jdoe. The Quota Specifier column in the quota report shows the same information as the Quota Target column in the quotas file, with the following exception:

If a user or group is under default quota restrictions, `quota report` displays information about the user or group as if the user or group had an entry in the quotas file. In this case, the Quota Specifier column in the quota report is blank.

With a path argument, `quota report` displays information about all quotas that apply to the specified file.

Creating a user quota

Follow these steps to limit disk space used by a user named jdoe.

Step	Action
1	Decide who should be limited by a disk quota. In this example, the user is the owner of */vol/home/user/jdoe*. The filer restricts disk usage of files with the same UID as */vol/home/user/jdoe*.
2	Add the following line to the quotas file: **`/vol/home/user/jdoe  user  300M  20K`** Substitute the values you want for 300M and 20K. You can also use the UID of jdoe or the name jdoe if you have set up a *passwd* file on the filer or the NIS database. If you want the restrictions to apply to a volume other than the root volume, for example, the home volume, enter the following line: **`/vol/home/user/jdoe  user@/vol/home  300M  20K`** Substitute the values you want for 300M and 20K.

Step	Action
3	Use the `quota on` command, including the volume name. For example, if quotas are already on for the home volume, enter the following commands: **`quota off home`** **`quota on home`** Alternatively, if jdoe is already under quota restriction, for example, if his files were restricted by a default user quota, enter the following command so that the user quota you just created can take effect: **`quota resize home`**

Creating a group quota

Follow these steps to limit disk space used by a group named service.

Step	Action
1	Decide which group should be limited by a disk quota. In this example, the group is the owner of */vol/home/user/service*. That is, the filer restricts disk usage of files with the same GID as */vol/home/group/service*.
2	Add the following line to the quotas file: **`/vol/home/group/service  group  700M  100K`** Substitute the values you want for 700M and 100K. You can also use the GID of the group or the name of the group. If you want the restrictions to apply to a volume other than the root volume, for example, the home volume, enter the following line: **`/vol/home/group/service  group@/vol/home  700M  100K`** Substitute the values you want for 700M and 100K.

Step	Action
3	Use the `quota on` command. If quotas are already on for the *home* volume, enter the following commands: **`quota off home`** **`quota on home`** Alternatively, if the service group is already under quota restriction, for example, if the group's files are already restricted by a default group quota, enter the following command so that the group quota you just created can take effect: **`quota resize home`**

Removing quota restrictions

To remove quota restrictions, follow these steps.

Step	Action
1	Remove the appropriate line or lines from the *quotas* file.
2	Use the `quota off` and `quota on` commands, including the volume name. For example, enter the following commands: **`quota off home`** **`quota on home`**

When quotas are exceeded

This section describes how the filer responds when quotas are exceeded and what users see on the client systems.

Messages displayed by the filer when quotas are exceeded: When it receives a write request, the filer first determines whether the file to be written is in a qtree. If the write would exceed the qtree quota, the filer logs the following error message:

```
tid tree_id: tree quota exceeded on volume_name
```

If the qtree is not full but the write would cause either the user or group quota to be exceeded, the filer logs one of the following errors:

```
uid user_id: disk quota exceeded
gid group_id: disk quota exceeded
```

where *user_id* is the user's UID and *group_id* is the group's GID.

To the client, the filer returns an "out of disk space" error to the NFS write request or a "disk full" error to the CIFS write request. The following sections describe how the clients notify the users about quotas being exceeded.

Messages displayed on NFS clients: If a write from an NFS client to a filer causes a quota to be exceeded, the user experience depends on the operating system version and the application.

If a UNIX client mounts a filer without the `noquota` option, the `login` program on the client checks to see whether the user has reached the disk quota and file quota each time the user tries to log in to the client. The client displays a message to alert the user before displaying the system prompt if a quota was reached. In the following example, a user reached the disk quota on the filer mounted as */t/toaster* on a client named client2:

```
rlogin client2
You have mail.
Block limit reached on /t/toaster
client2%
```

Not all versions of UNIX perform the quota check as described in this section. Also, the exact message printed varies from one version to another.

If a write causes a quota to be exceeded, the error message seen by the user depends on the application. For example, on a SunOS 4.x client, if a user tries to save a file using `vi` when his or her disk quota is reached, the error message is

```
Disc quota exceeded [Warning - /t/toaster/home/jdoe/file1 is
incomplete]
```

Messages displayed on CIFS clients: If a write from a CIFS client to a filer causes a quota to be exceeded, the user experience depends on the operating system and the application. Following are two examples:

- An application might display a message as follows:

```
Cannot write file filename
```

- When a user tries to copy a file to the filer using the Explorer in Windows 95, the error is as follows:

```
Cannot create or replace filename: Cannot read from the source file
or disk.
```

Increasing the maximum number of files

About increasing the maximum number of files

Initially, the maximum number of files on the filer is set at one for every 32 KB of disk space. The number is increased automatically when you add a new disk. The increase is determined by the filer, and is not a user-specified value.

Unlike UNIX, which requires that you specify the maximum number of files in a file system when you create the file system, the filer enables you to use the `maxfiles` command to increase the number of files for each volume at any time.

Note

Use caution when increasing the maximum number of files because after you increase this number, you can never reduce it. As new files are created, the file system consumes the additional disk space required to hold the inodes for the additional files; there is no way for the filer to release that disk space. An inode is a data structure containing information about files.

Viewing the number of files in a volume

To see how many files are in a volume, use the `df -i` *`pathname`* command, which shows how many inodes have been used, or use the `maxfiles` *`volume`* command.

For example, both of the following commands show that the home volume has used 2,872 inodes:

```
df -i /vol/home

Filesystem     iused     ifree    %iused   Mounted on
/vol/home/      2872      118090      2%    /vol/home

maxfiles home

Volume home: maximum number of files is currently 120962 (2872
used)
```

Information generated by the `maxfiles` command is available through SNMP using the Network Appliance custom MIB. For more information about the MIB, see “About the Network Appliance custom MIB” on page 112.

The df command

About the df command

To verify the amount of free disk space on the filer, enter the `df` command on the filer. Information generated from the `df` command is also available through SNMP, using the Network Appliance custom MIB, which is described in "About the Network Appliance custom MIB" on page 112.

With the `-i` option, the command displays the number of used inodes and the number of available inodes. Following is an example of the `df -i` command:

```
df -i /vol/home

Filesystem    iused      ifree    %iused   Mounted on
/vol/home/    240843     121525     66%     /vol/home
```

The total amount of disk space shown in the `df` output is less than the sum of available space on all disks installed in the volume. Consider the following sample `df` command issued on a volume with seven 2-GB disks installed:

```
df /vol/vol0

Filesystem          kbytes    used       avail   capacity
/vol/vol0/          9886448   8426196    1460252 85%
/vol/vol0/.snapshot 1098492   427008     671484  39%
```

When you add the numbers in the kbytes column, notice that the sum, 10,984,940 KB (about 11 GB), is significantly less than the total disk space installed (14 GB). The following reasons account for the discrepancy:

- The parity disk, which is a 2-GB disk in this example, is not reflected in the output of the `df` command.
- As with the UNIX FFS (Fast File System), the filer reserves 10 percent of the total disk space for efficiency, which `df` does not count as part of the file system space.

The filer doesn't limit the size of a file system. The maximum size of the file system is restricted only by the amount of disk space installed in the volume. For individual files, the maximum size is 8 TB. Many NFS clients restrict the file size to 2 GB or 4 GB.

Using the df command with qtrees

When you enter a `df` command with a path name on a client, the command returns the amount of free space in the file system containing the path name. For example, if the filer is mounted on the client as */t/toaster*, the `df` command on the client displays the disk information about the */t/toaster* file system as follows:

```
df /t/toaster/engineering/jdoe

Filesystem            kbytes    used   avail capacity  Mounted on
toaster:/            2097151 1646923  450228    79%    /t/toaster
```

However, if you defined qtrees on the filer, the information about available space could be misleading because the actual space available might be less. For example, if */engineering* is a qtree with a disk quota of 1,800 MB, the space available in the */t/toaster/engineering* directory is less than that in the `df` command output shown in the preceding example.

If you have qtrees on the filer, Network Appliance recommends that you mount each qtree separately. For example, if the filer named toaster has two qtrees, */vol/home/engineering* and */vol/home/marketing*, mount *toaster:/vol/home/engineering* and *toaster:/vol/home/marketing* on two mount points, for example, */t/toaster/engineering* and */t/toaster/marketing*.

In this way, the filer takes the qtrees into consideration when responding to a `df` command from a client and returns the amount of free space in each qtree, as opposed to the space available in the entire file system.

Data Backup 14

Introduction to data backup

Meaning of data backup

Data backup means copying data from disk to tape. While the filer dump command enables you to copy data to standard output, this chapter mainly discusses how to copy data to tape.

Why you want to back up data from disk to tape

The following list describes the reasons for backing up data from disk to tape:

- You can restore data from tape if an application or a user inadvertently corrupts or deletes files that cannot be recovered from snapshots.
- You can store the backup tapes at an off-site archive to protect the data against natural disasters.
- After you reinstall the file system on the filer (for example, for migrating to larger disks or for converting a single-volume filer to a multivolume filer), you can restore data from tape.

Different methods for backing up the filer

The filer supports three backup methods:

- Using the filer `dump` command.
- Using an application on a computer that supports the Network Data Management Protocol (NDMP). For more information about NDMP, refer to Chapter 17, "Network Data Management Protocol (NDMP)."
- Using commands or a non-NDMP-compliant application from an NFS or a CIFS client.

Caution

If you use the `tar` or `cpio` command on an NFS client to back up the filer, be aware that some versions of these commands fail with file systems that contain long path names, unusual file names, or hard links.

Comparison of backup methods

The following table compares the different backup methods.

Filer dump command	NDMP applications	Commands or non-NDMP applications on NFS or CIFS client
Initiate the backup session from the filer console or from a trusted host through `rsh`.	Initiate the backup session from the computer running the NDMP application.	Initiate the backup session from the NFS or CIFS client.
Back up any files or directories.	Back up any files or directories.	Back up the files or directories that are mounted or shared by the client.
Back up data to tape drives attached to the filer being backed up, to another filer, or to another computer.	Back up data to tape drives attached to the filer being backed up, to the computer running the NDMP application, or to another filer.	Back up data to tape drives attached to the filer or to the client if the client supports the `rmt` protocol.
Preserve both NFS and CIFS attributes.	Preserve both NFS and CIFS attributes.	If backup is initiated on an NFS client, CIFS attributes are not preserved.

Scope of this chapter

This chapter describes how to use the `dump` command to back up the filer. This chapter does not describe how to use commands or applications on an NFS or a CIFS client to back up the filer.

This chapter also does not discuss how to enable NDMP on the filer so that you can use an NDMP-compliant application for backing up the filer.

How the dump command works

Purpose of the dump command

The `dump` command writes file system data from disk to tape in a format that enables you to restore the data to a filer using the filer's `restore` command or the Solaris `ufsrestore` command.

What the dump command can back up

The `dump` command can back up a file, a directory, a qtree, or an entire volume. In the `dump` command, you specify the complete path name to be backed up. In this chapter, this path name is referred to as the "dump path."

How the dump command uses snapshots to back up data

The dump path can exist in the filer's active file system or in a snapshot. If the dump path is in an active file system, the filer takes a snapshot of the active file system before it writes the data to tape. The snapshot capability ensures that the data written to tape is consistent. As a result, you need not take the filer or volume off-line before initiating the backup.

The `dump` command names each snapshot it creates *snapshot_for_backup.n*. The *n* at the end of the snapshot name is an integer starting at 0. Each time the `dump` command creates a snapshot, it increments the integer by 1. The filer resets the integer to 0 when it is rebooted.

The `dump` command automatically deletes the snapshot after it successfully finishes the backup.

When the filer executes multiple `dump` commands simultaneously, the `dump` commands create multiple snapshots. For example, if the filer is running two `dump` commands simultaneously, you find these snapshots in the volumes from which data is being backed up: *snapshot_for_backup.0* and *snapshot_for_backup.1*.

Caution

When the `dump` command is in progress, the filer does not allow you to delete the *snapshot_for_backup.n* file. If you are backing up data from an hourly, daily, or weekly snapshot, make sure that the snapshot scheduler does not delete the snapshot before the `dump` command is finished.

Metadata being backed up

In addition to backing up data within files, the `dump` command backs up these types of metadata:

- UNIX group ID, owner ID, and file permissions
- UNIX access time and modify time
- File type, including UNIX symbolic links and hard links
- File size
- DOS name, attributes, and create time
- Windows NT ACLs

Note

The CIFS attributes (DOS name, attributes, and create time, and Windows NT ACLs) can be restored only with the filer's `restore` command. You cannot use the Solaris `ufsrestore` command to restore these attributes, although the `ufsrestore` command can restore the data in CIFS-created files.

How to exclude certain types of data from the backup

Options in the `dump` command enable you to exclude certain types of data from the backup.

Windows NT ACLs: You can choose not to back up Windows NT ACLs if the data is used only by NFS clients.

Exclude list: You can choose to exclude files and directories from a backup if you do not need those files and directories again. For example, you can exclude temporary files generated by some applications or object files produced during program compilation.

When the `dump` command traverses the dump path, it compares file and directory names to each exclude string specified in the `dump` command. If the name exactly matches the string, the file is excluded from the backup.

Note

You can reduce the amount of backup data by using an exclude list in the `dump` command. However, an exclude list increases the amount of time needed to finish a `dump` command because the filer must compare each file name to the exclude list to determine whether the file should be backed up.

Devices used by the dump command

The `dump` command can back up data to these devices:

- Tape drives or tape stackers attached to the filer

- Tape drives or tape stackers attached to another computer, provided that the following requirements are met:
  - The computer supports the `rmt` protocol.
  - The filer has a trusting relationship with the remote computer to which the tape drive is attached. This relationship enables the filer to write to the tape drive. For example, if you want to back up the filer to a tape drive attached to another filer, include the filer in the */etc/hosts.equiv* file of the destination filer. If you want to back up the filer to a tape drive attached to a SunOS or Solaris computer, include the filer in the */.rhosts* file on the computer.
  - The filer can resolve the name of the computer to which the tape drive is attached using the information about the computer in the filer's */etc/hosts* file or in the DNS database.
- Standard output, provided that you enter the `dump` command through `rsh`. Because the console is not a standard output device, you cannot write to standard output if you enter the `dump` command on the console.

Incremental backups

You can specify the level of backup in a `dump` command ("dump level"), which ranges from level 0 to level 9. A level 0 backup is a full backup: It writes all data in the dump path to the backup media. Backups at dump level ranging from level 1 to level 9 are incremental backups. In an incremental backup, only files changed since the previous level are written to the backup media.

Where you can enter the dump command

You can enter the `dump` command through the console or through `rsh`.

Benefits of entering the dump command through rsh

Entering the `dump` command through `rsh` gives you these benefits:

- When the `dump` command is in progress, you can still use the console to manage the filer. If the `dump` command entered on the console is backing up a large number of files, you cannot use the console for a long time.
- You can start multiple `dump` commands through `rsh`.
- It is less likely to inadvertently terminate the `dump` command. If you enter a `dump` command on the console, it could be terminated by a Ctrl-C entered on a host connected to the filer using `telnet`.
- You can automate filer backups through shell scripts.

- You can write data to standard output.

Benefits of entering the dump command on the console

If you enter the `dump` command on the console, you can read and respond to screen messages displayed by the command. For example, the command might prompt you for another tape to complete the backup. A `dump` command entered through `rsh` aborts and does not generate any messages when the command needs user intervention.

Format of the backup data

About this section

When the filer executes the `dump` command, it displays messages showing the different passes of the `dump` command. This section discusses the format of the backup data and what data is written to tape in each pass of the `dump` command.

Backup data format

The backup data format is organized based on inodes. An inode for a file or a directory contains information for tracking the file's or the directory's type, time stamps, bad blocks, and so on.

On each tape, the `dump` command creates two maps:

- The first map shows which inodes are used in the directory to be dumped. The filer uses this map to determine which files have been deleted or moved between incremental dumps.
- The second map shows which inodes have been written to the tape. The filer uses this map to verify the accuracy of the restore operation when the backup data is restored.

The `dump` command writes files and directories to tape by inode numbers.

Five passes of the dump command

The `dump` command consists of five passes. After you enter the `dump` command, the filer displays messages showing which pass is in progress.

In passes 1 and 2, the filer traverses directories to search for files to be backed up. The filer backs up a file if the file meets these requirements:

- The file is included in the dump path.
- The file has changed since the previous backup at a lower dump level.
- The file is not excluded by the exclude list specified in the `dump` command.

Example: If you initiate a level-1 `dump` command to back up */vol/vol0*, the filer searches for files in */vol/vol0* that have changed since the previous level-0 backup.

In passes 1 and 2, the filer also creates the maps described in "Format of the backup data."

In passes 3 and 4, the filer writes the data to tape in ascending inode order.

In pass 5, the filer writes the ACL information to tape.

How the dump command writes and stores data on tape

About this section

The `dump` command transfers a number of blocks of data at a time to an output file on tape. This section provides information about how the command writes and stores these blocks of data, which helps you decide the `dump` command format appropriate for your backup.

Meaning of tape block

A tape block is 1 kilobyte of data. In the `dump` command, you can specify the number of tape blocks that are transferred in each write operation. This number is called the "blocking factor."

Caution

If you plan to restore the backed-up files on a computer other than the filer, make sure that the blocking factor you choose does not exceed the maximum blocking factor supported by that computer. If you use a filer to restore the backup data, the blocking factor must not exceed 63.

Meaning of tape file

A tape file is a dump output file on tape. It can also be an entire tape. You can back up a dump path to one tape file on a tape, multiple tape files on a tape, or multiple tape files on multiple tapes.

In the `dump` command, you can specify the maximum size of the tape file in terms of tape blocks. For example, if you want the maximum tape file to be 2 GB, specify 2,097,151. (That is, the largest tape file can contain 2,097,151 tape blocks, which are 1 kilobyte each.)

Caution

If the backed-up files are to be restored on a computer other than the filer, make sure that the tape file size you choose does not exceed the maximum tape file size supported by that computer. If you use a filer to restore the backup data, the tape file size can be as large as a single tape.

When the dump command writes to multiple tape files

When the data being backed up exceeds the capacity of a tape, the `dump` command automatically writes to the next tape file specified in the command. If the current tape file is the last tape file listed on the `dump` command, the filer prompts you to load another tape.

If you specify the maximum tape file size in the `dump` command, the command writes to the next tape file when the backup data reaches that size, regardless of the amount of space left in the current tape.

A tape can contain multiple tape files, but any tape must contain at least one complete tape file.

Different types of tape files

The `dump` command can use these types of tape files:

- Local or remote. (Refer to "Devices used by the dump command" on page 444 for more information about where the tape devices can be located.)
- Rewind or norewind. If the tape file is a rewind file, the filer rewinds the tape after it finishes writing the tape file. If you want the `dump` command to write multiple tape files on the same tape, specify norewind tape files in the `dump` command.

  The name of a norewind tape file begins with an `n` (for example, `nrst0a`); the name of a rewind tape file begins with an `r` (for example, `rst0a`).
- Unload/reload tape files. These tape files apply only if you are backing up to tapes in tape stackers. When you use these tape files, the `dump` command unloads the tape when it reaches the end of tape, and then the tape stacker reloads another tape.

  The name of an unload/reload tape file begins with a `u` (for example, `urst0a`).

  An unload/reload tape file cannot be a norewind tape file. For example, do not specify `unrst0a`.

Note

If the tape stacker is attached to a remote host, the `dump` command can reload the tape only if that host supports automatic reloading.

- Standard output, which is specified as `-` in the `dump` command.

Determining the amount of backup data

Description

Before you enter the `dump` command, you must determine the amount of backup data so that you can estimate the number of tape files and the number of tapes required for the backup.

The procedure for estimating the amount of data depends on whether the data is in a qtree.

Step for estimating the amount of data if you back up a qtree

Follow the step in this table to display the number of kilobytes used for the qtree.

Step	Action
1	Enter the `quota report` command.

Steps for estimating the amount of data if you back up data not in a qtree

The procedure for estimating the amount of backup data depends on whether the filer is mounted on an NFS client or is shared by a CIFS client.

If the filer is mounted on an NFS client: Follow these steps to determine the amount of backup data from an NFS client.

Step	Action
1	On the NFS client, change directory to the mount point.

Step	Action
2	Enter the following command for each directory you want to back up: `du -s pathname ...` **Example:** If the NFS client mounts the filer to */toaster* and you want to back up the */etc* and */home* directories, enter these commands: **`cd /toaster`** **`du -s etc home`** The command output shows the amount of space allocated for the directories. Refer to the documentation for your client system for interpreting the output because the output is written in different units (for example, 512-byte units or 1,024 units) depending on the operating system.

If the filer is shared by a CIFS client: Follow these steps to determine the amount of backup data from a CIFS client.

Step	Action
1	On the CIFS client, point to the shared file or directory that you want to back up.
2	Right-click to display the pull-down menu.
3	Click Properties to display the number of bytes used by the file or directory.

Determining the number of tapes for the backup

Description

You must determine the number of tapes required for the backup before entering the `dump` command for these reasons:

- You can ensure that the `dump` command will not run out of tapes and be incomplete.
- You can load all tapes required in the tape drives or stackers in advance for an unattended backup. If you do not load enough tapes before entering the `dump` command and you start the `dump` command from the console, the filer prompts you to load additional tapes. If you start the `dump` command through `rsh`, you do not see the prompts from the filer and the `dump` command cannot be completed because of the lack of tapes.

Prerequisites

You must meet these prerequisites before you can determine the number of tape files required:

- The filer must be able to display tape drive information. To learn how to display tape drive information, refer to "Displaying tape device information" on page 489.
- You must know the amount of data to be backed up. For information about determining the amount of backup data, refer to "Determining the amount of backup data" on page 451.

Steps

Follow these steps to determine the number of tapes required for the backup.

Step	Action
1	Determine the capacity of the tape drives you are using for the backup.
2	Determine the amount of data that needs to be backed up.

Step	Action
3	Determine the amount of space that will be left unused on a tape. For example, if the `dump` command specifies several tape files, the command automatically writes to the next tape file even though there is space left in the current tape. **Example:** The following `dump` command starts the second tape file on `rst1a` after writing 2,000,000 tape blocks to the first tape file, even though the tape in `rst0a` contains unused space: `dump 0ufB rst0a,rst1a 2000000 /vol/vol0`

Prerequisites for the dump command

About this section

The prerequisites for the `dump` command depend on the format of the `dump` command. For example, the prerequisites for backing up data to a remote tape drive and to a local tape drive are different. This section describes the general prerequisites that must be met regardless of the `dump` command format. It also describes the specific prerequisites for specific backup procedures.

General prerequisites

You must meet these prerequisites for the `dump` command to run successfully:

- You must have enough tape space for the backup. Refer to "Determining the number of tapes for the backup" on page 453 to determine the number of tapes required for the backup.
- The tape drive is available. That is, it is not currently used for data backup or recovery.

Prerequisites for backing up to a nonqualified tape drive

The tape drive must be included in the */etc/cloned_tapes* file.

Prerequisites for backing up to a remote tape drive

These are the prerequisites for backing up data to a tape drive attached to a remote host:

- You must know the maximum blocking factor supported by the remote host to which data is backed up. Some computers allow a blocking factor greater than the default used by the `dump` command (63). Make sure that the blocking factor you specify in the `dump` command does not exceed the maximum blocking factor supported by the remote host.
- If you use a blocking factor greater than 63, you must know the maximum blocking factor supported by the system that you plan to use for restoring the data. If the maximum blocking factor is 63 on the system for restoring data, that system cannot restore data from a tape file created with a blocking factor greater than 63.
- You must know the maximum tape file size supported by the remote host to which data is backed up. For example, some UNIX systems do not support

tape files larger than 2 GB. In the `dump` command, you need to specify the tape file size appropriate to your remote host.

- If you use a tape file size larger than 2 GB, you must know the maximum tape file size supported by the system that you plan to use for restoring the data. If the maximum tape file size is 2 GB on the system for restoring data, that system might not be able to restore data from a tape file greater than 2 GB.
- The remote host must support the `rmt` protocol.
- The filer being backed up must have a trusting relationship with the remote host.

Recommendations for performing a backup

About this section

Some recommendations in this section are applicable to all kinds of backups; some recommendations depend on your priorities, such as your need to minimize backup time or to minimize tape handling.

General recommendations

Follow the recommendations in this section when you back up data from the filer.

Avoid backing up too much data in a single dump command: The reasons are listed below:

- The `dump` command cannot be restarted. That is, if the `dump` command encounters an error, you cannot correct the error and proceed from the point where the command fails. You must start the command from the beginning.
- The backup takes a long time to finish. This leaves you with a long time period during which changed data cannot be written to tape by an incremental backup.

Store incremental backups for the same dump path on the same tape: If you want to do a subtree restore, you must restore from all incremental backups to restore the most recent data. Having all the incremental backups in the same tape minimizes tape management during the restore process.

Write down qtree information before backing up qtrees: Record the information about qtrees if your backup contains qtrees. This is because when you restore data from tape, the filer does not automatically re-create the qtrees. You must re-create the qtrees into which data is to be restored. Having a record of all qtrees enables you to re-create the qtrees quickly.

Recommendations for minimizing backup time and data loss

The recommendations in this section enable you to minimize both the time required to perform a backup and the possibility of data loss. This is because the shorter the time for the `dump` command to finish, the more incremental backups you can perform. A large number of incremental backups minimize the amount of unrecoverable data.

Note

There is a disadvantage to having a large number of incremental backups: When you restore data, you must restore from all the incremental backup tapes, which might take a long time.

Use multiple local tape drives: Attach the maximum number of tape drives to the filer. The filer can write faster to a local tape drive than to a tape drive attached to another system.

Organize data to be backed up: The first two passes of the `dump` command run faster if the dump path is one of these:

- A volume
- A qtree
- Data not belonging to any qtrees

Limit the amount of data in each backup: Limit the amount of data in a volume or qtree to be backed up to 200 GB.

Schedule the backups appropriately: Schedule backups when the load on the filer is light. Refer to "Example of backing up the entire filer" on page 468 for an example of executing one full backup and several incremental backups a night.

Avoid using an exclude list: If you exclude certain files from the backup, the `dump` command takes longer because it must compare each file name to the strings in the exclude list to determine whether it should back up the file.

Recommendations for minimizing downtime during data recovery

The method used for restoring data has a greater effect on downtime than does the `dump` command. However, it helps to keep the data in the dump path to 100 GB or less because a full restore is more efficient than a partial restore. For example, it takes less time to restore an entire 100-GB qtree than to extract 100 GB of data from a 500-GB backup.

Recommendations for minimizing the number of tape drives required

Because the filer supports the `rmt` protocol, several filers can share the same tape drive. Attach the tape drive to the filer with the most data to back up. Follow these guidelines if multiple filers back up to the same tape drive:

- Use a private network for the backup so that the traffic load on the network does not slow down the backup process.

- Schedule the `dump` command on each filer so that it starts only when no other filers are using the tape drive.

The dump command syntax

Command syntax

The `dump` command syntax is as follows:

`dump [ option argument ] path`

Rules for entering the dump command

The following list describes the rules for entering the `dump` command:

- *option* can be a list. You must list all options together; do not separate the options by commas or spaces.
- *argument* can be a list of arguments, each of which is associated with an option. The arguments are separated by spaces.
- If an option takes a list of arguments, the arguments in the list are separated by commas.
- List the arguments in the same order as you list the options.
- The dump level must be the first option. You can type other options in any order.
- *path* is the complete path name of the file or directory to be backed up by the `dump` command.
- Always precede the volume name by `/vol/` even though the volume is a root volume. This is because between different levels of backups, you might have changed the root volume.

Note

Some options do not have any arguments.

Example of a simple dump command

The following example illustrates how you use the `dump` command syntax:

```
dump 0ufb rst0a 63 /vol/vol0/
```

Options: The following list describes the options:

- `0` is the dump level.
- `u` updates the */etc/dumpdates* file.
- `f` specifies the tape file.
- `b` specifies the blocking factor.

Arguments: The `0` and `u` options do not have arguments.

The argument for the `f` option is `rst0a` (the tape file name), and the argument for the `b` option is `63` (blocking factor). The `rst0a` and `63` arguments must be listed in the same order as the `f` and `b` options.

Path: The path name is */vol/vol0/*, which means that all files and directories in the *vol0* volume are backed up to tape.

Note

The `dump` command consists of more options than those described in this section. All options are described in greater detail in the next section.

Descriptions of dump options

The following table describes the meanings of the `dump` command options.

Option	Meaning
Dump level	It is mandatory. It can be a number from 0 to 9. Level 0 is a full backup; levels 1 through 9 are for incremental backups.
`A`	The `dump` command does not back up Windows NT ACLs.
`b`	It is the blocking factor. It takes a number as an argument, which is the number of 1-KB blocks in each write. If you use a local tape drive for backup, the number should be from 4 to 64. If you use a remote tape device for backup, the number must be from 4 to 256. The `dump` command generates an error message if you specify a number that is out of range for your tape drive. The default argument is 63.
`B`	It takes a number as an argument, which is the number of tape blocks in a tape file. The `dump` command writes the specified number of tape blocks to a tape file before starting a new tape file. The argument must be equal to or larger than the argument to the `b` option. If you do not specify this option, the `dump` command writes data until it reaches the end of tape.
`f`	It is mandatory. It takes the tape file name as an argument. You can specify a comma-separated list of tape file names.
`l`	It backs up only specific files and directories in the dump path. It must be used with the `n` option.

Option	Meaning
`n`	It takes a string as an argument, which is the name of the backup to be recorded in the */etc/dumpdates* file. It is required if you use the `l` option.
`Q`	It takes a volume name as argument. The `dump` command backs up all data in the specified volume that does not reside in a qtree. You cannot perform incremental backups on data that is backed up with the `Q` option. However, you can use both the `Q` option and `u` option to record the backup in the */etc/dumpdates* file. The entry in the */etc/dumpdates* file enables you to keep a history of the backups.
`u`	The `dump` command updates the */etc/dumpdates* file, which contains the dump path, the dump level, and the creation time of the snapshot used by the `dump` command. You must use this option if you plan to perform incremental backups in the future.
`X`	It takes a string as argument. You can specify a comma-separated list of strings. If the name of a file in the dump path matches one of the strings, the `dump` command excludes that file from the backup. Each string for the `X` option applies to files at every directory under the dump path. The following list describes the rules for specifying the strings: ◆ To exclude a file, the name of the file must match the string exactly. For example, if you specify `core`, only those files whose names are `core` are excluded. A file named `a.core` is not excluded. ◆ You can use the asterisk (*) as a wildcard character. ◆ The wildcard character must be the first or last character of the string. Each string can contain up to two wildcard characters. For example, you can specify `*.core`, `core.*`, or `*core.*`, but not `core*.1`. ◆ Because the strings in the list are comma-separated, if you want to exclude files whose names contain a comma, precede the comma in the string with a backslash. ◆ You can specify up to 32 strings for the `X` option.

Using the dump command to back up data to tape

Description

You can enter the `dump` command at any time to back up data in a specified path. After the `dump` command is finished, the data in the path is written to tape.

Prerequisites

You must meet all the prerequisites described in "Prerequisites for the dump command" on page 455.

Restrictions

You can run up to four `dump` commands in parallel.

Steps

Follow these steps to execute a `dump` command.

Step	Action
1	Prepare the number of tapes required, following the steps described in "Determining the number of tapes for the backup" on page 453.
2	Determine and note the blocking factor used in the backup. For more information about blocking factors, refer to "Meaning of tape block" on page 449.
3	Determine and note the tape file size used in the backup. For more information about tape file size, refer to "Meaning of tape file" on page 449.
4	Enter the `dump` command in the appropriate format, using the examples that follow as a guide.

Examples of level-0 backups to a local tape file

The following list provides examples of level-0 backups to a local tape file.

- **`dump 0f rst0a /vol/vol1/users/tom/specs`**

  The command performs a full backup of the */vol/vol1/users/tom/specs* directory. After the `dump` command is finished, the filer rewinds the tape.

- **`dump 0uf rst0a /vol/vol1/users/tom`**

  The command performs a full backup of the */vol/vol1/users/tom* directory and records the backup in the */etc/dumpdates* file.

- **`dump 0fQ rst0a /vol/vol1`**

  The command performs a full backup of all data in the */vol/vol1* volume that does not belong in any qtrees.

Examples of backups to a remote tape file

- **`dump 0f netapp1:nrst0a /vol/vol1`**

  The command performs a backup to a tape drive attached to a filer named netapp1.

- **`dump 0f unix_machine:/dev/rst0 /vol/vol1`**

  The command performs a backup to a tape drive on a UNIX system.

Example of an incremental backup to a local tape drive

- **`dump 1uf nrst0a /vol/vol1`**

  The command performs a level-1 backup of the */vol/vol1* volume to `nrst0a`. This means that only files that have changed since the most recent level-0 backup are written to tape. After the command is finished, it does not rewind the tape.

Examples of backups to multiple tape files

- **`dump 0f rst0a,rst1a /vol/vol1`**

  The command backs up the */vol/vol1* volume to the tape on `rst0a`. If the volume exceeds the capacity of the tape on `rst0a`, the command writes the second tape file to the tape on `rst1a`. Otherwise, it does not use the tape on `rst1a`.

- **`dump 0f urst0a,urst0a /vol/vol1`**

  The command backs up the */vol/vol1* volume to the tape in `urst0a`. After the command finishes writing to the first tape drive, it does not rewind the tape. If the volume exceeds the capacity of the tape on `urst0a`, the command writes the second tape file to the same tape drive on `urst0a`. If the command needs to write the third tape file, it prompts you to load a new tape.

Note

In this example, if you know in advance that the backup requires two tape files, use `rst0a` as the second tape file name in the `dump` command so that the command automatically rewinds the tape after it is finished.

Example of backing up a directory from a snapshot

- **`dump 0f rst0a /vol/vol1/.snapshot/weekly.0/home/users`**

  From the weekly snapshot, the command performs a level-0 backup of the */vol/vol1/home/users* qtree.

Example of backups to a tape stacker

- **`dump 0f urst0a,urst0a,urst0a,nrst0a /vol/vol1`**

  The command backs up the */vol/vol1* volume to a tape stacker. After the command finishes writing a tape file, it unloads the tape and reloads another tape. After the command finishes writing the fourth tape file, it does not rewind the tape.

Example of backing up multiple files or directories in one dump command

- **`dump 0ufnl rst0a user.1.3.5 /vol/vol1/home`**

  The `n` option specifies the name of the backup (`user.1.3.5`), which is recorded in the */etc/dumpdates* file. If you do not specify a name, the `dump` command fails. The `l` option specifies that you interactively enter the names of individual files and directories to be backed up from the */vol/vol1/home* directory.

  The filer displays some messages and a prompt for the names of the files and directories you want to back up. Enter each name as a path name relative to the dump path in the `dump` command. Do not specify ".." or specify a directory that contains symbolic links. To end the list of names, use a blank line or press Ctrl-D. Then the `dump` command displays messages about the progress of each pass of the command.

  The following example shows how to enter the relative path names for the directories to be backed up. The example ends the list of path names with a blank line.

```
DUMP: creating "snapshot_for_backup.0" snapshot.
creating.....................................................
DUMP: Date of this level 0 dump: Tue Jun  3 12:47:14 1997
DUMP: Date of last level 0 dump: the epoch
DUMP: Dumping /vol/vol0/home to nrst0a
DUMP: mapping (Pass I) [regular files]
DUMP: Reading file names from standard input
user1
user3
user5
```

Example of backing up data without ACLs

- **`dump 0uAf rst0a /vol/vol1`**

  The command performs a level-0 backup of the */vol/vol1* volume. The `A` option means that the backup does not include any Windows NT ACL information. Use this option if the files in the volume are for NFS only.

Example of specifying a blocking factor

- **`dump 0ufb rst0a 32 /vol/vol1`**

  The command performs a level-0 backup of the */vol/vol1* volume. The command writes 32 KB of data at a time, enabling you to restore the data from systems that limit each write to 32 KB.

Example of specifying a tape file size

- **`dump 0ufB nrst0a,nrst0a 2000000 /vol/vol1`**

  The command performs a level-0 backup of the */vol/vol1* volume to the tape in `nrst0a` using a tape file size of 2,000,000 tape blocks. After writing 2,000,000 tape blocks to the first tape file, the command does not rewind the tape; it continues writing the second tape file on the same tape. If there is more data to be backed up after the command reaches the end of the second tape file, the command prompts you for a new tape.

Example of excluding files from a backup

- **`dump 0ufX rst0a tmp,*.o,core*,*backup*,usr\, /vol/vol1`**

  The command performs a level-0 backup of the */vol/vol1* volume, which excludes the files that meet one of these requirements:

  - The name is `tmp`.
  - The name ends in `.o` (for example, `program.o`).
  - The name begins with `core` (for example, `core.small`).
  - The name contains `backup` (for example, `spec.backup.1`).
  - The name is `usr,`.

Example of backing up to a tape stacker shared by multiple filers

This section provides a sample schedule for backing up the following filers to the same tape stacker:

- netapp1, which contains 3 volumes
- netapp2, which contains 2 volumes
- netapp3, which contains 2 volumes

The schedule uses the following assumptions:

- The level-0 backup takes 10 hours.
- The tape stacker is attached to netapp1, and the tape stacker contains two tape drives.
- Root on each filer has the permission to write to netapp1.

According to this schedule, the filers perform the following types of backups:

- The filers fully back up each volume once a week.
- The filers perform a level-1 backup and a level-2 backup for each volume every week.
- The filers do not rewind the tape for each level-1 backup so that the subsequent level-2 backup can be written immediately after the level-1 backup on the tape.

Day of week	**dump commands on each filer**
Sunday	netapp1: **`dump 0uf urst0a,urst0a,urst0a /vol/vol1`** netapp3: **`dump 1uf netapp1:nrst1a /vol/vol6`** netapp2: **`dump 2uf netapp1:urst1a /vol/vol4`**
Monday	netapp1: **`dump 0uf urst0a,urst0a,urst0a /vol/vol2`** netapp3: **`dump 1uf netapp1:nrst1a /vol/vol7`** netapp2: **`dump 2uf netapp1:urst1a /vol/vol5`**
Tuesday	netapp1: **`dump 0uf urst0a,urst0a,urst0a /vol/vol3`** netapp1: **`dump 1uf nrst1a /vol/vol1`** netapp3: **`dump 2uf netapp1:urst1a /vol/vol6`**

Day of week	dump commands on each filer
Wednesday	netapp2: **`dump 0uf netapp1:urst0a,urst0a,urst0a /vol/vol4`** netapp1: **`dump 1uf nrst1a /vol/vol2`** netapp3: **`dump 2uf netapp1:urst1a /vol/vol7`**
Thursday	netapp2: **`dump 0uf netapp1:urst0a,urst0a,urst0a /vol/vol5`** netapp1: **`dump 1uf nrst1a /vol/vol3`** netapp1: **`dump 2uf urst1a /vol/vol1`**
Friday	netapp3: **`dump 0uf netapp1:urst0a,urst0a,urst0a /vol/vol6`** netapp2: **`dump 1uf netapp1:nrst1a /vol/vol4`** netapp1: **`dump 2uf urst1a /vol/vol2`**
Saturday	netapp3: **`dump 0uf netapp1:urst0a,urst0a,urst0a /vol/vol7`** netapp2: **`dump 1uf netapp1:nrst1a /vol/vol5`** netapp1: **`dump 2uf urst1a /vol/vol3`**

Example of backing up the entire filer

This section provides a sample backup schedule for backing up a filer with seven volumes. The schedule uses the following assumptions:

- A low rate of change enables you to fit multiple incremental backups on one tape.

- Two incremental backups performed on each volume each week offer sufficient data protection due to a low rate of change.
- The level-0 backup takes 10 hours.

According to this schedule, the filer performs the following types of backups:

- The filer fully backs up each volume once a week.
- The filer performs a level-1 backup and a level-2 backup for each volume every week.
- The filer does not rewind the tape for each level-1 backup so that the subsequent level-2 backup can be written immediately after the level-1 backup on the tape.

Day of week	dump commands

Sunday	`dump 0uf urst0a,urst0a,urst0a /vol/vol1`
	`dump 1uf nrst1a /vol/vol6`
	`dump 2uf urst1a /vol/vol4`

Monday	`dump 0uf urst0a,urst0a,urst0a /vol/vol2`
	`dump 1uf nrst1a /vol/vol7`
	`dump 2uf urst1a /vol/vol5`

Tuesday	`dump 0uf urst0a,urst0a,urst0a /vol/vol3`
	`dump 1uf nrst1a /vol/vol1`
	`dump 2uf urst1a /vol/vol6`

Wednesday	`dump 0uf urst0a,urst0a,urst0a /vol/vol4`
	`dump 1uf nrst1a /vol/vol2`
	`dump 2uf urst1a /vol/vol7`

Thursday	`dump 0uf urst0a,urst0a,urst0a /vol/vol5`
	`dump 1uf nrst1a /vol/vol3`
	`dump 2uf urst1a /vol/vol1`

Friday	`dump 0uf urst0a,urst0a,urst0a /vol/vol6`
	`dump 1uf nrst1a /vol/vol4`
	`dump 2uf urst1a /vol/vol2`

Saturday	`dump 0uf urst0a,urst0a,urst0a /vol/vol7`
	`dump 1uf nrst1a /vol/vol5`
	`dump 2uf urst1a /vol/vol3`

Data Recovery 15

Introduction to data recovery

Why you want to restore data from tape

This section describes why you want to restore data to the filer from tape.

Files were deleted from disk but backed up to tape: For example, if you inadvertently delete a file and want to restore the file, you can recover the file from tape.

Files are corrupted: If some files are corrupted, you can restore the subtree containing the files.

No disk slots are available for expansion: If the filer runs out of storage space, you can do the following tasks:

- Back up the entire filer
- Replace the current disks with disks of greater capacity
- Set up volumes on new disks
- Restore the filer from tapes

The entire filer is damaged and unusable: If the entire filer is damaged and unusable, call Network Appliance Technical Support to determine whether you can repair the filer.

If you can repair the filer but some files are deleted during the repair process, you need only restore the missing files. Otherwise, you need to reinitialize the disks and restore all files from tape.

Caution

Before initializing disks on your filer, call Network Appliance Technical Support for instructions. Initializing disks destroys all data on your disks.

When you do not recover data from tape

Recover a deleted file from tape only if you cannot recover the deleted file from any snapshot on the filer. If the file is in a snapshot, copying the file to the active file system is faster than recovering the file from tape.

Different methods for recovering data

The filer supports these methods for recovering data that was backed up by the filer's `dump` command:

- Using the `restore` command on the filer to restore the file from a local or remote tape drive. If you use a remote tape drive, the host for the tape drive must support the `rmt` protocol.
- Using the UNIX `ufsrestore` command on a client that supports the SunOS 5.*x*/Solaris 2.*x* `ufsdump` format.
- Using third-party backup management software on a computer that supports the Network Data Management Protocol (NDMP) to restore data from a local or remote tape drive. For more information about NDMP, refer to Chapter 17, "Network Data Management Protocol (NDMP)."

Caution

Use the Solaris `ufsrestore` command only if your filer runs NFS exclusively. If your filer runs the CIFS protocol, do not use the Solaris `ufsrestore` command. Doing so results in data loss.

What data cannot be recovered

The `restore` command cannot recover data that was backed up by commands issued on other systems. For example, do not try to use the filer's `restore` command to recover a tape file created by the Solaris `ufsdump` command.

UNIX file permissions and Windows NT ACLs

A file recovered by the `restore` command has the same UNIX file permissions and Windows NT ACLs as it did when it was backed up. If you restore a file that has only UNIX file permissions to an NTFS qtree or volume, the file has no Windows NT ACLs. The filer uses only the UNIX file permissions on this file until you create Windows NT ACLs on it.

Scope of this chapter

This chapter describes how to use the filer's `restore` command for data recovery. It does not describe how to use data recovery commands or programs on other systems. Refer to the documentation accompanying the other systems for information about their data recovery commands or programs.

The restore command syntax

The restore command syntax

The syntax for the `restore` command is as follows:

`restore [`*`function_key`*`] [`*`options`*`] [`*`arguments`*`] [`*`subtree`*`]`

Rules for using the restore command

Follow these rules when you enter the `restore` command:

- Specify no more than one function key.
- Specify multiple options without intervening spaces.
- Enter the arguments for each option in the order that you specify the options. Separate each argument from the next with a space.
- Place the subtree parameter after the last argument.

The restore command function keys

The following table describes the function keys for the `restore` command.

Key	Meaning
`r`	Rebuilds the file system or subtree. If you are applying incrementals, this must be the only option specified.
`R`	After the `restore` command is interrupted, this function key restarts data recovery from the last tape file used successfully by the command.
`t`	Lists all the file names on a tape. If you specify a path name, only the files in the path name are listed.
`x`	Extracts an individual file or subtree from the backup.

The restore command options

The following table describes the options for the `restore` command.

Option	Argument	Meaning
`b`	*number*	Specifies the blocking factor. Use the same argument for the `b` option in the `dump` command.

Option	Argument	Meaning
`f`	*tape_file*	Specifies the name of the tape file. If you specify "-", `restore` reads from standard input.
`D`	*pathname*	Specifies the absolute path name of a directory into which the files are restored. Without the path name, the files are restored to the directory from which they were backed up. If you created a backup before you installed multiple volumes on the filer, specify the path name into which the backup is restored. For example, if you backed up the */home* directory when the filer contained a single volume and you want to restore */home* to */vol/engineering/home*, specify */vol/engineering/home* as the target path name in the `restore` command. Otherwise, the *home* backup is restored to the */home* directory of the root volume.
`s`	*number*	Specifies the number of the file if multiple tape files exist on a tape. File numbering starts at 1.
`v`		Specifies that `restore` takes place in verbose mode. That is, `restore` displays each file name preceded with its file type. The filer restores files faster without the `v` option.
`y`		Specifies that `restore` not ask whether it should abort when getting an error. That is, if there are bad blocks, `restore` skips over them and continues. This option is particularly useful if you use `restore` through `rsh`. This is because without the `y` option, if `restore` through `rsh` encounters a read error, it aborts immediately.
`A`		Specifies that `restore` does not restore Windows NT ACLs.

Using the restore command

Description

Use the `restore` command if you want to recover data that was backed up by the filer's `dump` command. After you recover the data, the files contain the same data as they did when you ran the `dump` command. You can use the `restore` command at any time.

Restrictions

This section describes the restrictions of the `restore` command.

The i function key of the Solaris ufsrestore command: The `restore` command is similar to the Solaris `ufsrestore` command, except that the `restore` command doesn't support the `i` function key of `ufsrestore`. This function key enables you to specify interactively individual files and directories to be restored from a tape. However, you can restore individual files and directories using the `x` argument of the filer `restore` command.

Incremental-only restores: If you want to restore all files in a backup using the `restore r` command, you must begin the restoration from a level 0 backup. You can, however, restore a specified file or directory from an incremental backup tape when you use the `restore x` command.

Parallel restores: The filer supports up to three simultaneous `restore` commands.

Prerequisites

This section describes the requirements that you must meet for the `restore` command to be completed successfully.

- The space required for the `restore` command to be completed is about 100 MB more than the amount of data to be restored. The command aborts if it runs out of space. If you want to perform a full restore, the additional space should be in the directory that is the root of the backup. If you want to restore only some data in a backup, the additional space should be in the */etc/tmp* directory of the volume where data is to be restored.
- The `restore` command does not restore qtree information. Before you restore a volume containing qtrees, create the qtrees. Refer to "Performing a full restore of a volume containing qtrees" on page 478 for the procedure for restoring qtrees.

- Before you perform a full restore, make sure that the directory into which you restore data does not contain the *restore_symboltable* file. The `restore` command uses the *restore_symboltable* file for incremental restores and for resuming an interrupted restore. If the *restore_symboltable* file exists in the directory because of an unsuccessful restore, remove the file before starting the full restore, or the full restore fails.

Where you can enter the restore command

You can enter the `restore` command through the console or through `rsh`. Entering the `restore` command through `rsh` gives you these benefits:

- When the `restore` command is in progress, you can still use the console to manage the filer.
- You can start multiple `restore` commands through `rsh`.
- It is less likely that someone inadvertently terminates the `restore` command. If you enter a `restore` command on the console, it could be terminated by a Ctrl-C entered on a host connected to the filer using `telnet`.

However, if you enter the `restore` command on the console, you can read and respond to screen messages displayed by the command. For example, the command might prompt you for another tape to complete the recovery.

Steps

Follow these steps to restore files to the filer.

Step	Action
1	Place the tape containing the first tape file of the backup in the tape drive.
2	Enter the `restore` command as follows: **`restore [`*`function_key`*`] [`*`options`*`] [`*`arguments`*`] [`*`subtree`*`]`**
3	If prompted, insert the next tape in the backup.
4	Repeat Step 3 until the restore is complete.

Performing a full restore of a volume containing qtrees

Description

The `restore` command does not restore qtree information. For example, if you need to restore a volume or a part of a volume that contains qtrees, you must first create the qtrees into which data is to be restored.

Steps

Follow these steps to fully restore a volume containing qtrees.

Step	Action
1	Obtain the record of the qtree information as of the last volume backup.
2	Create all qtrees according to the record.
3	Set the security style of each qtree to the value that is the same as when the volume was backed up.
4	If you want to recover the entire volume with one `restore` command, go to Step 5. Using one `restore` command requires less time for recovering the entire volume than using multiple `restore` commands. If you want to use multiple `restore` command to recover all qtrees in the volume, go to Step 6. Using multiple `restore` commands enables you to quickly recover those qtrees that are needed the most.
5	Enter the `restore` command to recover the entire volume. **Example:** **`restore rfD rst0a /vol/vol0`**
6	Enter the `restore` command to restore each qtree in the volume. **Examples:** **`restore xfD rst0a /vol/vol0/NTusers /NTusers`** **`restore xfD rst0a /vol/vol0/UNIXusers /UNIXusers`**

Examples of the restore command

Example of restoring a subtree

The following steps show how to restore a subtree named */vol/vol0/home* from a local tape drive. The subtree was backed up as a subtree, not as a directory within a subtree.

Step	Action
1	Enter the `sysconfig -t` command on the filer to determine the name of the tape device.
2	Delete all files in the subtree before the restore.
3	Install the tape that contains the level 0 backup for the *home* subtree in a local tape drive.
4	Enter the following command if the subtree is backed up to only one tape: **`restore rfD rst0a /vol/vol0/home`** You do not have to specify the `D` option and */vol/vol0/home* if the destination directory for the restore is the same as the directory that was backed up. This command restores the backup in `rst0a` to the */vol/vol0/home* directory, using the same block size as when the subtree was backed up to the tape. If the backup of the subtree is contained in multiple tape files, `restore` prompts you for the next tape file when appropriate.
5	After the restore is finished, remove the tape from the drive. Install another tape that contains the next lowest level of incremental backup. **Example:** If you have tapes containing level 1, level 2, and level 3 backups of */vol/vol0/home*, load the tape with the level 1 backup in the drive and repeat the `restore` command.
6	After the last incremental backup is restored, from a client, remove the *restore_symboltable* file in the directory that you just restored.

Example of restoring the entire filer

To restore the entire filer, repeat the procedure described in this section for each volume. This section assumes that you already initialized the disks on your filer.

If there is one backup for each volume: If you used one `dump` command to back up each volume, follow the procedure in "Example of restoring a subtree" on page 479 to restore each volume. The only difference is that you use */vol/volume_name* as the directory to which the backup is restored.

If each volume was backed up as subtrees or qtrees: If you used separate `dump` commands to back up files, directories, and qtrees that make up each volume, restore each file, directory, and qtree. For example, if the root volume contains two directories, */vol/vol0/etc* and */vol/vol0/home*, and you used the `dump` commands to back them up separately, follow the steps in the following table to restore the entire volume. In this example, the files are restored from a local tape drive.

Step	Action
1	If you are not restoring qtrees, go to Step 4. If you want to restore qtrees, go to Step 2.
2	Create directories on the filer using the `qtree` command that is used as the top of the subtrees to be restored. **Examples:** `qtree etc` `qtree home`
3	Set the security for each qtree created in Step 2.
4	Install the tape that contains the level 0 backup for the */vol/vol0/etc* subtree in the local tape drive. **Note** The `restore` command doesn't support incremental-only restores. You must begin the restoration from a level 0 dump.
5	Enter the `sysconfig -t` command on the filer to determine the name of the tape device.

Step	Action
6	Enter the following command on the filer: **`restore rfD rst0a /vol/vol0/etc`** This command restores the backup in `rst0a` to the */vol/vol0/etc* directory. If the backup for the subtree is in multiple tape files, `restore` prompts you for the next tape volume when appropriate. See "Examples of restoring from multiple tapes" on page 481 for a sample screen display when more than one tape contains the dump volume.
7	After the restore is finished, remove the tape from the drive. Install another tape that contains the next lowest level of incremental dump. **Example:** If you have dump tapes containing level 1, level 2, and level 3 dumps of */vol/vol0/etc*, load the tape with the level 1 dump in the drive and repeat the `restore` command. Repeat this step until the backup of the highest dump level, level 3 in this example, is restored.
8	From the client, remove the *restore_symboltable* file in the directory that you just restored.
9	Repeat Steps 4 through 8 for the */vol/vol0/home* subtree.

Examples of restoring from multiple tapes

This section provides examples illustrating how to restore */vol/vol0* that was backed up to multiple tapes using the following command:

```
dump 0fB rst0a,rst1a 600 /vol/vol0
```

Restoring the volume to a directory from multiple tapes using two tape drives: The following command restores the volume to the */vol/vol0/myexample* directory from two tapes that are in different tape drives:

```
restore rfD rst0a /vol/vol0/myexample
```

When the `restore` command prompts you for the next tape drive, enter the name of the tape drive, for example, `rst1a`.

Restoring a volume to a directory from multiple tapes using one tape drive: If the two tapes use the same tape device, for example, `rst0a`, remove the tape currently in the tape drive and load the next tape. Then accept the default tape device name, for example, `rst0a`, when the filer prompts you for the name of the device for the second tape.

Example of restoring a named file from multiple tapes

If you want to restore a specific file or directory from a subtree, use the `x` option in the `restore` command. Specify the path name of the file or directory relative to the subtree that was backed up.

Example: If the *vol0* volume was backed up as a subtree and you want to restore all contents of the */vol/vol0/test* directory, specify */test* as the path name in the `restore` command.

Example: The subtree containing the directory to be restored was backed up to two tapes:

```
restore xf rst0a /test
```

In this command, you must specify the tape device for the first tape; it is `rst0a` in this example. This is necessary because `restore` needs to read information about the directory structure of the subtree from the first tape before restoring the data.

The filer then displays the following messages to let you specify the tape devices containing the multiple tape volumes. When the filer prompts you for the tape volume number, start with the last tape volume, as shown in this example:

```
You have not read any tapes yet.
Unless you know which volume your file(s) are on you should start
with the last volume and work towards the first.
Specify next volume #: 2
Mount tape volume 2
Enter “none” if there are no more tapes
otherwise enter tape name (default: rst0a) rst1a
You have read volumes: 2
Specify next volume #: 1
Mount tape volume 1
Enter “none” if there are no more tapes
otherwise enter tape name (default: rst1a) rst0a
```

After prompting for the tape devices, the filer also displays the following question:

```
set owner/mode for '.'? [yn] y
```

To keep the original owner and permission modes for the restored files and directories, enter **y**.

Example of listing files

The following example lists the names of all files backed up to `rst0a`:

```
restore tf rst0a
```

The following example lists the names of all files in the */vol/vol0/home* directory on `rst0a` after you backed up the */vol/vol0* volume:

```
restore tf rst0a /home
```

Restarting the restore command

Description

Restart the `restore` command if data recovery is interrupted for reasons such as a power outage or a Ctrl-C inadvertently entered by someone on the filer console. When you restart the `restore` command, the filer restarts data recovery from the last tape file that was successfully used for restoring data. You do not have to start the command from the beginning of the backup.

Restrictions

The `restore` command has the following restrictions:

- You can restart a `restore` command only if the backup consists of multiple tape files. If the backup contains only one tape file, simply reenter the `restore` command to start the recovery from the beginning of the backup.
- You can restart a `restore` command only if the command is for a full restore. If the `restore` command is for extracting an individual file or subtree from a backup (that is, if you use the `x` function key), you cannot restart the command.

Steps

Enter the same `restore` command as the one that was interrupted with the following changes:

- Replace the `r` function key with the `R` function key to indicate that you want to restart the `restore` command.
- Use the appropriate tape file name for the `f` function key. The tape file should be the one used by the original `restore` command when it was interrupted.

Example

In this example, the following `restore` command was entered on the filer console for performing a full restore from three tape files in a tape stacker:

```
restore rfD urst0a,urst0a,rst0a /vol/vol0
```

When the `restore` command is recovering data from the second tape file, it is interrupted by a Ctrl-C entered on the filer console. Because the second tape file is on an unload/reload device, the filer closes the file and ejects the tape when the `restore` command is interrupted.

To restart the `restore` command, reload the second tape file and enter the following command.

```
restore RfD urst0a,rst0a /vol/vol0
```

The `restore` command restarts data recovery from the beginning of the second tape file of the backup.

How to use a filer tape drive to restore files to another system

About this section

This section describes the requirements for using commands such as `ufsrestore` on another system to restore data from a tape drive attached to the filer. It also describes how to specify the tape drive in the command for restoring data.

This section does not discuss the exact commands that you need to enter on the other system. Refer to the documentation for the other system to learn about its commands for restoring data.

Requirements

The filer with the tape drive must allow access from the other system. That is, the host or filer from which a `dump` (or `ufsdump`) or `restore` (or `ufsrestore`) command is issued must be listed in the following configuration files:

- The */etc/hosts.equiv* file on the filer with the tape drive. Alternatively, you can specify both the host and user in */etc/hosts.equiv.* The filer */etc/hosts.equiv* file contains entries in this format:

  ```
  hostname [ username ]
  ```

  For more information about */etc/hosts.equiv*, see the *na_hosts.equiv*(5) man page.
- The */etc/hosts* file on the filer with the tape drive or in the DNS database if the filer is using DNS.

Also, the filer with the tape drive must be added to the other system's */etc/hosts* file.

Format for specifying filer tape drive

In the commands for dumping and restoring, specify the filer tape drive in the following format:

```
filer:device_name
```

Tape Device Management 16

Introduction to tape device management

Why you want to manage a tape device

You need to manage a tape device when you back up data from the filer to tape or when you recover data from tape to the filer.

Scope of this chapter

This chapter discusses the following topics:

- Displaying information about tape devices attached to the filer
- Positioning tapes

How the filer displays information about various tape drives

Introduction

This section describes how the filer displays tape drive information about different types of tape drives. You need the tape information for planning your backup.

This section does not describe how to display information about tape drives on remote hosts. The filer can only display information about local tape drives.

Qualified tape drives

Qualified tape drives are the tape drives that Network Appliance tests with the filer. These tape drives are on the Network Appliance price list.

Nonqualified tape drives

Nonqualified tape drives are the tape drives that Network Appliance has not tested with the filer. The filer can boot with nonqualified tape drives. However, for the filer to write to the drives, you must include information about the tape drive in the */etc/cloned_tapes* file.

The */etc/cloned_tapes* file enables the filer to register the drive as a clone of a qualified drive. If this file does not exist when the filer boots with a nonqualified tape drive, the filer creates the file.

Format of the /etc/cloned_tapes file

Specify each entry of the */etc/cloned_tapes* file in one of the following formats:

- `clone_vendor_id clone_product_id EMULATES vendor_id product_id`
- `clone_product_id EMULATES product_id`

Example of /etc/cloned_tapes entry

The following example is an entry of the */etc/cloned_tapes* file:

```
QUANTUM DLT9000 EMULATES QUANTUM DLT7000
```

Displaying tape device information

Description

Before you use the `dump` command to back up to a tape device attached to the filer, verify that the filer detects the device. Also, verify the device name to be included in the `dump` command.

Step for displaying information about qualified tape devices

Follow the step in this table to display information about qualified tape devices.

Step	Action
1	Enter the following command: **`sysconfig -t`**

Steps for displaying nonqualified tape devices

Follow the steps in this table to display information about nonqualified tape devices.

Step	Action
1	If the filer has accessed the tape drive through the `dump` or `mt` command, go to Step 3. Otherwise, go to Step 2.
2	Enter the following command to access the tape device: `mt -f` *`device`* `status`
3	Enter the following command: **`sysconfig -t`**

Steps for displaying information about tape stackers

Follow the steps in this table to display information about tape stackers.

Step	Action
1	If the tape stacker's autoload setting is On, go to Step 2. If the tape stacker's autoload setting is Off, go to Step 4.

Step	Action
2	Turn off the autoload setting of the tape stacker.
3	Reboot the filer.
4	Enter the following command: **`sysconfig -m`**

Displaying tape device information along with other filer information

You can use the `sysconfig -v` command to display the filer hardware information. Look for the information about the SCSI adapter to which the tape devices are attached.

Example of the sysconfig -t command for a qualified tape drive

The following example of the `sysconfig -t` command displays information about a qualified tape drive:

```
sysconfig -t
Tape drive (6.4)  Digital DLT7000
rst0l  -  rewind device,        format is: 81633 bpi 40 GB (w/comp)
nrst0l -  no rewind device,     format is: 81633 bpi 40 GB (w/comp)
urst0l -  unload/reload device, format is: 81633 bpi 40 GB (w/comp)
rst0m  -  rewind device,        format is: 85937 bpi 35 GB
nrst0m -  no rewind device,     format is: 85937 bpi 35 GB
urst0m -  unload/reload device, format is: 85937 bpi 35 GB
rst0h  -  rewind device,        format is: 85937 bpi 50 GB (w/comp)
nrst0h -  no rewind device,     format is: 85937 bpi 50 GB (w/comp)
urst0h -  unload/reload device, format is: 85937 bpi 50 GB (w/comp)
rst0a  -  rewind device,        format is: 85937 bpi 70 GB (w/comp)
nrst0a -  no rewind device,     format is: 85937 bpi 70 GB (w/comp)
urst0a -  unload/reload device, format is: 85937 bpi 70 GB (w/comp)
```

Examples of the sysconfig -t command for a nonqualified tape drive

The following example of the `sysconfig -t` command displays information about a nonqualified tape drive that the filer has not registered as a clone:

```
sysconfig -t
Tape drive (6.5) DLT9000 (Non-qualified tape drive)
```

The following example of the `sysconfig -t` command displays information about a nonqualified tape drive that the filer has registered as a clone:

```
sysconfig -t
```

```
Tape drive (6.5) DLT9000 emulates Digital DLT7000.
```

Example of the sysconfig -m command

The following example of the `sysconfig -m` command displays information about a tape library attached to the filer:

```
sysconfig -m
Media changer (6.6)  BHTi     Quad 7
mc0 - media changer device
```

In this example, a tape library with SCSI ID 6 is attached to slot 6 of the filer.

Example of the sysconfig -v command

The following example shows a tape stacker with SCSI ID 6 and a tape drive with SCSI ID 4 attached to slot 6 of the filer:

```
slot 6: SCSI Host Adapter 6 (QLogic ISP 1040B)
   Firmware Version 2.26           Clock Rate 60MHz.
   6: BHTi    Quad 7          1.41
   4: QUANTUM DLT7000         1B41
   In-Band Enclosure Services
```

Using the mt command to control tape devices

The mt command syntax

You can control tape devices with the `mt` command. The syntax of the `mt` command is as follows:

`mt [-f|-t] tapedevice command [ count ]`

This section discusses only the `eom`, `fsf`, `rewind`, `offline`, and `status` commands. Keep in mind the following information:

- The `-f` and `-t` options are interchangeable as far as the filer is concerned. Only the `-f` option is shown here.
- For additional information about controlling tape devices and detailed information about the `mt` command, consult the *mt*(1) man page.
- For information about the format of a filer tape device name, consult the *tape*(4) man page.

Moving a tape to the end of data

You can append material, such as a dump, on a tape device. To do so, use the `eom` command, which moves the tape to the end of data (end of media if the tape is full), as in the following example:

```
mt -f nrst0a eom
```

After this command, you can write to the remainder of the tape. Make sure that there is enough tape for your additional data.

Appending a dump

To skip over a previously created dump file and append a dump on a tape device, use the `fsf` command, as in the following example:

```
mt -f nrst0a fsf 1
```

Rewinding a tape

To rewind a tape, use the `rewind` command, as in the following example:

```
mt -f nrst0a rewind
```

Taking a tape drive off-line

To rewind the tape and, if appropriate, take the tape drive off-line by unloading the tape, use the `offline` command, as in the following example:

```
mt -f nrst0a offline
```

Displaying status information

To display status information about the tape unit, use the `status` command, as in the following example:

```
mt -f nrst0a status

Tape drive: Quantum DLT7000
Status: ready, write enabled
Format: 85937 bpi 70 GB (w/comp)
fileno = 0  blockno = 95603  resid = 0
```

Network Data Management Protocol (NDMP) 17

Introduction to NDMP

Advantages of NDMP

You can use third-party backup management programs that support NDMP Version 1, 2, or 3 to back up or restore the filer. Using NDMP-compliant programs offers these advantages over using the filer's `dump` and `restore` commands:

- You can execute more parallel backup and recovery operations. The filer supports up to 100 NDMP sessions.
- You can choose a backup program that best meets your administrative needs. For example, you can use an NDMP-compliant backup program that provides a graphical user interface for scheduling backups.
- NDMP-compliant programs can back up the filer to remote tape drives faster than the filer `dump` command.
- NDMP-compliant programs can back up the filer to tape libraries.
- It is easier to use NDMP to manage tape devices than to use the filer `dump` and `restore` commands.

Scope of this chapter

This chapter describes how to configure the NDMP server on the filer so that a computer can run an NDMP-compliant application to back up the filer or to restore data to the filer.

Tape devices that can be used by the filer's NDMP server

Tape devices

The filer can read from or write to these devices when using NDMP:

- Tape drives or libraries attached to the filer
- Tape drives or libraries attached to the workstation that runs the backup management program
- Tape drives or libraries attached to a workstation or filer on your network

Restrictions for tape libraries

The following restrictions apply when you use NDMP for backing up the filer:

- You can attach up to 10 tape libraries to the filer to be backed up by NDMP.
- If you want to use NDMP to manage your tape library, you must set the tape stacker's autoload setting to Off. If the autoload setting is On, the filer uses the tape stacker in the same way as it uses a tape drive.

Naming convention for tape libraries

When you use NDMP to back up the filer to attached tape libraries, use the following names for the tape libraries:

- `/dev/mc`*n*
- `/dev/spt`*n*

In the actual name, *n* is a number (for example, `/dev/mc0` or `/dev/spt0`). To verify the name for an attached tape library, use the `sysconfig -m` command.

Typical application of NDMP

The following figure illustrates how you can back up the filer to various devices using NDMP.

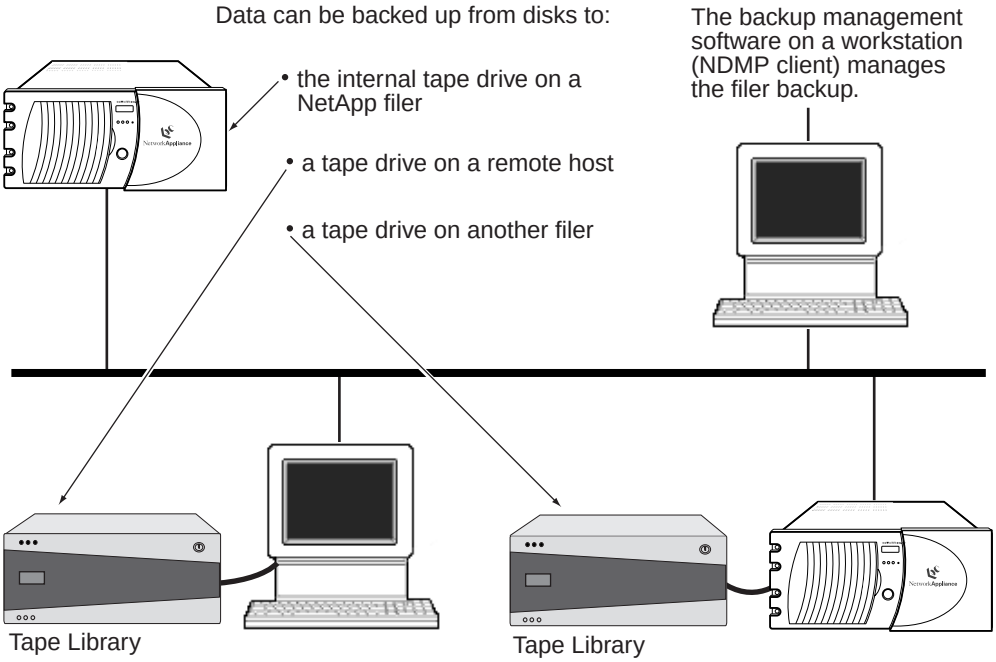
Data can be backed up from disks to:
• the internal tape drive on a NetApp filer
• a tape drive on a remote host
• a tape drive on another filer
The backup management software on a workstation (NDMP client) manages the filer backup.
Tape Library
Tape Library

The ndmpd command

Purposes of the ndmpd command

The `ndmpd` command enables NDMP support on the filer and displays information about backup and restore operations that use NDMP.

The ndmpd command syntax

The syntax for the `ndmpd` command is as follows:

`ndmpd [ on | off | status | probe [`*`session`*`]| kill [`*`session`*`] ]`

The ndmpd command options

The following table describes the `ndmpd` command options.

Option	Description
`on`	Enables NDMP. **Running NDMP after each reboot:** If you want the filer to run NDMP after each reboot, put the `ndmpd on` command in the */etc/rc* file.
`off`	Disables NDMP. After you enter `ndmpd off`, the filer continues processing existing requests but rejects new ones.
`status`	Displays the current state of each NDMP session. Sessions are numbered 0 to 99.
`probe`	Displays detailed information about the NDMP session specified. The information is for debugging. Using the `probe` option without a `session` parameter displays NDMP information for all sessions.
`kill`	Terminates the specified session. Use this option if you want to clear a hung NDMP session.

Example of status option

In the following example, the `ndmpd status` command displays information about session 4:

```
ndmpd status
ndmpd ON.
Session: 4
  Active
  version:                3
  Operating on behalf of primary host.
  tape device:    not open
  mover state:    Idle
  data state:     Idle
  data operation: None
```

Example of probe option

In the following example, the `ndmpd probe 0` command shows detailed status of active session 4:

```
ndmpd probe 0
ndmpd ON.
Session: 4
  isActive:               TRUE
  protocol version:       3
  effHost:                Local
  authorized:             FALSE
  client addr:            10.10.10.12:47154
  spt.device_id:          none
  spt.ha:                 -1
  spt.scsi_id:            -1
  spt.scsi_lun:           -1
  tape.device:            not open
  tape.mode:              Read only
  mover.recordSize:       0
  mover.recordNum:        0
  mover.state:            Idle
  mover.mode:             Read
  mover.pauseReason       N/A
  mover.haltReason        N/A
  mover.dataWritten:      0
  mover.seekPosition:     0
  mover.bytesLeftToRead:  0
  mover.windowOffset:     -1
```

Log messages generated by NDMP

Where log messages are saved

When an NDMP session is in progress, the filer generates messages about the session and saves them to a file named *ndmplog.yyyymmdd* in the */etc* directory. In the actual command, *yyyy* is the year, *mm* is the month, and *dd* is the date. For example, the log file generated on December 7, 1998, is named *ndmplog.19981207*.

If multiple NDMP sessions take place on the same day, the filer saves the information about all sessions to the same log file.

Maximum number of log files on the filer

The filer keeps a maximum of eight log files. They are for the past seven days and the current day. The filer deletes a log file when it is eight days old.

Considerations when using NDMP on the filer

Port used by the NDMP daemon

The NDMP daemon listens on port 10,000 of the filer. Do not use any third-party network management programs that try to use this port, or NDMP does not function properly.

Maximum length of the password used for MD5 authentication

If you use MD5 authentication for NDMP, the filer only accepts passwords that contain eight characters or fewer.

Volume Copy Using the vol copy Command Set 18

About this chapter

Overview of volume copy

The filer can copy data from one volume to another volume at a given time. The result of copying a volume is an off-line volume containing the same data as the source filer at the time you initiated the copy operation.

Chapter contents

This chapter discusses the following topics about copying volumes:

- "Introduction to the filer's commands for copying volumes" on page 504
- "Requirements and recommendation for copying a volume" on page 506
- "Details about copying one volume to another" on page 509
- "Management of a volume copy operation when it is in progress" on page 514

This chapter does not discuss SnapMirror, which replicates data in one volume on another volume (the replica) and periodically updates the replica to reflect incremental changes on the source filer. For more information about SnapMirror and about the differences between copying a volume and replicating a volume, refer to Chapter 19, "Data Replication Using SnapMirror."

Introduction to the filer's commands for copying volumes

Purposes of the vol copy command set

The `vol copy` command set enables you to copy one volume to another. The commands in this command set control copying both data in the active file system and data in snapshots from one volume to another. The source and destination volumes of the copy can reside on the same filer or on different filers.

For more information about snapshots, see Chapter 11, "Snapshots."

When to copy volumes

The following table describes some situations where you might find copying volumes useful.

Situation	Reasons for copying one volume to another
You want to copy data from one filer to another regularly to ensure high data availability.	After you copy the data, clients can switch to the destination filer in the following scenarios: ◆ When you shut down the source filer for software or hardware upgrades. ◆ If a network client process accidentally deletes a large number of files on the source filer, clients can continue to have access to the files when you are restoring the files to the source filer. ◆ If the source filer is not available for reasons such as natural disasters, you can put the destination filer on-line to continue file service.
You want to migrate data from one filer to another.	The destination filer has more storage or is a model that supports newer technology, such as FC-AL disks.

Situation	Reasons for copying one volume to another
You want to move a volume from one set of disks to another on the same filer.	◆ Splitting a volume. **Example:** You can copy the *vol0* volume to the *vol1* volume and then delete duplicated files and directories in these volumes so that the original contents of *vol0* are split into two volumes. ◆ Expanding filer storage. **Example:** You have eight 4-GB disks for the *vol0* volume and four 9-GB spare disks. You can migrate *vol0* to the four 9-GB disks and replace all 4-GB disks with larger capacity disks.

Benefits of the vol copy command set

Although you can copy data on the filer using client programs such as `cpio` or using the filer's `dump` and `restore` commands, the `vol copy` command set offers the following benefits:

- When a `vol copy` command reads and writes data, the filer does not traverse directories on the filer. Data is copied block for block directly from the disks, which means that the filer can finish the copying faster than it could with other methods.
- Using a `vol copy` command, the filer preserves the snapshot data of the source volume. If, in the future, users might need to use snapshots that were taken before data was copied from one volume to another, for example, if users accidentally delete files and need to recover them, use a `vol copy` command for migrating data.

Requirements and recommendation for copying a volume

Requirements for copying a volume

The filers involved in a volume copy operation must meet several requirements for data to be copied successfully. The following list is a brief description of these requirements. The rest of this section provides more detailed information about verifying whether the source and destination volumes meet these requirements.

- The source volume must be on-line and the destination volume must exist and be off-line.
- The destination volume must not be the root volume because the destination volume must be off-line when the filer executes the `vol copy` command, and a root volume must always be on-line.
- The capacity of the destination volume must be greater than or equal to the capacity of the source volume.
- The destination volume must not contain data that you want to preserve.
- The source and destination filers must have a trusting relationship with each other.
- The localhost interface must be specified as a trusted host if data is copied on the same filer.

Verifying the status of each volume

The destination volume must exist before you enter the `vol copy start` command to start copying a volume. If the volume does not exist, the command does not create the volume and the command returns an error. See "Creating volumes" on page 89 for information about how to create a volume.

After you verify that the destination volume exists, check the status of the source volume and the destination volume.

Checking the status of a volume: The source volume must be on-line and the destination volume must be off-line. To verify whether a volume is on-line or off-line, use the following command syntax:

`vol status` *`volume_name`*

Changing the status of a volume: If you need to change the status of a volume, use one of the following command syntaxes:

- `vol offline` *`destination_volume`*
- `vol online` *`source_volume`*

Note

The `vol offline` command takes effect only after you reboot the filer.

Verifying the size of each volume

The capacity of the destination volume must be greater than or equal to the capacity of the source volume, regardless of how many snapshots you try to copy. To verify the capacity of a volume, follow these steps:

1. Enter the following command syntax:

   `df /vol/`*`volume_name`*

   The command displays information about disk space in the active file system and in the snapshot for the specified volume.

2. Add the numbers in the kbytes column in the `df` output. The result is the number of kilobytes of data that the volume can hold.

Verifying the contents of the destination volume

If the destination volume is not a new volume, make sure that the destination volume does not contain data that you might need in the future. This is because after the filer starts copying the source volume, it overwrites the entire destination volume. That is, all data in the active file system and in the snapshot of the destination volume is lost after the filer starts copying the data.

Verifying the relationship between filers

If the source and destination volumes in a volume copy operation reside on two filers, the filers must have a trusting relationship with each other. That is, you must specify each filer as a trusted host in the */etc/hosts.equiv* file of the other filer. The */etc/hosts.equiv* file contains a list of host names, each of which is on a separate line.

Verifying localhost as a trusted host

If you want to copy data between volumes on the same filer, you must specify localhost on the filer as a trusted host in the filer's */etc/hosts.equiv* file. Localhost is the interface through which the filer sends packets to itself.

If you have not already listed localhost as a trusted host, follow these steps to include localhost in */etc/hosts.equiv*:

1. Enter the loopback address, which is 127.0.0.1, in the */etc/hosts* file.

2. Type `localhost` on a separate line in the */etc/hosts.equiv* file.

Recommendation for copying a volume

To avoid performance problems when copying to a different filer, you might want to set up a private network for copying between the source and destination filers. This is because when a filer copies data between two volumes, it floods the network with packets. Users of the filers involved in a volume copy operation might notice a degradation in response time during the copy.

Details about copying one volume to another

Command syntax for copying one volume to another

To copy one volume to another, use the following command syntax:

`vol copy start [ -S | -s snapshot ] source destination`

Specifying the snapshots to copy: The `-S` and `-s` arguments specify the snapshots to copy. The following table describes the snapshots to copy and the resulting snapshots on the destination volume, depending on the argument.

Argument used	Snapshots to copy from the source volume	Snapshots in the snapshot file system of the destination volume
None	The snapshot taken after you enter the `vol copy start` command	A snapshot named *snapshot_for_backup.n*, where *n* is a number starting at 0
`-S`	All snapshots in the snapshot file system of the source volume and the snapshot taken after you enter the `vol copy start` command	All snapshots in the source volume and *snapshot_for_backup.n*, where *n* is a number starting at 0
`-s` followed by the name of the snapshot	The specified snapshot	The specified snapshot

Note

The `vol copy start -S` command does not copy any snapshots that are created when the copying is in progress. For example, if the copying lasts from 11:45 p.m. to 1:00 a.m. the next day, and the filer creates a snapshot named *nightly.1* at midnight, the filer does not copy the *nightly.1* snapshot.

Specifying the volumes involved in the copy: The *source* and *destination* arguments are the names of the volumes. If a volume is on a different filer, precede the volume name with the filer name and a colon. For examples illustrating how to specify volume names, see "Examples of the vol copy start command" on page 510.

Where to enter the vol copy start command

If the copying takes place between two filers, you can enter the command on either the source or destination filer. You cannot, however, enter the command on a third filer that does not contain the source or destination volume.

Examples of the vol copy start command

The following table shows several examples of the `vol copy start` command.

If you want to...	Enter...
Copy all snapshots from the *vol0* volume to the *vol1* volume on the same filer	`vol copy start -S vol0 vol1`
Copy a nightly snapshot from the *vol0* volume to the *vol1* volume on the same filer	`vol copy start -s nightly.1 vol0 vol1`
Create a snapshot in the *vol0* volume to be copied to the *vol1* volume on the same filer	`vol copy start vol0 vol1`
Copy all snapshots from the *vol0* volume to the *vol1* volume on a different filer named filerA	`vol copy start -S vol0 filerA:vol1`

Results of the vol copy start command

The `vol copy start` command generates volume copy operations and produces screen messages showing the progress of the operations.

Volume copy operations: Each `vol copy start` command generates two volume copy operations, as described in the following list:

- One operation is for reading data from the source volume. Screen messages displayed by a `vol copy` command refer to this operation as the `volcopy dump` operation.
- One operation is for writing data to the destination volume. Screen messages displayed by a `vol copy` command refer to this operation as the `volcopy restore` operation.

The filer assigns a volume copy operation number to each operation.

When to use the volume copy operation number: You need the volume copy operation number if you want to stop a volume copy operation or change the volume copy operation speed.

For more information about obtaining the volume copy operation number, see "Checking the status of a volume copy operation" on page 514.

Screen messages from the vol copy command: When the filer is copying a volume, the filer displays messages indicating the percentage of the volume copy operation completed and the number of minutes remaining. When the filer finishes copying all data, it displays the filer prompt.

Maximum number of simultaneous volume copy operations

Each filer supports up to four simultaneous volume copy operations. Whether a filer can execute a `vol copy start` command depends on how many volume copy operations are already in progress on the filer or filers specified in the `vol copy start` command, as illustrated in the following examples.

Example: You can enter the following two `vol copy start` commands on a filer to copy volumes locally:

```
vol copy start vol0 vol1
vol copy start vol2 vol3
```

When these commands are in progress, you cannot enter additional `vol copy start` commands because four volume copy operations are already running on the filer. Two of the operations are for reading the *vol0* and *vol2* volumes, and two of the operations are for writing the *vol1* and *vol3* volumes.

Example: Suppose you enter the following three `vol copy start` commands on a filer named filerA to copy volumes to another filer named filerB.

```
vol copy start vol0 filerB:vol0
vol copy start vol1 filerB:vol1
vol copy start vol2 filerB:vol2
```

When these commands are in progress, filerA runs three volume copy operations to read the volumes, and filerB runs three volume copy operations to write the volumes.

You can enter an additional `vol copy start` command to copy between filerA and filerB because the command adds one more volume copy operation to each filer.

However, you cannot enter an additional `vol copy start` command to copy volumes locally on either filerA or filerB. This is because the additional command would create two volume copy operations, one for reading and one for writing, on the filer that performs the local copying. The filer cannot support these two additional volume copy operations because three operations are already in progress.

Possible errors

If your filer does not meet a requirement described in "Requirements and recommendation for copying a volume" on page 506, the `vol copy start` command generates one or more error messages. The following table explains the meanings of the possible error messages.

Error message	Meaning

`Permission denied.`
 	
`VOLCOPY: Could not connect to filer 127.0.0.1.`
	The source filer does not have permission to copy to the destination filer.
	
	Action: Make sure that the filers have a trusting relationship with each other. If both the source volume and destination volume are on the same filer, remember to include localhost in the */etc/hosts.equiv* file.

Error message	Meaning
`VOLCOPY: volcopy restore: volume is online, aborting`	The destination volume is on-line. **Action:** Take the destination volume off-line and reboot the filer that contains the destination volume.
`VOLCOPY: volcopy restore: volume is too small, aborting`	The destination volume is smaller than the source volume. **Action:** Add more disk space to the destination volume or choose another destination volume of sufficient capacity.

Management of a volume copy operation when it is in progress

Checking the status of a volume copy operation

You can use the following command syntax to check the status of one or more volume copy operations:

`vol copy status [ operation_number ]`

This command displays the status for a specified volume copy operation. If you do not specify the operation number, the command displays the status of all volume copy operations in progress. In the command output, the operations are differentiated from one another with unique volume copy operation numbers, ranging from 0 to 3. For more information about volume copy operation numbers, see "Results of the vol copy start command" on page 510.

Where to enter the vol copy status command: If you start a volume copy operation from the filer's console, you can enter the `vol copy status` command only through `rsh` when the copy operation is in progress. This is because you do not have access to the filer prompt on the console when the filer is copying the volume.

If data is being copied between two filers, you can enter this command through an `rsh` connection to either filer. The operation numbers displayed on the source filer and the destination filer are different because the reading and the writing are considered two different operations.

Example of a vol copy status command

The following example illustrates a `vol copy start` command that copies the *vol0* volume to the *vol1* volume on the same filer:

```
vol copy start -S vol0 vol1

Copy Volume: vol0 on machine 127.0.0.1 to Volume: vol1
Reading the dump stream
VOLCOPY: Starting on volume 1.
This dump contains 257 blocks
10:04 pm : volcopy restore 1 : begun.
10:04 pm : volcopy restore 1 : 5 % done. Estimate 3 minutes
remaining.
.
.
.
10:04 pm : volcopy restore 1 : 95% done. Estimate 1 minutes
remaining.
```

Before the filer prompt is displayed again, you can use the `vol copy status` command on a trusted host of the filer, as shown in the following example:

```
rsh toaster vol copy status

10:04 pm : volcopy dump 0 : 99 % done. Estimate 1 minutes remaining.
10:04 pm : volcopy restore 1 : 99 % done. Estimate 1 minutes
remaining.
No operation 2 in progress.
No operation 3 in progress.
```

In this example, volume copy operation 0, shown as `volcopy dump 0` in the display, is for reading the data from the *vol0* volume; volume copy operation 1, shown as `volcopy restore 1` in the display, is for writing the data to the *vol1* volume.

Aborting a volume copy operation

To stop a volume copy operation, use the following command syntax:

`vol copy abort [ operation_number ]`

The *operation_number* parameter specifies the volume copy operation to be aborted. You can obtain the operation number from the `vol copy status` output.

Caution

If you specify no operation number in the command, the filer aborts *all* volume copy operations. It does not display a help string for the `vol copy abort` command.

If data is being copied between two filers, you can execute this command on either filer.

If you start the volume copying operation from the filer console, you can enter the `vol copy abort` command only through `rsh` because you do not have access to the filer prompt on the console during the copying.

Caution

An incomplete volume copy operation leaves unusable data in the destination volume.

Controlling the speed of a volume copy operation

You can control the speed of a volume copy operation in two ways.

- Before you enter the `vol copy start` command, use the following command syntax to set the default speed for all volume copy operations:

  `options vol.copy.throttle` *`value`*

  The *value* variable specifies the speed, which ranges from 10 (full speed) to 1 (one-tenth of full speed).

- If a volume copy operation is in progress, use the following command syntax to set the speed of a specific operation:

  `vol copy throttle [` *`operation_number`* `]` *`value`*

  In this case, because a volume copy operation is in progress and you do not have access to the filer prompt, you must enter the `vol copy throttle` command through `rsh`.

  The *operation_number* parameter specifies the volume copy operation whose speed you want to adjust. If you do not specify an operation number, the command applies to all volume copy operations that are in progress. The *value* variable specifies the speed, which ranges from 10 (full speed) to 1 (one-tenth of full speed).

Note

The speed for reading data from the source volume and the speed for writing data to the destination volume can be different. In this case, the smaller of the two values determines the time required for the filer to finish copying the data.

Displaying the default speed for copying a volume

Before starting the `vol copy start` command, you can verify the default speed for all volume copy operations using the following command:

```
options vol.copy.throttle
```

It displays the value (1 through 10) to be used by all volume copy operations. The value of the `vol.copy.throttle` option was set at 10 at the factory.

Example of controlling the speed of copying a volume

The following example illustrates changing the speed of all volume copy operations in progress to one-tenth of full speed through `rsh`:

```
rsh toaster vol copy throttle 1

volcopy operation 0: Throttle adjusted from 100% to 10%.
volcopy operation 1: Throttle adjusted from 100% to 10%.
```

Data Replication Using SnapMirror 19

About this chapter

Overview of SnapMirror

The SnapMirror feature replicates data from one volume (the source) to another volume (the mirror) and periodically updates the mirror to reflect incremental changes on the source. The result of this process is an on-line, read-only volume that contains the same data as the source volume at the time of the most recent update. SnapMirror requires a license code.

Chapter contents

This chapter discusses the following SnapMirror topics:

- "Purposes of SnapMirror" on page 518
- "How SnapMirror works" on page 520
- "Snapshots created during data replication" on page 522
- "How SnapMirror works with quotas" on page 525
- "How SnapMirror works with the dump command" on page 526
- "The /etc/snapmirror.allow file" on page 527
- "The /etc/snapmirror.conf file" on page 528
- "Replicating a volume" on page 530
- "Disabling data replication for the entire filer" on page 534
- "Resuming data replication for the entire filer" on page 535
- "Disabling data replication for one volume" on page 536
- "Checking data replication status" on page 538
- "Converting a mirror to a regular volume" on page 540
- "Differences between the vol copy command and SnapMirror" on page 542

This chapter does not discuss the `vol copy` command, which copies a volume to another but does not periodically updates the copy. For more information about the `vol copy` command, refer to Chapter 18, "Volume Copy Using the vol copy Command Set."

Purposes of SnapMirror

Why you want to replicate a volume

You want to replicate a volume if you are in any of the situations described in "When to copy volumes" on page 504. Because SnapMirror offers additional advantages, there are other situations where you want to replicate a volume, as described in the following table.

Situation	How data replication helps
Remote access to data: Users who need read access to a volume are distributed over a large geographical area.	You can replicate the source volume on other filers that are geographically closer to the users. Users accessing a local filer can read the data faster than they could if they connected to a distant filer.
Load balancing: A large number of users need only read access to a volume.	Replicating a volume on multiple filers enables you to distribute the load.
Backup: You need to reserve all processing and networking resources on a filer for serving NFS and CIFS requests.	After replicating data on the source filer, you can back up the data in the mirror to tape. This means the source filer does not have to allocate resources for performing backups.
Data migration: You want to migrate data from one filer to another without interrupting network service.	After replicating the source filer, you can export or share the source filer with read-only permissions so that network clients cannot make additional changes to the source filer. Because the source and destination filers contain identical data, you can change the status of each mirror on the destination filer to a regular volume. For more information about how to convert a mirror to a regular volume, refer to "Setting volume options" on page 90.

Situation	How data replication helps
Disaster recovery: You want to provide immediate access to data after some disaster has caused a volume or a filer to be unavailable.	You can change the status of the mirror to a regular volume. Users can immediately have access to the same data as that on the source volume. The `snapmirrored` volume option controls whether a volume is a mirror or a regular volume. For more information about volume options, refer to "Setting volume options" on page 90.

How SnapMirror works

Command and configuration file for controlling SnapMirror

To use SnapMirror for replicating a volume, use the `vol snapmirror` command set and two configuration files: */etc/snapmirror.conf* and */etc/snapmirror.allow*. The command set is described in "Replicating a volume" on page 530. The configuration files are described in "The /etc/snapmirror.allow file" on page 527 and "The /etc/snapmirror.conf file" on page 528.

How the filer creates a baseline version of the mirror

To use SnapMirror, you create an off-line volume to be used as the destination for data replication. The destination volume is referred to as the "mirror." The mirror can reside on the same filer as the source volume or on another filer.

To replicate data for the first time, the filer transfers all data in all snapshots in the source volume to the mirror. After the filer finishes transferring the data, it brings the mirror on-line. This version of the mirror is the base line for future incremental changes.

How the filer updates the mirror

To make incremental changes on the mirror, the filer takes regular snapshots on the source volume according to the schedule specified in the */etc/snapmirror.conf* file. By comparing the current snapshot with the previous snapshot, the filer determines what changes it needs to make to synchronize the data in the source volume and the data in the mirror. For example, if a file is deleted from the source volume, SnapMirror deletes the corresponding file in the mirror the next time it updates the mirror.

For more information about how SnapMirror uses snapshots to transfer data from the source volume to the mirror, refer to "Snapshots created during data replication" on page 522.

Number of volume copy operations SnapMirror generates

When SnapMirror transfers data from one volume to another, it generates two volume copy operations in the same way as the `vol copy` command. Refer to "Volume copy operations" on page 511 for the meaning of a volume copy operation.

What happens after you replicate a volume

After you replicate a volume, the active file system and all snapshots in the source volume are available on the mirror. As with any volume, you can export the mirror for NFS mounting or add a share corresponding to this volume for CIFS sharing.

SnapMirror makes all changes to the mirror at the same time. If you have an open file in the mirror while the filer is updating the mirror, you do not see the changes immediately. After the mirror update is finished, if you open the file, you see the changes. This is similar to the situation where another user changes a file in a regular volume when you are reading the file. You can see the changes the next time you open the file.

Differences between a mirror and a regular volume

The following list describes the differences between a mirror and a regular volume:

- A mirror has snapmirrored status, which means that it contains read-only data and network clients cannot write data to the mirror.
- The filer does not create automatic snapshots on the mirror based on the snapshot schedule.
- You cannot use the `qtree create` command to create qtrees on a mirror. However, if qtrees exist in the source volume, the filer mirrors the qtrees to the mirror.
- You cannot enable quotas on a mirror.

Snapshots created during data replication

Naming conventions for snapshots used by SnapMirror

SnapMirror creates snapshots on the source volume, which are copied to the mirror. The snapshot name is in the following format:

filer_volume.number

filer is the host name of the destination filer.

volume is the name of the destination volume (the mirror).

number is the number for the snapshot starting at 1.

SnapMirror automatically deletes old snapshots that are no longer necessary for data replication.

Example

The following example describes the snapshots that are created on the source volume and copied to the mirror.

In this example, data is replicated from */vol/vol1*of filerA to */vol/vol2* of filerB.

To create a baseline version of a mirror, filerA creates a snapshot file named */vol/vol1/filerB_vol2.1* on filerA. All snapshots in */vol/vol1* of filerA, including */vol/vol1/filerB_vol2.1*, are transferred to */vol/vol2* of filerB.

For example, the `snap list` command on filerB generates the following display after */vol/vol1/filerB_vol2.1* is transferred from filerA to filerB.

```
Volume vol2
working.....

%/used    %/total  date          name
-------- --------  ------------  --------
0% ( 0%) 0% ( 0%)  Nov 17 10:50  filerB_vol2.1
1% ( 0%) 0% ( 0%)  Nov 17 10:00  hourly.0
1% ( 0%) 0% ( 0%)  Nov 17 00:00  nightly.0
1% ( 0%) 0% ( 0%)  Nov 15 16:00  hourly.1
1% ( 0%) 1% ( 0%)  Nov 15 15:00  hourly.2
2% ( 0%) 1% ( 0%)  Nov 15 14:00  hourly.3
2% ( 0%) 1% ( 0%)  Nov 15 13:00  hourly.4
2% ( 0%) 1% ( 0%)  Nov 15 12:00  hourly.5
```

When it is time to update the mirror, another snapshot is created on filerA. The `snap list` command on filerA generates the following display after */vol/vol1/filerB_vol2.2* is created on filerA:

```
Volume vol1
working....
%/used   %/total     date          name
-------- ----------  ------------  --------
0% ( 0%) 0% ( 0%)    Nov 17 10:52  filerB_vol2.2 (busy)
0% ( 0%) 0% ( 0%)    Nov 17 10:51  filerB_vol2.1
1% ( 0%) 0% ( 0%)    Nov 17 10:00  hourly.0
1% ( 0%) 0% ( 0%)    Nov 17 00:00  nightly.0
1% ( 0%) 0% ( 0%)    Nov 15 16:00  hourly.1
1% ( 0%) 1% ( 0%)    Nov 15 15:00  hourly.2
```

After *filerB_vol2.2* is transferred from filerA to filerB, both the *filerB_vol2.1* and *filerB_vol2.2* snapshots exist on filerB.

On filerA, however, *filerB_vol2.1* is no longer needed and is deleted; only *filerB_vol2.2* remains.

Consequences of deleting a required snapshot

The filer fails to replicate data if it cannot find the required snapshots in the source and destination volumes. This section provides an example illustrating what happens if you inadvertently delete a snapshot required for data replication.

In this example, data is replicated from */vol/vol1* of filerA to */vol/vol2* of filerB. Suppose you use the `snap delete` command on filerA to delete the *`filerB_vol2.2`* snapshot. When filerB tried to update the mirror according to the schedule in the *snapmirror.conf* file, it could not find the snapshot to determine what incremental changes are required in the mirror.

As a result, filerB displays the following error message.

```
The source filer requires a complete transfer. The destination
volume is online and must be offline for a complete transfer.
```

FilerA displays the following error message:

```
Destination filerB could not accept a complete transfer.
```

These messages mean that filerB tries to create a baseline version of the mirror, but doing so requires that the mirror on filerB be taken off-line.

To proceed, use the `vol offline` command to take the mirror off-line and reboot the filer. At the next scheduled mirror update, the filer creates a baseline version of the mirror.

How SnapMirror works with quotas

Quotas on the mirror

Quotas are always disabled on a mirror, regardless of whether quotas are enabled on the source volume. If you try to enable quotas on a mirror, the filer generates an error message.

How to apply the same quota restrictions on the former mirror

After you convert a mirror to a regular volume, you can enable quotas on it.

However, if the former mirror and the source volume reside on different filers, they might be under the effect of different quota restrictions because there is one */etc/quotas* file for each filer. If you want the same quota restrictions to be applied on both volumes, make sure that the filer on which the former mirror resides has a */etc/quotas* file containing all entries from the */etc/quotas* file used by the source volume.

Refer to "Converting a mirror to a regular volume" on page 540 for information about changing the status of a mirror.

How SnapMirror works with the dump command

How to back up data in the mirror

If you want to back up data from a mirror, the data must be in an existing snapshot, not the active file system, on the mirror. This is because before the `dump` command writes data from an active file system to tape, the filer must create a snapshot in the volume containing the data. The filer cannot create snapshots on the mirror, which is a read-only volume, to be used by the `dump` command.

You can back up any snapshot displayed by the `snap list` command on the mirror. You can also create a snapshot on the source volume, replicate the snapshot to the mirror, and use the `dump` command to back up this snapshot from the mirror to tape.

Effect of the dump command on the mirror update schedule

Running the `dump` command on a mirror might change the mirror update schedule in the following ways:

- If the mirror is scheduled to be updated when a `dump` command is in progress on the mirror, the filer delays replicating data from the source volume until the `dump` command is finished.
- If you start the `dump` command when the filer is replicating data to the mirror, the filer stops the replication and does not update the mirror. It restarts replication after the `dump` command is finished.

Delaying the mirror update is necessary because when the `dump` command backs up data from a snapshot, the snapshot must exist until the command is completed. During a mirror update, the filer deletes existing snapshots on the mirror and copies new ones from the source. This process causes the snapshot being backed up to disappear and the `dump` command to terminate.

The /etc/snapmirror.allow file

Purpose of the snapmirror.allow file

The */etc/snapmirror.allow* file on the source filer specifies the host names of filers that are allowed to replicate data from the source filer.

The filer is not shipped with a default */etc/snapmirror.allow* file. You must use a text editor to create the file if you want to use SnapMirror.

When you can modify the snapmirror.allow file

You can modify the */etc/snapmirror.allow* file at any time.

Format of the snapmirror.allow file

Each entry in the */etc/snapmirror.allow* file contains the host name of the filer that can replicate data from the source filer. Type each entry on a separate line.

Example: If you want to replicate volumes locally on filerA, enter this line in the */etc/snapmirror.allow* file on filerA:

```
filerA
```

Example: If you want to replicate volumes from filerA to filerB, enter this line in the */etc/snapmirror.allow* file on filerA:

```
filerB
```

The /etc/snapmirror.conf file

Purpose of the snapmirror.conf file

The */etc/snapmirror.conf* file resides on the destination filer. It controls where data is copied and how often a mirror is updated.

The filer is not shipped with a default */etc/snapmirror.conf* file. You must use a text editor to create the file if you want to use SnapMirror.

When you can modify the snapmirror.conf file

You can modify the */etc/snapmirror.conf* line at any time.

Format of the snapmirror.conf file

Each entry of the */etc/snapmirror.conf* file is in the following format:

source_filer:source_vol destination_filer:destination_vol argument schedule

Meaning of each field in a snapmirror.conf entry

The following list describes the meaning of each field in a *snapmirror.conf* entry:

- *source_filer* is the host name of the filer from which data is replicated.
- *source_vol* is the source volume name. For example, for the *vol1* volume, type **`vol1`**. Do not type the full path name (*/vol/vol1*).
- *destination_filer* is the host name of the filer to which data is replicated.
- *destination_vol* is the destination volume (mirror) name. For example, for the *vol1* volume, type **`vol1`**. Do not type the full path name (*/vol/vol1*).
- *argument* is kbs (kilobytes per second), the maximum speed at which data is transferred. Enter a value greater than or equal to 11. By default, the filer transfers the data as fast as it can. Enter a dash (-) to indicate that you want to use the default value for *argument*.

Note

The actual data transfer speed might be limited by factors such as network bandwidth. For example, if you specify a large value such as 1,000,000, the filer might still transfer the data at 2,000 kilobytes per second.

- *schedule* is the schedule used by the destination filer for updating the mirror.

You must not leave any field in a *snapmirror.conf* entry blank. If you do not want to specify the maximum speed for data transfer, specify a dash as the *argument*.

Rules for specifying the update schedule

The schedule in each */etc/snapmirror.conf* entry contains four fields:

- minute
- hour
- day of month
- day of week

The fields are separated from each other by a space. If a field contains more than one value, the values are separated from each other by a comma. A field containing an asterisk (*) means that the field is irrelevant. If you specify an asterisk in each field of the schedule, the filer updates the mirror every minute.

The update schedule is mandatory. The filer generates an error message if a */etc/snapmirror.conf* entry does not contain a schedule.

Example: **`filerA:vol1 filerB:vol2 kbs=2000 45 10,11,12,13,14,15,16 * 1,2,3,4,5`**

In this example, */vol/vol1* on filerA is replicated to */vol/vol2* on filerB. Data is replicated at a maximum rate up to 2,000 kilobytes per second. FilerB updates */vol/vol2* at 10:45 a.m., 11:45 a.m., 12:45 p.m., 1:45 p.m., 2:45 p.m., 3:45 p.m., and 4:45 p.m., Monday through Friday. The asterisk in this example means that the mirror update schedule is not affected by the day of month.

When changes to snapmirror.conf take effect

If SnapMirror is enabled, the changes take effect within two minutes. If SnapMirror is disabled, the changes take effect immediately after you enter the `vol snapmirror on` command to enable the feature.

Recommendation

You can create a single */etc/snapmirror.conf* file for your site and copy it to all the filers that use SnapMirror. The */etc/snapmirror.conf* file can contain entries pertaining to other filers. For example, the */etc/snapmirror.conf* file on filerB can contain an entry for replicating a volume from filerC to filerD. When filerB reads the */etc/snapmirror.conf* file, it simply ignores this entry.

Replicating a volume

Description

Replicate a volume if you are in any situation described in "Purposes of SnapMirror" on page 518. You can replicate a volume at any time. After you replicate a volume, you have a mirror that contains the same data as the source volume.

Prerequisites

The following prerequisites must be met before you can replicate a volume:

- You must purchase the SnapMirror license. If the source volume and the mirror are on different filers, you must purchase a license and enter the SnapMirror license code for each filer. Refer to "Enabling services" on page 67 for information about how to enter a license code.
- You must create an off-line volume to be used as the mirror. SnapMirror does not automatically create a volume. Refer to "Creating volumes" on page 89 for information about how to create a volume; refer to "Making a volume inactive" on page 90 for information about how to take a volume off-line.
- The source volume must be on-line. Refer to "Reactivating an off-line volume" on page 90 for information about how to put a volume on-line.
- The source volume must not be a mirror.
- The capacity of the mirror must be greater than or equal to the capacity of the source volume. The configuration of the volumes, however, can be different. For example, the RAID group size can be different for the two volumes. Refer to "Adding disks to a volume" on page 89 for information about how to add disks to a volume.
- The mirror must not be the root volume.

Restrictions

Each filer supports up to four simultaneous volume copy operations. For information about volume copy operations, refer to "Number of volume copy operations SnapMirror generates" on page 520.

Cautions

- SnapMirror creates snapshots in the source volume, which are copied to the mirror. Do not delete these snapshots because incremental changes to the

mirror depend on the snapshots. If the filer cannot find the required snapshot, it cannot perform incremental changes to the mirror.

Refer to "Snapshots created during data replication" on page 522 for an example about how snapshots are created on both the source volume and the mirror.

- The mirror must have snapmirrored status. Do not disable the snapmirrored status using the `volume options` command. Otherwise, the mirror becomes a regular, writable volume. Disable the snapmirrored status only when you no longer need to mirror incremental changes from the source volume.

Recommendations

Follow these recommendations to minimize confusion or to replicate data efficiently:

- When specifying the mirror update schedule in the */etc/snapmirror.conf* file, stagger the update times instead of starting multiple mirror updates at the same time. If the filer does not have enough resources to perform all scheduled mirror updates, it postpones some updates until it has enough resources to do so. As a result, the filer might need to perform subsequent updates at times that are different from those you specify in the */etc/snapmirror.conf* file.
- During data replication, the filer copies data from all snapshots from the source volume. Therefore, the filer must preserve all snapshots when data replication is in progress, even though the snapshot schedule might call for the deletion of some snapshots. If you want the filer to delete snapshots at the exact time specified by the `snap sched` command, schedule the mirror updates to be different from the snapshot deletion times.
- The filer transfers data faster if the source volume and the mirror have the following characteristics:
  - They contain disks of the same size.
  - They contain RAID groups of the same size.
  - They contain the same number of RAID groups.
- If the source volume has quotas enabled and you want to apply the same quota restrictions to the mirror after converting the mirror to a regular volume, make sure that the filer on which the mirror resides contains a copy of the */etc/quotas* file from the source filer. Whenever you make a change to the */etc/quotas* file on the source filer, make the same change to the copy of the */etc/quotas* file on the destination filer.

Steps

Follow these steps to replicate a volume.

Step	Action
1	Add the host name of the destination filer to the */etc/snapmirror.allow* file on the source filer.
2	Edit the */etc/snapmirror.conf* file on the destination filer to specify the volume to be replicated and the schedule at which the mirror is updated.
3	Enter the `vol snapmirror on` command on both the source filer and destination filer to enable SnapMirror. **Result:** The filer reads the */etc/snapmirror.conf* file. If the filer is a destination filer specified in the */etc/snapmirror.conf* file, it establishes a connection with the source filer at the time of the scheduled mirror update. If a baseline version of the mirror does not exist, the filer takes a snapshot of the source volume at the time of the scheduled mirror update and transfers all data in the snapshot from the source volume to the mirror. The filer does not display any messages if it can transfer the data successfully. You can check the data replication status following the instructions in "Checking data replication status" on page 538. The filer makes subsequent updates to the mirror according to the schedule specified in the */etc/snapmirror.conf* file. **Note** The `vol snapmirror on` command does not persist across filer reboots. Put the command in the */etc/rc* file if you want the command to remain in effect after the filer is rebooted. If the */etc/rc* file does not contain a `vol snapmirror` command, data replication is disabled.

Step	Action
4	If the source volume and the mirror reside on different filers, and if the source volume has quotas enabled, copy the */etc/quotas* file from the source filer to a file on the destination filer. Whenever you make a change to the */etc/quotas* file on the source filer, make the same change to the copy on the destination filer so that the destination filer always contains a record of all the quota entries used by the source filer.

Disabling data replication for the entire filer

Description

You can disable data replication for the entire filer if you decide that data replication is no longer necessary. You can disable the feature at any time, even when replication is underway. The mirror remains unchanged after you disable replication.

Steps

Follow these steps to disable volume replication for the entire filer.

Step	Action
1	Enter the `vol snapmirror off` command on the destination filer to disable SnapMirror. **Result:** If the filer is currently transferring data from one volume to another, the transfer stops immediately. The destination volume remains the same as before the transfer. The snapshot taken in the source volume for the data transfer remains, but is deleted and replaced by a new snapshot the next time the mirror is updated. The filer stops reading the */etc/snapmirror.conf* file every minute. Entering the `vol snapmirror off` command on the destination filer does not affect SnapMirror on the source filer. Other filers can continue to replicate data from the source filer.
2	If the `vol snapmirror on` command is in the */etc/rc* file, remove the command. **Note** There is no need to enter the `vol snapmirror off` command in the */etc/rc* file. By default, data replication is disabled.

Resuming data replication for the entire filer

Description

After you disable data replication, you can resume it at any time. When data replication resumes, the filer copies the data from the source volume to the mirror that has changed since the last update.

Prerequisites

In addition to the requirements described in "Prerequisites" on page 530, you must meet these requirements before you can resume data replication:

- The mirror must have snapmirrored status. If you converted the mirror to a regular volume, the filer cannot resume data replicating on the regular volume. If you want to use the regular volume as a mirror again, follow the procedure in "Replicating a volume" on page 530 so that the filer can create a baseline version of the mirror. Do not try to assign the snapmirrored status to the volume, because you cannot set the `snapmirrored` volume option to On.
- If the mirror never lost the snapmirrored status but you took the mirror off-line, you must put it back on-line.

Step

Follow this step to resume data replication for the entire filer.

Step	Action
1	Enter the `vol snapmirror on` command on the destination filer to enable SnapMirror. **Result:** The filer reads the */etc/snapmirror.conf* file to determine whether it needs to create a baseline version of a mirror or to update a mirror.

Disabling data replication for one volume

Description

You can stop data replication for a particular volume if you decide that there is no need to update the mirror. For example, you might want to change the mirror to a regular volume.

You can disable data replication at any time, even when data transfer is underway. The destination volume remains the same as before the transfer. The snapshot taken in the source volume for the data transfer remains, but is deleted and replaced by a new snapshot the next time the mirror is updated.

Steps to disable data replication for one volume

Follow these steps to disable data replication for one volume. These steps disable volume replication until you reenable it.

Step	Action
1	Comment out the entry in the */etc/snapmirror.conf* file by preceding the entry with a pound sign (#).
2	Enter the `vol snapmirror on` command to make the filer reread the */etc/snapmirror.conf* file. **Result:** If the filer is currently transferring data from that volume, the transfer stops immediately.

Steps to disable data replication while data transfer is in progress

Follow these steps to disable data replication while data transfer is in progress. These steps stop the replication temporarily; at the time of the next scheduled update, data replication starts up again.

Step	Action
1	Enter the `vol copy status` command to display the operation numbers of the volume copy. **Result:** Because SnapMirror uses the `vol copy` command to transfer data, the `vol copy status` command displays all volume copy operations, including the one for replicating a volume.

2	Enter the `vol copy abort` command to terminate the volume copy operation.
	
	Result: The filer stops copying data and displays messages similar to these:
	
	`image operation 0 is now being aborted`
	`image operation 1 is now being aborted`

Checking data replication status

Description

If you want to know how much data has been copied to the mirror or whether a filer is using resources for replicating data, you can use the `vol snapmirror status` command at any time to check the status of data replication.

Prerequisite

SnapMirror must be enabled before you can check the status of data replication. Otherwise, the following error message is displayed:

```
Snapmirror is turned off.
```

Step

Follow this step to check the status of data replication.

Step	Action
1	Enter the `vol snapmirror status` command. **Result:** The filer displays a message showing whether a transfer is in progress or how much data replication has been completed.

Examples

The following examples describe how the `vol snapmirror status` command displays the status of data replicating.

When no data replication is in progress: The following display shows that currently no data is being copied from the *vol0* volume to the *vol1* volume on filerA:

```
Source        Dest            Status
filerA:vol0  filerA:vol1      Idle
```

When data replicating is in progress: The following display shows that the filer has just begun transferring data from the *vol0* volume to the *vol1* volume on filerA:

```
Source        Dest            Status
filerA:vol0  filerA:vol1    Transferring
```

The following display shows that the filer finished transferring 26% of the data from the *vol0* volume to the *vol1* volume on filerA:

```
Source        Dest            Status
filerA:vol0  filerA:vol1    Transferring (26% complete)
```

Converting a mirror to a regular volume

Description

If you use SnapMirror for data migration, after you synchronize the data between the source volume and the mirror, convert the mirror to a regular volume. If you use SnapMirror for disaster recovery, convert the mirror to a regular volume after the source volume becomes unavailable. After the conversion, you can export the volume or create a share for the volume so that network users can write to it in the same way as they did to the source volume.

Prerequisite

You must meet this prerequisite if the source volume and the mirror reside on different filers and you want the same quota restrictions to be applied after converting the mirror to a regular volume:

The destination filer must have a */etc/quotas* file that includes all entries from the */etc/quotas* file used by the source filer. If SnapMirror is used for data migration, you can copy the */etc/quotas* entries from the source filer to the */etc/quotas* file of the destination filer before you convert the mirror to a regular volume. However, if SnapMirror is used for disaster recovery, on the destination filer, keep a copy of all */etc/quotas* entries used by the source filer at all times so that you can apply the entries to the destination volume when the source filer becomes unavailable.

Steps

Follow these steps to convert a mirror to a regular volume.

Step	Action
1	If you want to enable quotas after converting the mirror to a regular volume, go to Step 2. Otherwise, go to Step 4.
2	Edit the */etc/quotas* file on the destination filer so that after the conversion, the former mirror has the same quota restrictions as the source volume. If the source volume uses per-volume quotas, replace the source volume name with the mirror name in the quota entries.

Step	Action
3	Enter the following command to convert the mirror to a regular volume: `vol options` *`volume_name`* `snapmirrored off` If you want to enable quotas on the former mirror, go to Step 4. Otherwise, the procedure is complete.
4	Enter the following command to enable quotas on the former mirror: `quota on` *`volume_name`*

Differences between the vol copy command and SnapMirror

Differences

This following table describes the differences between the `vol copy` command and SnapMirror.

The vol copy command	SnapMirror
It is a standard Data ONTAP feature that requires no license codes.	It requires a license code.
The result of copying a volume is an off-line volume.	The result of replicating a volume is a mirror, which is an on-line, read-only volume.
It does not copy incremental changes from the source volume.	It periodically copies incremental changes from the source volume.
It requires that you enter the destination filer name in the */etc/hosts.equiv* file.	It requires that you enter the destination filer name in the */etc/snapmirror.allow* file.
It does not involve a configuration file that controls how a volume is copied.	It requires that you specify an entry in the */etc/snapmirror.conf* file to control the volume to be replicated, the data transfer rate, and the mirror update schedule.

System Information and Performance 20

Chapter contents

This chapter discusses the topics described in the following table.

Read...	To learn about...
"Displaying the Data ONTAP version" on page 544	How to use the `version` command to display the Data ONTAP version.
"Displaying filer configuration information" on page 545	How to use the `sysconfig` command to display filer configuration information.
"Displaying volume information" on page 547	How to use the `vol` command to display information about a volume's configuration.
"Displaying filer statistics" on page 548	How to use the `sysstat` and `uptime` commands to display filer statistics.
"Displaying network statistics" on page 550	How to use the `netstat` command to display network statistics.
"Displaying interface statistics" on page 551	How to use the `ifstat` command to display interface statistics.
"Improving filer performance" on page 561	How to improve filer performance by adding hardware or setting software options.

Displaying the Data ONTAP version

How to display the Data ONTAP version

To display the version of Data ONTAP currently running on a filer, use the `version` command. The display shows the version number and the date of the version, as follows:

```
version

NetApp Release 5.1: Fri May 12 03:06:00 PDT 1998
```

Displaying filer configuration information

Use the sysconfig command

The `sysconfig` command displays information about the filer's hardware configuration. The exact types of information displayed depend on the command options.

Displaying disk information using sysconfig -d

The `sysconfig -d` command displays product information about each disk in the filer.

Displaying RAID information using sysconfig -r

The `sysconfig -r` command displays RAID configuration information about the parity disk, data disks, and hot spare disks, if any. This information is useful for the following purposes:

- Locating a disk referenced in a screen message. Refer to "Using disks of various sizes" on page 75 for more information about disk identifiers.
- Determining how much space on each disk is available to the filer. Refer to "Understanding usable space on each disk" on page 76 for more information about disk capacity.
- Determining the status of the disk operations, such as RAID scrubbing, reconstruction, parity verification, adding a hot spare, and disk failure.

You can obtain the information displayed by `sysconfig -r` from SNMP, using the Network Appliance custom MIB. For information about SNMP, see "Using SNMP" on page 111.

Displaying tape drive information using sysconfig -t

The `sysconfig -t` command displays device and configuration information for each tape drive on the system. Use this command to determine the capacity of the tape drive and the device name before you use the `dump` and `restore` commands.

Displaying overall filer information using sysconfig -v

The `sysconfig -v` command displays the system's RAM size, NVRAM size, and information about devices in all expansion slots. This information varies according to the devices on the filer. You can specify a slot number to display information about a particular slot. Slot numbers start at 0, where slot 0 is the system board.

Displaying overall filer information using sysconfig

If you enter `sysconfig` without any options, information similar to what you get with `sysconfig -v` is displayed, but the information is abbreviated. When you report a problem to Network Appliance, provide the information displayed by `sysconfig -v`. This information is useful for diagnosing system problems.

Displaying volume information

Use the vol status command

The `vol status` command displays information about a volumes's configuration. The types of information displayed depend on the command options. When you specify a volume, the information for that volume is displayed; when you do not specify a volume, the status of all volumes in the filer is displayed.

Displaying volume state information with vol status

With no options, the `vol status` command displays a one-line synopsis of volume states. This includes the volume name, whether it is on-line or off-line, and other states, for example, partial, degraded, and so on.

Displaying disk information using vol status -d

The `vol status -d` command displays information about disks. The disk information is the same as the information from the `sysconfig -d` command.

Displaying RAID information using vol status -r

The `vol status -r` command displays a list of the RAID information. The display is the same as the `sysconfig -r` display.

Displaying RAID information for each group using vol status -v

The `vol status` -v command displays information about each RAID group.

Displaying filer statistics

Use the sysstat and uptime commands

You use the `sysstat` and `uptime` commands to display filer statistics.

About the sysstat command

The `sysstat` command displays information about CPU utilization, file operations, read and write operations on disks and tapes, and the age of data in the cache buffer. By default, the filer displays statistics every 15 seconds; you can specify the interval, in seconds, at which statistics are displayed.

The `sysstat` output is particularly useful for revealing file access patterns on your filer, from which you can determine whether you should install more NVRAM, system memory, or disks.

Depending on the applications that access the filer, the computers on the network, and the network configuration, the following suggested actions might improve your filer's performance:

- Typically, 2 MB of NVRAM is sufficient for a system with one SCSI channel. However, if the Disk Kb/s write column shows that the filer writes to disks almost every second, adding NVRAM might help improve system performance. If there are short spurts of writes, adding NVRAM does not improve system throughput. See the hardware guide for your filer model for information about how much additional NVRAM you can install.
- The data in the Cache age column indicates how fast read operations are cycling through system memory. If the values in the Cache age column are consistently below 5, the filer might benefit from more system memory. Cache age shows the age of the oldest read-only blocks in memory. A low cache age means that the filer is retrieving information from disk instead of from memory. For information about adding memory, see "About the uptime command" on page 549. If the CPU utilization percentage in the CPU column is low but you are not getting the expected performance, you might need to add disks to the system.
- If CPU utilization is high, you might need to add another filer.
- If the network traffic shown in the Net Kb/s column is at or near the capacity of the network interfaces (1 MB per Ethernet adapter and 10 MB per FDDI adapter), you might need to add another network adapter card to the system.

About the uptime command

The `uptime` command prints the current time, the length of time the system has been up, and the total number of NFS operations the system has performed since it was last booted.

Example: An example of the display is

```
uptime

8:54am up 2 days 22:23, 3122520 NFS ops
```

Displaying network statistics

Use the netstat command

You use the `netstat` command to display network statistics.

About the netstat command

The `netstat` command displays network-related data in various output formats.

The `netstat -i` and `netstat -I` options show the state of all network interfaces or one specific interface, respectively. The `netstat -r` command shows the filer's routing table.

For information about troubleshooting network problems, see "Network problems" on page 587. For more information about the `netstat` command, see the *na_netstat*(1) man page.

Displaying interface statistics

Use the ifstat command

The `ifstat` command prints per-interface statistics not reported by commands such as `netstat`. This includes some statistics maintained by the networking code, as well as statistics maintained by the driver and by the networking card.

The output of the `ifstat` command might contain many fields, because different types of interfaces, for example, Ethernet, Gigabit (GB) Ethernet, FDDI, and ATM, have different statistics.

ifstat syntax: The syntax for the `ifstat` command is

```
ifstat [-z]  -a | interface
```

The `-z` option "zeros" (or clears) the statistics. The `-a` option lists statistics for all the filer's interfaces. The interface option indicates the type of interface for which you want statistics.

Explanation of interface statistics

Ethernet: The following table describes the statistics in the RECEIVE section of the `ifstat` command output when you use the command on an Ethernet interface.

Statistic	Meaning
`Packets`	Number of packets received on the interface.
`Bytes`	Number of bytes received on the interface.
`Errors`	Number of errors during Ethernet frame reception, including all kinds of receive errors.
`No buffers`	Number of received packets dropped due to the unavailability of buffers.
`Length err`	Number of frames truncated due to the shortage of receive descriptors.
`Runt frames`	Number of runt frames.

Statistic	Meaning
`Long frames`	Number of frames received that exceeded the maximum Ethernet-specified size of 1,518 bytes.
`CRC error`	Number of CRC errors that occurred on the received frames.
`H/w overflow`	Number of frames discarded because of receive FIFO overflow.
`Process stop`	Number of times the receive process stopped.
`List overflow`	Number of frames dropped due to the unavailability of descriptors.
`Process reset`	Number of times the receive process was reset.
`Rst frame drops`	Number of frames discarded due to the resetting of the receive process.

The following table describes the statistics in the TRANSMIT section of the `ifstat` command output when you use the command on an Ethernet interface.

Statistic	Meaning
`Packets`	Number of packets attempted to be transmitted.
`Bytes`	Number of bytes attempted to be transmitted.
`Errors`	Number of hardware errors encountered while attempting to transmit.
`Collisions`	Number of collisions that occurred while transmitting frames.
`Late collisions`	Number of collisions aborted due to a late collision.
`Excess coll`	Number of times transmission was aborted due to excessive collisions.
`Queue full`	Number of times the queue was full.
`List full`	Number of frames that were dropped due to the unavailability of descriptors.

Statistic	Meaning
`No carrier`	Number of times the carrier signal was not present during transmission.
`Underflow`	Number of times the transmitter aborted the message because data arrived late from memory.
`Defer`	Number of times transmission had to be deferred.
`Time out`	Number of times the transmit jabber timer expired.
`Stopped`	Number of times the receive process stopped.
`List underflow`	Number of frames that had to be dropped due to the unavailability of descriptors.
`Loss of carrier`	Number of times the carrier was lost.
`No buffers`	Number of times buffers were unavailable.
`Requeue`	Number of times the frame was requeued due to list underflow.
`Threshold up`	Number of times the transmit threshold was increased.
`Threshold dn`	Number of times the transmit threshold was decreased.
`Threshold lvl`	Current threshold level.

The following table describes the statistics in the DEVICE section of the `ifstat` command output when you use the command on an Ethernet interface.

Statistic	Meaning
`Interrupts`	Number of times the Ethernet device interrupted the host.
`Resets`	Number of times the Ethernet device was reset.

The following table describes the statistics in the LINK INFO section of the `ifstat` command output when you use the command on an Ethernet interface.

Statistic	Meaning
`Auto`	Auto-Negotiation state.
`Mediatype`	Media type, such as twisted pair.
`Link Partner`	The Auto-Negotiation capability of the remote end. It is unknown if Auto-Negotiation is disabled.
`Link State`	Link status.

GB Ethernet: The following table describes the statistics in the RECEIVE section of the `ifstat` command output when you use the command on a GB Ethernet interface.

Statistic	Meaning
`Packets`	Number of packets received on the interface.
`Bytes`	Number of bytes received on the interface.
`Errors`	Number of errors during Ethernet frame reception, including all kinds of receive errors.
`Queue full`	Number of received packets dropped due to the unavailability of buffers.
`Unicast packets`	Number of unicast packets received.

The following table describes the statistics in the TRANSMIT section of the `ifstat` command output when you use the command on a GB Ethernet interface.

Statistic	Meaning
`Packets`	Number of packets attempted to be transmitted.
`Bytes`	Number of bytes attempted to be transmitted.
`Errors`	Number of hardware errors encountered while attempting to transmit.

Statistic	Meaning
`Collisions`	Number of collisions that occurred while transmitting frames.
`Unicast packets`	Number of unicast packets transmitted.

The following table describes the statistics in the DEVICE section of the `ifstat` command output when you use the command on a GB Ethernet interface.

Statistic	Meaning
`Received errors`	Number of errors encountered during reception by the interface.
`Transmit errors`	Number of errors encountered during transmission by the interface.
`Collisions`	Number of collisions encountered during transmission by the interface.

FDDI: The following table describes the statistics in the RECEIVE section of the `ifstat` command output when you use the command on an FDDI interface.

Statistic	Meaning
`Packets`	Number of packets received on the interface.
`Bytes`	Number of bytes received on the interface.
`Errors`	Number of errors during Ethernet frame reception, including all kinds of receive errors.
`Queue full`	Number of received packets dropped due to the unavailability of buffers.
`H/w faults`	Number of hardware faults.
`SA matches`	Number of SA matches.
`DA nonmatches`	Number of DA non-matches.

Statistic	Meaning
`PDQ aborts`	Number of PDQ aborts.
`Invalid lengths`	Number of invalid lengths.
`Fragments`	Number of fragments.
`Format errs`	Number of frames received that contained an odd number of symbols; that is, the number of non-integral octets.
`MAC resets`	Number of MAC resets.
`CRC Errors`	Number of frames received with a bad frame check sequence.
`Status Errors`	Number of frames received that had the E (error detected) indicator set or a missing E indicator.
`Long Frames`	Number of frames received that were larger than the maximum legal size.
`Short Frames`	Number of short frames received.
`Copied buffers`	Number of copied buffers.
`No buffers`	Number of times there were insufficient buffers.
`LLC errors`	Number of LLC errors.
`Blk check errs`	Number of block check errors.
`Length errors`	Number of length errors.
`Overruns`	Number of times that the interface received a frame loss indication from the hardware on the interface.

The following table describes the statistics in the TRANSMIT section of the `ifstat` command output when you use the command on an FDDI interface.

Statistic	Meaning
`Packets`	Number of packets attempted to be transmitted.
`Bytes`	Number of bytes attempted to be transmitted.

Statistic	Meaning
`Errors`	Number of hardware errors encountered while attempting to transmit.
`Queue full`	Number of received packets dropped due to unavailability of buffers.
`No buffers`	Number of times buffers were not available.
`No buf resets`	Number of times the process was reset due to the unavailability of buffers.
`No descriptors`	Number of times descriptors were unavailable.
`I/f down`	Number of times the interface was down.
`Underruns`	Number of times an underrun occurred because the interface's transmit FIFO became empty due to insufficient memory. That is, the device tried to transmit but could not get data into memory fast enough.
`Failures`	Number of frames that were transmitted with an error.

The following table describes the statistics in the DEVICE section of the `ifstat` command output when you use the command on an FDDI interface.

Statistic	Meaning
`Interrupts`	Number of times the FDDI device interrupted the host.
`Spurious ints`	Number of spurious interrupts received from the device.
`Resets`	Number of times the device was reset due to unrecoverable hardware faults.
`Non-existing mem`	Number of times the device was asked to address nonexisting memory. This might indicate a software or hardware problem.
`Parity errs`	Number of parity errors detected on the interface.

Statistic	Meaning
`Host parity errs`	Number of parity errors detected on the host.
`Spurious Unsols`	Number of unknown unsolicited error events reported by this link entity.
`Ring inits in`	Number of ring claim processes or ring beacon processes initiated by this link entity.
`Ring beacon out`	Number of ring beacon processes initiated by this link entity.
`Ring beacon in`	Number of ring beacon processes recognized by this link entity.
`Dupl. address`	Number of times the duplicate address test failed.
`Dupl. token`	Number of times this link entity detected a duplicate token.
`Ring purge errs`	Number of times the MAC PurgeError Event bit was set.
`FCI strip errs`	Number of times a Frame Content Independent Strip operation was terminated.
`Trace init`	Number of PC Traces initiated by this station.
`Selftest failures`	Number of failed self- tests during device resets.
`PC trace rcvd`	Number of PC traces received by this station.
`Config changes`	Number of config changes.
`LEM rejects`	Number of LEM rejects.
`Elastic buf errs`	Number of ELM Chip EBUFF_ERROR events.
`Link test rejects`	Number of link test rejects.

ATM: The following table describes the statistics in the RECEIVE section of the `ifstat` command output when you use the command on an ATM interface.

Statistic	Meaning
`Packets`	Number of packets received on the interface.
`Bytes`	Number of bytes received on the interface.
`Errors`	Number of errors during Ethernet frame reception, including all kinds of receive errors.
`Vpi out of range`	Number of packets received on an out-of-range VPI.
`Vci out of range`	Number of packets received on an out-of-range VCI.
`Packets dropped`	Number of received packets dropped because of errors reported by the adapter.
`No connection`	Number of packets received for a disconnected or nonexistent virtual circuit.
`Small buffers`	Number of small received buffers supplied to the adapter.
`Large buffers`	Number of large received buffers supplied to the adapter.

The following table describes the statistics in the TRANSMIT section of the `ifstat` command output when you use the command on an ATM interface.

Statistic	Meaning
`Packets`	Number of packets attempted to be transmitted.
`Bytes`	Number of bytes attempted to be transmitted.
`Errors`	Number of hardware errors encountered while attempting to transmit.
`Queue full`	Number of packets dropped because they could not be put in the transmit queue.
`List full`	Number of packets that could not be transmitted because the transmit queue was full.
`Iface unavailable`	Number of packets that could not be transmitted because the referenced interface was down.

For more information about the `ifstat` command, see the *na_ifstat*(1) man page.

Improving filer performance

About this section

This section describes configuration procedures that might improve your filer's performance.

Limiting directory file size

An extremely large directory file can use up most of the filer's CPU cycles when a user enters an `ls` command in the directory. The limit on directory file size is set by the `options wafl.maxdirsize` command. The default limit, 10 MB, should prevent the system from hanging. A directory of this size accommodates approximately 300,000 files with short file names.

The `wafl.maxdirsize` option takes the maximum number of kilobytes as its argument. When you reset the maximum directory size, the argument that you supply is rounded up to the next highest 4K boundary. If a user tries to create a file that would cause the directory to grow larger than the specified size limit, the user's command fails.

Balancing NFS traffic on network interfaces

To balance network traffic among different interfaces, attach multiple interfaces on the filer to the same physical network. For example, if two Ethernet interfaces on the filer named toaster are attached to the same network where four clients reside, specify in */etc/fstab* on client1 and client2 that these clients mount from *toaster-0:/home*. Specify in */etc/fstab* on client3 and client4 that these clients mount from *toaster-1:/home*. This scheme can balance the traffic among interfaces if all clients generate about the same amount of traffic.

The filer always responds to an NFS request by sending its reply on the interface on which the request was received.

Avoiding access time update for inodes

If your applications do not depend on having the correct access time for files, you can disable the update of access time (`atime`) on an inode when a file is read. To prevent updates, turn the `no_atime_update` option On. Consider turning this option On if your filer has extremely high read traffic, for example, on a news server used by an Internet provider, because it prevents inode updates from contending with reads from other files.

Caution

If you are not sure whether your filer should maintain an accurate access time on inodes, leave this option at its default, Off, so that the access time is updated.

Improving performance on directory lookups

Turning the `nfs.big_endianize_fileid` option On improves performance on directory lookups for clients that use the file ID in the file handle as a hash key in certain ways. Enable this option if your NFS clients are mainly running HP-UX or IRIX.

Note

If you turn the `big_endianize_fileid` option On, all NFS clients that have mounted directories from the filer must unmount and remount them; otherwise, they get “stale file handle” errors on all references to files already opened on the filer until they unmount and remount all directories.

Improving read-ahead performance

If the file access patterns of your clients are random (nonsequential), turning minimal read-ahead On might improve performance. By default, the filer uses aggressive read-ahead, which enhances sequential access, which is more commonly used by UNIX clients and applications. To specify minimal read-ahead, turn the `minra` option On. By default, the option is Off and the filer does very aggressive read-ahead.

Troubleshooting 21

Troubleshooting topics

Use the information in the following table to determine which troubleshooting topics to read.

Read...	To learn about...
"Getting technical assistance" on page 566	Getting technical assistance.
"Booting from system boot diskette" on page 567	Booting from diskette to recover from a lost password, configuration file errors, and some disk problems.
"Restarting a shut down filer" on page 570	Restarting a filer when it has shut down automatically.
"NVRAM problems" on page 571	How the filer handles NVRAM low battery conditions and inconsistent NVRAM contents.
"Volume problems" on page 573	Types of volume problems.
"Failed mounts and stale file handles" on page 574	How to fix mount and stale file handle problems caused by changing a volume's name.
"Volume name problems" on page 575	Using the rules for naming volumes, and about volume naming error messages.
"Disk problems" on page 576	Types of disk problems.
"Disk failure without a hot spare disk" on page 577	How the filer handles a disk failure when a hot spare disk is not available.
"Disk failure with a hot spare disk" on page 578	How the filer handles a disk failure when a hot spare disk is available.
"Disk errors" on page 579	Error messages the filer displays when it discovers certain disk problems.

Read...	To learn about...
“Inconsistent file system” on page 581	What to do if the filer displays a message that a file system is inconsistent.
“Disk operations in Maintenance mode” on page 582	Using maintenance-mode operations to correct disk problems.
“Configuration problems” on page 583	Booting the filer and recovering from errors in configuration files.
“How to reset the filer password” on page 585	Resetting the administration password for the filer.
“How to initialize all disks and create a new file system” on page 586	Initializing all the filer disks because you want to redeploy the filer or because Network Appliance Technical Support instructs you to do so.
“Network problems” on page 587	Diagnosing network problems.
“NFS problems” on page 589	Diagnosing NFS problems.
“Windows access problems” on page 591	Diagnosing problems with accessing the filer from Windows.
“UNIX cpio problems” on page 595	How 16-bit versions of the `cpio` command might cause problems when copying large files.
“UNIX df problems” on page 596	How the `df` command on some UNIX systems might report free space incorrectly.
“Cluster messages” on page 597	What you can do to fix cluster problems.
“Serious error messages” on page 603	What to do if the filer halts with a serious error message.

Related information

Hardware problems are described in the hardware guide for your filer model. See the appropriate hardware guide for information about system LEDs; LCD error messages, if your filer has an LCD, and system start-up messages.

Getting technical assistance

Information to note before calling for support

If you encounter problems with the filer that you cannot solve, you might need to contact your service provider or Network Appliance Technical Support for assistance. If you do, have the following information available:

- the system serial number
- the system configuration as reported by the `sysconfig -v` command (if the filer still responds to the command)
- any diagnostic messages that were reported (Diagnostic messages are in the */etc/messages* file on the root volume.)

How to contact Network Appliance

The following table provides telephone and FAX numbers, and a URL that you can use to contact Network Appliance Technical Support.

Telephone	FAX	Web
For U.S. customers: 888 463-8277 For international customers: +1 408 367-3000 (ask for Technical Support)	(408) 367-3151	http://www.netapp.com/support/form.html

Booting from system boot diskette

Boot from diskette to correct some types of problems

You might need to reboot the filer from the system boot diskette to correct configuration problems, recover from a lost password, or correct certain disk configuration problems.

Restrictions on booting from diskette

Do not boot the filer with a pre-5.0 system boot diskette. The attempt to boot will fail because pre-5.0 software cannot recognize system data from a filer running Data ONTAP 5.0 or later software.

If your file is using clustering, do not boot with a pre-5.2 system boot diskette. Releases earlier than 5.2 do not support clustering.

If you need to run an earlier release of Data ONTAP on the filer, contact Network Appliance Technical Support, as described in "Getting technical assistance" on page 566, for instructions about how to downgrade your filer.

Procedure for booting from diskette

Complete the steps in the following table to display the boot menu from a diskette. Booting from the diskette takes longer than booting from the hard disk.

Step	Action
1	Insert the diskette labeled System Boot Disk 1 into the filer's diskette drive. **Note** The Data ONTAP 5.2 system boot diskettes are specific to the model of filer that you are updating. You received diskettes appropriate to the model you are updating. If you have different models, be sure to check the diskette labels to make sure you're using the correct diskettes for the filer.
2	From the system console, enter the `reboot` command or, if the system is powered off, power it on. **Result:** The filer begins the boot process.

Step	Action
3	When the filer's LCD prompts you to, remove the diskette and insert the diskette labeled System Boot Disk 2 into the filer's diskette drive.

4	Press the Enter key on your console.
	
	Result: The filer boots and, if it can, displays the boot menu, shown below:
	
	`1) Normal Boot`
	`2) Boot without /etc/rc`
	`3) Change Password`
	`4) Initialize all disks`
	`5) Maintenance mode boot`
	
	`Selection (1-5)?`

Step	Action	
5	Choose one of the boot types shown in the following table by entering the corresponding number.	
	Boot type	**Description**
	Normal Boot (1)	Use Normal Boot to run the filer normally, but from a diskette.
	Boot without */etc/rc* (2)	Use Boot without */etc/rc* to troubleshoot and repair configuration problems. **Note** Booting without */etc/rc* causes the filer to use only default options settings, disregard all options settings you put in */etc/rc*, and disable some services, such as `syslog`.
	Maintenance mode boot (5)	Use Maintenance mode boot to go into Maintenance mode and perform some volume and disk operations and get detailed volume and disk information. Maintenance mode is special for the following reasons: ◆ Most normal functions, including file-system operations, are disabled. ◆ A limited set of commands is available for diagnosing and repairing disk and volume problems. ◆ You exit Maintenance mode with the `halt` command.

Restarting a shut down filer

Procedure for restarting filer after unexpected shutdown

Complete the steps in the following table to restart your filer if it shuts down unexpectedly.

Step	Action	
1	Write down the messages displayed on the console and the message on the LCD, if your filer has an LCD.	
2	**If the console is...**	**Then...**
	Showing the `mon>` prompt	Enter **`b`** (for boot) to reboot the filer.
	Showing the `ok` prompt	Enter **`boot`** to reboot the filer.
	Unresponsive	Reset the filer by turning it off, leaving it off for 30 seconds, and turning it back on.
	Results: The filer boots and displays the system prompt on the console.	
3	If the filer still does not boot, display the boot menu from diskette, as described in "Booting from system boot diskette" on page 567, then choose Normal Boot.	

NVRAM problems

Types of NVRAM problems

This section describes two types of NVRAM problems:

- NVRAM battery failure
- Inconsistent NVRAM contents

What F220, F330, and F540 filers do when NVRAM battery is low

This section applies only to the NetApp F220, NetApp F330, and NetApp F540 filers.

If the filer detects low batteries on the NVRAM card, it gives advance warning and automatically shuts down every 24 hours. You can change the time interval with the `raid.timeout` option to the `options` command.

If the filer is in degraded mode and the NVRAM batteries are low, it fails to boot. Replace the NVRAM batteries immediately.

Very rarely do NVRAM battery failures cause data loss. When the filer executes the `halt` command, it flushes the contents of NVRAM to disk and turns off the NVRAM so that no battery power is used.

How the filer handles inconsistent NVRAM contents

The filer performs a number of checks to ensure that the NVRAM contents are consistent. If the contents are inconsistent, the filer performs one of following actions:

Inconsistency due to improperly updated volume: If the inconsistency is due to a failure to update a volume properly, the filer displays a message suggesting that you halt the filer with the `halt` command, take the offending volume off-line in Maintenance mode, and reboot the filer.

Inconsistency due to log updates for off-line volume: If the inconsistency is due to the log having updates for an off-line volume, the filer asks whether to discard them.

Inconsistency due to other reasons: If there are many inconsistencies that cannot be repaired, the filer discards the inconsistent contents and creates a core dump file. The file requests received during the last few seconds before the filer shuts down are lost. This does not cause the file system to become inconsistent, but files written during the last 10 seconds before shutdown might contain old or

incorrect data. Also, because the parity of some recently written stripes might be incorrect, the filer must do a parity check on the entire RAID array. The parity check, and any correction, is performed on-line; that is, the filer conducts the test in the background while continuing otherwise normal operation.

This kind of data error can happen only when you boot a system after a failure or after you turn off a filer without using the `halt` command.

Volume problems

Types of volume problems described

This section describes the following types of volume problems:

- Failed mounts and stale file handles
- Volume name problems

Failed mounts and stale file handles

Changing volume names can cause mount and file handle problems

If mounts fail and clients see stale file handles, it might be because you renamed a volume and it is not exported anymore. This is because the new volume name was not in the */etc/exports* file on the root volume.

Procedure for fixing the /etc/exports problem

Complete the step in the following table to fix the problem.

Step	Action
1	Edit the */etc/exports* file on the filer default volume to change the old volume name to the new volume name. **Results:** The problem is resolved.

Volume name problems

Volume naming rules

A valid volume name has the following characteristics:

- The prefix is followed by either a letter or an underscore (“_”).
- It contains only letters, digits, and underscores.
- It is not longer than 255 characters.

Examples of volume names

Examples of valid volume names are

- *_tech_pubs*
- *SW_Engineering*
- *Dept_32*

Error messages about volume names

You might get an error message that contains one of the following phrases:

- invalid volume name
- unrecognized volume name
- illegal volume name

If you get one of these messages, take one of the following actions:

- Type correctly the name of an existing volume.
- Type a valid volume name.

Disk problems

Types of disk problems described

This section describes how the filer reacts to a

- Disk failure without a hot spare disk
- Disk failure with a hot spare disk

Disk failure without a hot spare disk

About this section

This section describes how the filer reacts to a disk failure when a hot spare disk is available.

Filer runs in degraded mode

If a disk fails, the filer continues to run without losing any data but has a somewhat degraded performance.

Caution

Replace the disk as soon as possible, because a second disk failure could cause the filer to lose the entire file system.

Filer logs warning messages in /etc/messages

The filer logs a warning message in the */etc/messages* file on the root volume every hour after a disk fails.

Filer shuts down automatically after 24 hours

To ensure that you notice the failure, the filer automatically shuts itself off in 24 hours, by default, or in a period you set with the `raid.timeout` option to the `options` command. You can restart the filer without fixing the disk, but it continues to shut itself off periodically until you repair the problem.

Caution

Check the */etc/messages* file on the root volume once a day for important messages. You can automate checking this file with a script on a remote host that periodically searches the file and then alerts you.

Filer reconstructs data after disk is replaced

After you replace a disk, the filer detects the new disk when the system boots. The filer starts file service and reconstructs the missing data in the background with minimum interruption to service.

Disk failure with a hot spare disk

About this section

This section describes how the filer reacts to a disk failure when a hot spare disk is not available.

Filer replaces disk with spare and reconstructs data

If a disk fails, the filer

- replaces the failed disk with a hot spare disk
- reconstructs the missing data on the hot spare disk in the background, so that the interruption to file service is minimized
- logs the activity in the */etc/messages* file on the root volume

The filer does not shut down automatically.

Caution

After the filer is finished reconstructing data on the hot spare disk, replace the failed disk with a new hot spare disk as soon as possible so that there is always a hot spare disk available in the system. For information about replacing a disk, refer to "Disk management tasks" on page 85.

If a second disk fails and there is no hot spare disk available, contact Network Appliance Technical Support, as described in "Getting technical assistance" on page 566.

Related information

In addition to disk failure and hot spare disk replacement activity, the */etc/messages* file on the root volume logs any failure in a periodic check of the hot spare disk. For more information, refer to the *na_sysconfig*(1) and *na_messages*(5) man pages.

Disk errors

Types of disk errors described

The filer displays error messages when the following disk problems occur:

- A disk does not exist.
- A disk is in use.
- Disks are missing.

Error message: Nonexistent disks

If you get a message indicating that a disk does not exist, especially if you are adding a disk to a volume or a RAID group, make sure that

- The disk is specified correctly.
- The specified disk is a spare.

Error message: Disk in use

If you are adding a disk to a volume and you get a message indicating that a disk is in use, the disk you specified might already be a system disk. Make sure that

- The disk is specified correctly.
- The specified disk is a spare.

Error message: System cannot boot because disks are missing

You might get a message similar to the following:

```
The system cannot boot with more than one disk missing from a RAID
group.
```

This message indicates that a volume might be missing some disks because either not all the disks in a volume were transferred to a new filer or disks were damaged.

Make sure that all the disks in a volume were transferred. If the problem persists, it might mean that disks were damaged. Complete the steps in the following table to resolve the problem.

Step	Action
1	Display the boot menu from diskette, as described in "Booting from system boot diskette" on page 567, then choose Maintenance mode boot.
2	Use the `vol offline` command to take the volume off-line.
3	Reboot the filer; the volume specified in Step 2 is off-line.
4	Add missing disks, if possible, then bring them on-line; otherwise, follow these steps: 1. Replace any broken disks. 2. Destroy the old volume. 3. Create a new volume. 4. Use the `restore` command to restore the contents of the old volume from a backup tape.

Inconsistent file system

Inconsistencies seldom occur

The file system rarely becomes inconsistent. However, an inconsistent file system can be a result of combined disk and NVRAM failure.

Contact technical support if an inconsistency occurs

If your file system becomes inconsistent, contact Network Appliance Technical Support for assistance, as described in "Getting technical assistance" on page 566.

Disk operations in Maintenance mode

Maintenance mode operations

Maintenance mode enables you to perform the following operations to troubleshoot disk problems:

- Obtain detailed device information for each disk with the `disk_list` command.
- Check access to a particular disk with the `disk_check` command.
- Erase a disk label with the `disk_erase_label` command.

Note

For information about Maintenance mode, see "Booting from system boot diskette" on page 567.

Displaying detailed disk information

The `disk_list` command displays detailed device information for each disk on the system, such as drive type, and firmware revision level.

Checking access to a disk

The `disk_check` command checks access to a particular disk, issuing read requests to the disk. The disk is active for approximately 15 seconds. During this time, you can watch the disk's activity LED to verify that the disk is accessed.

Erasing a disk label

The `disk_erase_label` command erases a disk label on the specified disk drive.

Caution

Use the `disk_erase_label` command carefully because after a disk label is erased, RAID treats the disk as uninitialized and no longer recognizes it as a member of a RAID array.

Typically, use this command to remove an old RAID label on a previously used disk that you are adding to an existing RAID group. For example, use this command to erase the disk label when you move a disk from one filer to another.

Configuration problems

The /etc/rc, /etc/exports, and /etc/hosts files can contain errors

Configuration problems usually occur in one of the three configuration files on the root volume:

- */etc/rc*
- */etc/exports*
- */etc/hosts*

This section describes common configuration problems.

What to do when the filer is not accessible from the administration host

If you can access the filer from the console but not from the administration host, the filer's */etc/hosts* file on the root volume might have an IP address that is unreachable. Complete the steps in the following table to fix this problem.

Step	Action
1	Log in to the filer from the system console.
2	At the filer prompt, enter the following commands, replacing the information shown in italics with values appropriate for your filer: **`ifconfig`** ***`if`*** **`mediatype`** ***`type IP_address`*** **`netmask`** ***`netmask`*** **`exportfs -i -o root=`***`adminhost_IP_addr /`* **`nfs on`**
3	Mount the root file system from the administration host.
4	Edit the configuration files as described in Chapter 2, "Filer Administration Basics."

Filer runs setup when /etc/rc is damaged or missing

If the */etc/rc* file on the root volume is accidentally deleted, the filer automatically runs `setup` the next time it is booted.

If the system does not respond to network requests after a boot, check the console to make sure that the system is not waiting for your input.

If the filer cannot boot from the hard disk because of damaged configuration files, you can boot it from the system boot diskette. In this case, you must manually initialize the filer and correct the configuration. See "Booting from system boot diskette" on page 567 for information about booting from the system boot diskette.

How to recover from configuration errors if NFS is the only licensed protocol

If you are running NFS only, complete the steps in the following table to recover from configuration errors in the */etc/rc* file.

Step	Action
1	Display the boot menu from diskette as described in "Booting from system boot diskette" on page 567.
2	Enter **2** to choose Boot without */etc/rc*. **Note** Booting the filer without the */etc/rc* file on the root volume automatically disables CIFS service. You cannot do this procedure using CIFS.
3	At the filer prompt, enter the following two commands, replacing the variables shown in italics with values appropriate for your filer: **`ifconfig`** ***`if`*** **`mediatype`** ***`type IP_address`*** **`netmask`** ***`netmask`*** **`exportfs -i -o access=adminhost_IP_addr,root=adminhost_IP_addr nfs on`**
4	Mount the root file system from the administration host.
5	Edit the configuration files.
6	Remove the system boot diskette from the disk drive and reboot the filer to test the new configuration files.

Note

Although you can correct network problems using various keyboard commands, correct the */etc/rc* file on the root volume so that it initializes the system correctly if there is a power outage or system software failure.

How to reset the filer password

Reset the password if you forget it

If you forget your filer password, reset the password by using the system boot diskette. To avoid security problems, take care to limit access to the system boot diskette.

Procedure for resetting the password

Complete the steps in the following table to reset the filer password.

Step	Action
1	Reboot from diskette as described in "Booting from system boot diskette" on page 567.
2	When the boot menu appears, enter **3** to choose Change Password.

3	When the filer prompts you for a new password, enter it at the prompt.
	
	Results: The system prints the following message:
	
	`Password Changed`
	`Hit Return to reboot:`

4	Remove the diskette from the filer's diskette drive and reboot the filer by pressing the Enter key.

How to initialize all disks and create a new file system

Initializing all disks erases all data

You might need to initialize all disks and create a single new file system in the following circumstances:

- You decide to redeploy an existing filer and need to completely reconfigure it.
- Network Appliance Technical Support advises you that the only way to recover from an error is to initialize all disks.

Caution

Initializing all disks causes all existing data to be lost.

Procedure for initializing all disks

Complete the steps in the following table to initialize all disks.

Step	Action	
1	**If the console is...**	**Then...**
	Displaying the filer prompt, for example, `toaster>`.	Place the system boot diskette into the diskette drive of the filer and enter **`reboot`**.
	Not displaying the filer prompt.	Reboot from diskette as described in "Booting from system boot diskette" on page 567.
2	When the boot menu appears, enter **4** to choose Initialize all disks. **Results:** The filer initializes all the disks and creates a single-volume file system.	

Network problems

Detect network problems using ping at the filer console

You can detect network problems by going to the filer console and using the `ping` command.

What the ping command does

The `ping` command checks whether the filer can communicate with other hosts on the network and that other hosts can communicate with the filer.

How to troubleshoot network problems

If the filer should be able to connect with a host but `ping` does not respond with a message indicating that the host is alive, complete the steps in the following table to troubleshoot the problem.

Step	Action
1	Check that the network cable is tightly connected to the proper interface connector.
2	Use the `ifconfig` command to verify that the IP address and netmask are set correctly and that the up and running flags are displayed.
3	Use the `arp -a` command to confirm that the filer has the correct IP-to-Ethernet address map for the host you are trying to reach.
4	Use the `netstat -r` command to examine the routing tables.

Step	Action
5	Use the `netstat -i` command to check for excessive errors on the interfaces. If you see excessive input errors (ierrs) or output errors (oerrs), check the network connections on both ends of the connection. Bad transceivers or network hubs can sometimes introduce errors into the network. Collisions reported by `netstat` are a concern only if the filer detects a substantial percentage of collisions as compared to the total packet throughput. **Note** The goal is to keep collisions below 5 percent, but a network can operate properly, but slowly, with collision rates as high as 30 percent.
6	Use the `routed status` command to determine the status of the default router. **Note** An improperly set up network router can also cause network problems. If a router is not working correctly or is not configured with the filer's address, clients or hosts on the other side of the router cannot access the filer through that router.
7	If you are using the filer on a CIFS network and you are experiencing difficulty accessing the filer using name-based IP operations, for example, ping toaster, create static mappings on your WINS servers for each of the filer's interfaces.

Contact technical support about other network problem

If you encounter other problems, contact Network Appliance Technical Support for assistance, as described in "Getting technical assistance" on page 566.

NFS problems

Client's inability to mount directories indicates NFS problems

NFS problems are indicated when the filer and the client can communicate with each other using the `ping` command and the client can connect to the filer using `telnet`, but the client cannot mount volumes or directories from the filer.

How to troubleshoot NFS problems

Complete the steps in the following table to troubleshoot NFS problems.

Step	Action
1	Make sure that the filer is licensed for NFS by entering the `license` command at the filer prompt. If the following message appears `nfs not licensed` or if a message with protocols other than NFS appears but NFS is absent, you must get a license for NFS. **Note** For information about how to get a license, contact Network Appliance Technical Support, as described in "Getting technical assistance" on page 566.
2	Make sure that the filer can correctly look up the client host name.
3	Make sure that NFS service has been turned On using the `nfs on` command.
4	Make sure that the filer and the client are using correct IP addresses and names.

Step	Action
5	Sometimes a client can see the filer but gets a Permission Denied message when requesting a mount. If this happens, follow these steps: ◆ Make sure that you defined the file systems correctly in the filer's */etc/exports* file on the root volume and that you ran the `exportfs` command on the filer. ◆ On certain clients, the mount request does not come from the root user using a privileged port. The filer denies such mount requests by default to ensure secure access. To grant such mount requests, enter the following `options` command: **`options nfs.mount.rootonly off`** To make this change permanent, add the preceding command to your */etc/rc* file on the root volume.

Windows access problems

Kinds of access problems

This section describes preliminary troubleshooting steps, then describes how to troubleshoot the following problems:

- "Filer can't register with the Windows NT domain" on page 592
- "Incorrect password or unknown username" on page 593
- "Users can't map a drive" on page 594

Preliminary troubleshooting steps

Complete the steps in the following table to begin to troubleshoot Windows problems.

Step	Action
1	Make sure that the filer is licensed for CIFS by entering the `license` command at the filer prompt. If the following message appears `CIFS not licensed` or if a message with protocols other than CIFS appears but CIFS is absent, you must get a license for CIFS. **Note** For information about how to get a license, contact Network Appliance Technical Support, as described in "Getting technical assistance" on page 566.
2	If you are authenticating through a Windows NT domain, make sure that the filer has an account in the domain. **Note** You can verify that the filer is registered with a domain controller by using the `cifs testdc` command.
3	Make sure that CIFS service has been properly configured with the `cifs setup` command.

Filer can't register with the Windows NT domain

If you are using WINS, use the table in "Using Wins" to troubleshoot the problem; if you are not using WINS, use the table in "Not Using WINS" on page 593.

Using WINS: Use the following table if you are using WINS.

<table>
<tr><td colspan="5">Is the WINS server working?</td></tr>
<tr><td colspan="4">Yes</td><td>No</td></tr>
<tr><td colspan="3">Is the domain controller running?</td><td colspan="2" rowspan="6">Get the WINS server working.</td></tr>
<tr><td colspan="2">Yes</td><td>No</td></tr>
<tr><td colspan="2">Can you ping the domain controller?</td><td rowspan="5">Start the domain controller.</td></tr>
<tr><td>Yes</td><td>No</td></tr>
<tr><td colspan="2">Can you ping the filer?</td><td rowspan="3">Check network connec-tivity</td></tr>
<tr><td>Yes</td><td>No</td></tr>
<tr><td>Contact Network Appliance Technical Support, as described in "Getting technical assistance" on page 566.</td><td>Check net-work con-nectivity</td></tr>
</table>

Not Using WINS: Use the following table if you are not using WINS.

<table>
<tr><th colspan="3">Is the PDC or BDC on the same subnet as the filer, and is the subnet connected to the first configured network interface card on the filer?</th></tr>
<tr><th colspan="2">Yes</th><th>No</th></tr>
<tr><td colspan="2">Is there a computer account created for the filer in the Windows NT domain?</td><td rowspan="2">Put the domain controller on the same subnet as the filer, and connect the subnet to the first configured network interface card on the filer.

Without a WINS server, the filer can talk only to the domain controller by broadcast. The filer broadcasts only from the first configured network interface. The domain controller must be on this subnet.</td></tr>
<tr><th>Yes</th><th>No</th></tr>
<tr><td>Contact Network Appliance Technical Support, as described in “Getting technical assistance” on page 566.</td><td>Create a computer account for the filer in the domain.</td></tr>
</table>

Incorrect password or unknown username

Use the following table to troubleshoot the problem.

<table>
<tr><th colspan="4">Are users in the same domain as the filer?</th></tr>
<tr><th colspan="2">Yes</th><th colspan="2">No</th></tr>
<tr><th colspan="2">Are you using a virtual LAN?</th><th colspan="2">Is there a trust relationship between the filer and the domain?</th></tr>
<tr><th>Yes</th><th>No</th><th>Yes</th><th>No</th></tr>
<tr><td>Have the users log in with DOMAINNAME \ USERNAME</td><td>Enter the users into the /etc/passwd file.</td><td>Have the users log in with DOMAINNAME \ USERNAME</td><td>Establish a trust relationship between the user’s account domain and the filer’s domain.</td></tr>
</table>

Users can't map a drive

If users get an Access Denied message, use the following decision table to troubleshoot the problem.

<table>
<tr><th colspan="3">Is the filer in the same domain as the PDC?</th></tr>
<tr><th colspan="2">Yes</th><th>No</th></tr>
<tr><td colspan="2">Does the user have access rights to the share?</td><td rowspan="3">Use Server Manager to add the filer to the domain or establish a trust relationship between the domains.</td></tr>
<tr><td>Yes</td><td>No</td></tr>
<tr><td>Contact Network Appliance Technical Support, as described in "Getting technical assistance" on page 566.</td><td>Give the user rights to the share.</td></tr>
</table>

UNIX cpio problems

The cpio version should support 32-bit inode definition numbers

If you copy large amounts of data using the UNIX `cpio` utility, some files might be copied incorrectly. This happens in UNIX versions of `cpio` that still use a 16-bit inode definition number. Large file systems require a 32-bit inode definition number. The problem generally occurs only on file systems with a large number of hard links.

Why the problem occurs: Some versions of `cpio` work by copying each file with hard links once and then re-creating the hard links. Trouble occurs because the inode number is assumed to be 16 bits. If another file has a matching low-order 16 bits, an internal number collision occurs, which `cpio` does not recognize. The `cpio` utility then overwrites the first file and creates files that no longer contain the original data.

Ask UNIX provider whether cpio version supports 32-bit inode definition numbers

Check with your UNIX provider to see whether your version of `cpio` has this problem. If you use SunOS 4.x, ask for the Sun patch 100556-01. This patch works around the problem by requiring that the files have the same UID, GID, mode, mod time, inumber, and device before concluding that they are the same file.

UNIX df problems

The df version must support large file systems

Some UNIX versions of the `df` command have file system limits considerably smaller than the file system size supported by the filer. This can cause the UNIX `df` command to show an incorrect and useless amount of filer disk space in use or remaining. However, the disk space you installed in your filer is fully available and you can use it.

Enable NFS option to avoid displaying useless data

To avoid a useless display of disk space on a client system that uses NFS version 2, enable the `nfs.v2.df_2gb_lim` option, as described in "Configuring filer options" on page 52. For additional information, refer to the *na_options*(1) man page.

DOS, Windows, and Macintosh clients might have display problem

Some DOS, Windows, or Macintosh clients might have a display problem similar to UNIX systems; in these cases, enable the `nfs.v2.df_2gb_lim` option.

Filer df command always shows correct disk space

At all times, the `df` command entered on the filer correctly shows the amount of disk space used and remaining.

Qtrees affect disk space displayed by df

If a directory in a qtree is mounted and a client issues a `df` command on something under that mount point, the command shows the smaller of the client's file system limit or the filer disk space. This makes the qtree look fuller according to the client `df` command than it actually is.

Filer quota report command always displays correct usage

The filer `quota report` command shows the correct usage within that qtree.

Cluster messages

Kinds of cluster messages

This section lists a console message, LCD messages, and system messages that indicate conditions that you can take steps to address.

What to do when you see a message that indicates trouble

If a filer or cluster ceases to function properly, save or record as many messages as you can. If the message is not listed here, contact Network Appliance Technical Support, as described in "Getting technical assistance" on page 566.

Where cluster LCD messages appear

LCD messages appear on two lines on the filer's LCD. Cluster messages, if they appear at all, appear on the first line of the LCD; other messages appear on the second line.

Cluster LCD messages that indicate trouble are accompanied by an amber LED indicator.

Cluster LCD messages that you can address

The following messages pertaining to clusters can appear on the LCD of a takeover filer. These are messages that you can take steps to address

Message	Explanation	What you do
`cf disabled`	Takeover has been disabled.	Follow these steps: 1. Unless you want takeover to be disabled, issue the `cf enable` command. 2. Check the interconnect cables to make sure that they are seated properly. 3. Issue the `cf status` command to verify the link.

Message	Explanation	What you do
`Iconnect down!`	The interconnect link is broken.	Follow these steps: 1. Check the cluster interconnect cables to make sure that they are seated properly. 2. Issue the `cf status` command to verify the link.
`Iconnect error`	One of the interconnect links is broken.	Follow these steps: 1. Check the interconnect to make sure that it is functioning properly. 2. Issue the `cf status` command to verify the link.
`L: vol` *`volume`* *`n`*`% rebuilt`	The filer has taken over its partner and one of its own volumes is reconstructing itself.	Wait until the message disappears.
`P: vol` *`volume`* *`n`*`% rebuilt`	The filer has taken over its partner and one of its partner's volumes is reconstructing itself.	Wait until the message disappears.
`Partner down`	The partner is down and takeover is disabled.	If you want to take over the partner, issue the `cf enable` command, then the `cf takeover` command.
`Takeover`	The filer has taken over its partner.	Perform a giveback.
`Takeover in` *`nnn`*	The filer is counting down seconds until it takes over its partner.	Unless you want takeover to occur, issue the `cf disable` command.

The following message pertaining to a cluster can appear on the LCD of a failed filer.

Message	Explanation	What you do
`taken over`	The filer has been taken over.	Perform maintenance on the filer, then perform a giveback from the takeover filer.

Where syslog messages appear

Syslog messages appear by default in the */etc/messages* file. For additional information about syslog messages, see the Message logging section of Chapter 2 of the *System Administrator's Guide*.

Syslog messages that you can address

The following syslog messages are messages that you can take steps to address. By default they appear on the system console.

Message	Explanation	What you do
`cf.timed.enable: timed requires java which is not enabled on this system`	Essential Java files are missing from the filer.	Follow these steps: 1. Install the Java .zip files from the distribution CD-ROM onto the */etc/java* directory of your filers. 2. Reboot.
`Interconnect is down`	The cluster interconnect is not working properly. Takeover capability is disabled.	Follow these steps: 1. Check the cluster interconnect cables to make sure that they are seated properly. 2. If the connections are secure, and the message returns, check for Cluster Interconnect adapter messages in the */etc/messages* file, then contact Network Appliance Technical Support, as described in "Getting technical assistance" on page 566.

Message	Explanation	What you do
`No disks were detected on Fibre Channel port B; this filer will be unable to takeover correctly`	The B-port loop has not been connected to the local filer, which means that the local filer cannot take over its partner.	Follow these steps: 1. Check the B-port loop connections to the partner's disk shelves and make sure that they are securely connected. 2. Verify that the LRC cards are properly seated.
`Root volume` *`volume`*`, RAID group` *`n`* `is missing` *`n`* `disks and is unusable.`	A disk shelf or adapter failed, or the root volume's disks might be spread between A-port and B-port loops.	Check the disk shelves and make sure that they are wired properly.
`Volume` *`volume`* `is missing a data disk and has dirty parity.`	In this situation, the filer is refusing to perform a giveback because a volume is in degraded mode with dirty parity.	**Caution** You can lose data if you perform a giveback with a volume that is missing a data disk and has dirty parity. Follow these steps: 1. Wait until one if these events happen: ❖ Parity is clean. ❖ The data disk is reconstructed. 2. Verify that one of the events in Step 1 has happened using the `sysconfig -r` command. 3. Perform a giveback.
`Volume` *`volume`* `is missing` *`n`* `disks from RAID group` *`n`* `and is being taken offline for this boot.`	The volume's disks might be spread between A-port and B-port loops.	Follow these steps: 1. Verify that the disk shelves are wired properly. 2. Reboot the filer to clear this condition.

Message	Explanation	What you do
`WARNING cluster monitor fast timeout was blocked for` *n* `secs.`	The Data ONTAP kernel is supposed to schedule the cluster monitor's fast timeout thread at regular intervals to ensure that mailbox data is read and written within specific time frames. If a filer does not receive the mailbox data because the fast timeout thread is blocked, it might assume that its partner is down and attempt to perform a takeover.	If this message happens frequently, gather system utilization data (`sysstat`, profiling data, and so on), then contact Network Appliance Technical Support, as described in "Getting technical assistance" on page 566 so that they can characterize the system.
`X link is down`	The Cluster Interconnect has detected a failure in one of its two communication channels. This circumstance does not disable takeover but deserves investigation. **Note** This circumstance can also occur temporarily during the booting of a clustered filer.	Check the X (top) cable connections to make sure that the cables are seated properly.
`Y link is down`	The Cluster interconnect has detected a failure in one of its two communication channels. This circumstance does not disable takeover but deserves investigation. **Note** This circumstance can also occur temporarily during the booting of a cluster filer.	Check the Y (bottom) cable connections to make sure that the cables are seated properly.

About console messages

Some messages appear on the console but are not written to the */etc/messages* file. Most of them are self-explanatory.

Console message

The following table explains the console message “The cluster partner appears not to be operational.”

Message	Explanation	What you do
`The cluster partner appears not to be operational. You will be asked whether you want to continue. If you answer "yes", the existing cluster monitor disk state will be overwritten and this node will be rebooted. Answering "no" will halt this node with no modification to the cluster monitor disk state.` `WARNING: Answering "yes" while the cluster partner is operational will have unexpected and potentially catastrophic results: YOUR FILESYSTEMS MAY BE DESTROYED`  `Do you wish to continue [y/n]?`	The local filer has been taken over but it appears its partner is not operational. This message is usually preceded by the following message: `This node was previously declared dead. Pausing to check cluster partner status ... partner not responding.` `Thu Aug 13 16:58:37 GMT [kbd_login_mgr]: Login from console`	Follow these steps: 1. Verify that the partner is operational or repair it. 2. Reboot it. 3. Perform a giveback from the partner. Giveback is the only method to ensure that data is consistent after a takeover. 4. If you cannot repair the partner, perform one of the following actions: ❖ Contact Network Appliance Technical Support, as described in “Getting technical assistance” on page 566. **Caution** If you enter **y** and override the takeover, client applications might lose data that was written before the reboot. ❖ Enter **y** to continue and override the takeover and allow the local filer to reboot.

Serious error messages

Panic messages mean serious problems

If your filer has a serious problem, such as a problem with the hardware or a bug in the system software, it issues a system panic message similar to the following one:

```
PANIC: system hung (NS0)!
Volume: volname
Version: verno
```

Message Component	Description
`system hung (NS0)`	Indicates the panic class of the message and is significant. The actual text of the message varies with circumstances.
`volname`	Is the name of the volume.
`verno`	Is the version number.

What to do after a panic message

Complete the steps in the following table when your system issues a panic message.

Step	Action
1	Write down the panic message.
2	Call Network Appliance Technical Support immediately, as described in "Getting technical assistance" on page 566.
3	Provide the panic message to Network Appliance Technical Support.

SecureAdmin management 22

SecureAdmin overview

What SecureAdmin does

The SecureAdmin™ feature enables you to use encrypted passwords and login names to administer filers. When encrypted, passwords and logins become impervious to being discovered over a network, thus improving filer security. Anything sent to or from any console session or remote shell session is encrypted.

How you use SecureAdmin

You can use SecureAdmin only for administrative tasks. Although all information exchanged between a filer and a client is encrypted, because you cannot write files directly to the filer, SecureAdmin can protect only logins and passwords.

Available only in the U.S. and Canada

SecureAdmin is available only in the U.S. and Canada. Users in other parts of the world cannot install or use SecureAdmin.

SecureAdmin concepts

Supported protocol and version

SecureAdmin provides a secure remote shell and interactive network sessions. It supports SSH 1.0 clients only.

SecureAdmin supports only password authentication and DES (Data Encryption System) and 3DES session encryption. DES is an encryption algorithm, and 3DES is an encryption algorithm that relies on triple encryption.

Note

If you run the `secureadmin` command, other protocols might appear, but you should disregard any references to other protocols.

SecureAdmin administration

A daemon (`sshd`) runs the SecureAdmin feature. You administer SecureAdmin with the `secureadmin` command. You can perform the following operations on SecureAdmin:

- Set up and start SecureAdmin, as described in "Setting up and starting SecureAdmin" on page 609.
- Reinitialize SecureAdmin and generate new keys, as described in "Reinitializing SecureAdmin" on page 610.
- Disable SecureAdmin, as described in "Disabling SecureAdmin" on page 612.
- Restart SecureAdmin, as described in "Restarting SecureAdmin" on page 613.
- Get the status of SecureAdmin, as described in "Determining whether SecureAdmin is running" on page 614.
- Display the version of SecureAdmin you are running, as described in "Displaying the version of SecureAdmin" on page 615.

SecureAdmin uses RSA encryption

SecureAdmin uses public-key encryption provided by RSA for host authentication and session encryption.

Host and server key properties

Host and server keys have the following properties:

- They can range in strength from 384 through 1,024 bits. The higher the strength the greater the security, but the longer it takes to initiate a new SecureAdmin session, and the higher the likelihood of lowering filer performance.
- The two keys must differ by at least 128 bits.
- The default values are as follows:
  - Host key: 768 bits
  - Server key: 512 bits

SecureAdmin keys

SecureAdmin performs encryption with three keys: a host key, a server key, and a session key, as described in the following table. The filer and clients generate keys with DES, 3DES, or proprietary algorithms.

Key	Description
Host key	One of the keys that the filer uses in encrypting and decrypting the session key. The filer generates the host key when you run the SecureAdmin `setup` program.
Server key	One of the keys that the filer uses in encrypting and decrypting the session key. The filer generates the server key when the following occurs: ◆ An hour elapses ◆ You start SecureAdmin ◆ The filer reboots
Session key	A key generated by the client. The session key encrypts all the data in a network session.

How SecureAdmin uses keys

The following process table shows how the filer and an SSH 1.0 client use keys.

Stage	What the client does	What the filer does
1	The client sends an SSH request to the filer.	The filer receives an SSH request from a client.
2		The filer sends to the client the public portion of the host and server keys.
3	The client stores the public portion of the host key.	
4	The client generates a session key using client-specific procedures.	
5	The client encrypts the session key using the public portion of the host and server keys and sends it to the filer.	
6		The filer decrypts the session key using the private portions of the portion of the host and server keys.
7	The filer and the client exchange information that they encrypt and decrypt using the session key.	

Setting up and starting SecureAdmin

Description

Use this procedure to set up and start SecureAdmin. It generates a host key and a server key, which are used in encrypting the key used during a session. You complete this procedure after you install SecureAdmin.

Prerequisites

You can accept the default sizes of 768 bits for the host key and 512 bits for the server key, or supply a size for the host key and a size for the server key, as described in "Host and server key properties" on page 607. The keys can range from 384 through 1,024 bits. The keys must differ by at least 128 bits.

Steps

To set up and start SecureAdmin, complete these steps.

Step	Action
1	Enter the following command: `secureadmin setup` **Result:** The SecureAdmin `setup` program starts.
2	When prompted, enter a size for the host key.
3	When prompted, enter a size for the server key.
4	When prompted, confirm the parameters that you specified. **Result:** SecureAdmin generates the host key in the background, and, after a minute or two, the `setup` program sends a `syslog` message that announces that SSH Setup is done.
5	After the `syslog` message is generated, make the keys take effect by entering the following command: `secureadmin enable ssh`

Reinitializing SecureAdmin

Description

Use this procedure to generate new host and server keys. You perform this procedure when you want new keys or when you want to change the sizes of keys. After you complete this procedure, SecureAdmin is reinitialized and new keys are generated to your new specification.

Step

To reinitialize SecureAdmin, complete these steps.

Step	Action
1	Enter the following command: **`secureadmin -f setup`**
2	When prompted, enter a size for the host key.
3	When prompted, enter a size for the server key.
4	Cancel the existing keys by stopping the `sshd` daemon by entering the following command: **`secureadmin disable ssh`**
5	Make the keys take effect by entering the following command: **`secureadmin enable ssh`**

Sample output

Some clients might send the following warning when they detect the change in the host key. You need not do anything in response to it.

```
@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@
@       WARNING: HOST IDENTIFICATION HAS CHANGED!         @
@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@
IT IS POSSIBLE THAT SOMEONE IS DOING SOMETHING NASTY!
Someone could be eavesdropping on you right now (man-in-the-middle
attack)!
It is also possible that the host key has just been changed.
Please contact your system administrator.
Add correct host key in /u/mlam/.ssh/known_hosts to get rid of this
message.
```

```
Agent forwarding is disabled to avoid attacks by corrupted servers.
Are you sure you want to continue connecting (yes/no)?
```

Disabling SecureAdmin

Description

Use this procedure to turn off SecureAdmin. You turn off SecureAdmin when you want transactions between your filer and your client not to be protected by encryption. You can perform this procedure any time after installing SecureAdmin on your client. After you complete this procedure, passwords and login names are sent in clear text.

Step

To turn off SecureAdmin, complete this step.

Step	Action
1	Enter the following command: **`secureadmin disable ssh`**

Restarting SecureAdmin

Description

Use this procedure to restart SecureAdmin. You restart SecureAdmin when you want administrative transactions between your filer and your client to be protected by encryption. You can restart SecureAdmin any time after disabling SecureAdmin on your client. After you complete this procedure, passwords and login names are sent in encrypted text.

Step

To restart SecureAdmin, complete this step.

Step	Action
1	Enter the following command: `secureadmin enable ssh`

Determining whether SecureAdmin is running

Description

Use this procedure to determine whether SecureAdmin is running, and therefore whether administrative transactions between the filer and a client are being encrypted. You can perform this procedure any time after installing SecureAdmin on your client.

Step

To find out whether SecureAdmin is running, complete this step.

Step	Action	
1	Enter the following command: `secureadmin status` **Result:** One of the following displays appears.	
	Status	**Display**
	On	`ssh  - running`
	Off	`ssh  - inactive`

Displaying the version of SecureAdmin

Description

Use this procedure to display the version of SecureAdmin you are running. You can display the version of SecureAdmin any time after starting SecureAdmin.

Step

To display the version of SecureAdmin you are running, complete this step.

Step	Action
1	Enter the following command: **`secureadmin version`** **Result:** The following display appears: `SecureAdmin (tm) 1.0`

Detailed Options Information 23

About options

About setting detailed information

Options are also described in other sections of this guide and in the man pages. The *na_options*(1) man page contains information about `options` command options. The *na_vol*(1) man page contains information about the `vol options` command options.

Option values

The default value of an option is listed is below the option name.

The following conventions apply to default values listed:

- "None" means that there is no default value.
- If the default is "On," the other possible value is "Off."
- If the default is "Off," the other possible value is "On."

The values for On and Off are not case-sensitive. If you do not supply a parameter in the `options` command, the command prints the current values of all available options.

Option groups

The sections that follow provide a brief description of each option grouped by function. The sections areas follows:

- "Autosupport options" on page 618
- "CIFS options" on page 620
- "DNS options" on page 624
- "HTTP options" on page 625
- "NFS and (PC)NFS options" on page 627
- "NIS options" on page 629
- "RAID options" on page 630
- "Timed options" on page 631
- "Volume options" on page 633
- "Miscellaneous options" on page 635

Autosupport options

What the autosupport options do

The autosupport options control whether and how the filer sends automatic status messages.

For more information about autosupport, see "Sending automatic email to Network Appliance" on page 55.

The autosupport.doit option

Default: None

Description: Immediately sends an email message describing the status of the filer. A word entered as the value for the option is sent in the notification subject line and should describe the reason for the notification.

The autosupport.enable option

Default: On

Description: Enables the `autosupport` daemon, which sends automatic email messages to report the status of the filer.

The autosupport.from option

Default: autosupport

Description: Specifies the sender of the automatic email message. It should be the email address of the person whom Network Appliance contacts after receiving an automatic email message.

The autosupport. mailhost option

Default: administration_host

Description: Specifies the mail hosts that receive automatic email messages. Use a comma-separated list with no spaces.

The autosupport.noteto option

Default: None

Description: Specifies up to five recipients of an automatic short email message. Use a comma-separated list with no spaces.

The autosupport.to option

Default: autosupport@netapp.com

Description: Specifies up to five recipients of autosupport email messages. Use a comma-separated list with no spaces.

CIFS options

What the CIFS options do

The CIFS options control CIFS features on the filer.

The cifs.access_logging_enable option

Default: Off

Description: When On, enables the filer to process access logging, or auditing, information. The default is Off.

The cifs.access_logging.filename option

Default: None

Description: Specifies the active event log file. The file must be in an existing directory in a network share.

The cifs.bypass_traverse_checking option

Default: On

Description: When On, directories in the path to a file are not required to have the 'X' (traverse) permission.

Note

This option does not apply in UNIX qtrees or volumes.

The cifs.guest_account option

Default: None

Description: When you leave `cifs.guest_account` blank, a CIFS user can log in to the filer without an account in the password database, provided that a domain controller authenticates the user.

When you set `cifs.guest_account` to the name of an account in the NIS `passwd` map or */etc/passwd* (typically guest), guest access to the filer is enabled, and the user has the UNIX user ID and the group ID of the guest account. For more information, see "Enabling guest and generic access" on page 282.

The cifs.home_dir option

Default: None

Description: Specifies the complete path name of the “homes directory.” The directories under this path should have the names of users as their names. When a CIFS user connects to the filer and there is a directory name that exactly matches the user’s name, the user sees a share of that name (truncated to 12 characters) that is the user’s home directory. Only the user can access the home directory using this share. All other users are denied access.

The cifs.idle_timeout option

Default: 18000

Description: Specifies the number of seconds that elapse before the filer disconnects an idle session. The value can range from 600 through 4,000,000 (effectively infinite).

The cifs.netbios_aliases option

Default: None

Description: Specifies a list of alternative names for the filer. Use a comma-separated list of names.

The cifs.oplocks.enable option

Default: On

Description: When this option is On, the filer enables clients to use oplocks (opportunistic locks) on files. Oplocks provide a significant performance enhancement, but have the potential to cause lost cached data on some networks with impaired reliability or latency, particularly wide-area networks. In general, you should disable this option only if there are problems with databases and to isolate problems.

The cifs.perm_check _use_gid option

Default: On

Description: This option affects security checking for Windows clients of files with UNIX security where the requestor is not the file owner. In all cases, Windows client requests are checked against the share-level ACL. If the requestor is the owner, the User permissions determine the access.

If the requestor is not the owner and if perm_check_use_gid is On, files with UNIX security are checked using normal UNIX rules; that is, if the requestor is a member of the file's owning group, the Group permissions are used, otherwise the Other permissions are used.

The cifs.scopeid option

Default: blank

Description: Specifies a second element for a single-element NetBIOS computer name. This element is case-sensitive. You use this option to isolate a group of computers on a network that communicate only with other computers with the identical NetBIOS Scope ID.

This option is not recommended if you are using DNS for name resolution because NetBIOS Scope IDs and DNS are incompatible.

For more information, refer to Microsoft documentation.

The cifs.cifs.search_domains option

Default: None

Description: Specifies a list of domains that trust each other to search for a mapped account. The argument for the option is a comma-separated list that is searched in order. If no list is supplied, all domains are searched. You use this option to limit searches if you used an asterisk for a domain name in the usermap.cfg file.

The cifs.show_snapshot option

Default: FALSE

Description: Specifies whether to show the *~snapshot* snapshot directories in folders. To show the snapshot directories, set this option to TRUE.

The cifs.symlinks.cycleguard option

Default: On

Description: If an object being accessed by a CIFS client is a symbolic link, the `cifs.symlinks.cycleguard` option, when set to On, eliminates the possibility of cyclic directories. It does so by preventing the following of symbolic links that contain the "dot" (".") or "dot-dot" ("..") component—symbolic links that could refer to a directory higher in the same tree. With the

`cifs.symlinks.cycleguard` option set to Off, if you are careful, you can use symbolic links having “dot” or “dot-dot” components. For more information about this option, refer to ““How to prevent symbolic link cycling” on page 194.

The cifs.symlinks .enable option

Default: On

Description: When you set `cifs.symlinks.enable` to On (the default setting), if the object being accessed by a CIFS client is a symbolic link, the filer follows the link with the condition that the ultimate target turns out to reside within the originating share. This ensures that the client has access permission to the target. This applies both to relative symbolic links (links to paths beginning with a character other than / and treated as a path relative to the parent directory of the symbolic link) and absolute symbolic links (links to paths beginning with / and treated as a path relative to the root of the file system). For more information about this option, refer to “Managing symbolic links for CIFS access” on page 192.

DNS options

What the DNS options do

The DNS options control how the filer works with DNS.

For more information about DNS, see Chapter 4, “Network Administration.”

The dns.domainname option

Default: None

Description: Sets the DNS domain name to the specified domain name.

The dns.enable option

Default: Off

Description: Enables the DNS client on the filer. Before you enable DNS, you must set the DNS domain and the */etc/resolv.conf* file must exist.

HTTP options

What the HTTP options do

The HTTP options enable and control HTTP services.

For more information about HTTP on the filer, see Chapter 9, "HTTP Administration."

The httpd.admin.enable option

Default: On

Description: Enables HTTP access to the filer's on-line Help files and other files used by FilerView.

The httpd.enable option

Default: Off

Description: Enables the HTTP server.

The httpd.log.max_file_size option

Default: 2147483647 (2 GB - 1 byte)

Description: Specifies the number of bytes */etc/log/httpd.log*, the HTTP log file, can grow to. The maximum value is 500 GB.

The httpd.rootdir option

Default: None

Description: Specifies the root directory containing files and directories that HTTP transfers to clients.

The httpd.timeout option

Default: 900 seconds (15 minutes)

Description: Specifies the minimum amount of time, in seconds, before an idle HTTP connection times out.

The httpd.timewait.enable option

Default: On

Description: When you set this option to On, the filer drops an HTTP connection one minute after the client closes it. When you set this option to Off, the connection is not dropped and resources are consumed until the connection times out.

NFS and (PC)NFS options

What the NFS and (PC)NFS options do

The NFS and (PC)NFS options enable and control NFS services.

For more information about NFS, see Chapter 7, "NFS Administration."

The nfs.mount_rootonly option

Default: On

Description: Configures the filer to give NFS access only to root by requiring that mount requests come from privileged ports (ports 0 through 1023).

The nfs.per_client_stats.enable option

Default: Off

Description: Specifies whether the filer collects and displays NFS statistics from individual clients.

The nfs.tcp.enable option

Default: Off

Description: Specifies whether the filer supports NFS over TCP. Enable this option if a client has problems using NFS over UDP.

The nfs.v2.df.2gb.lim option

Default: Off

Description: Limits to 2 GB the response the filer gives to requests from NFS v2 clients regarding total space, free space, or available space. This option is necessary for some NFS clients to calculate the amount of free space accurately. Without this option, a file system with more than 2 GB of free disk space might appear to be full to the client that initiates the "file system statistics" request.

The nfs.v3.enable option

Default: On

Description: Specifies whether the filer supports NFS Version 3. Disable this option if a client has problems using NFS Version 3 and that client cannot be configured to use NFS Version 2.

The nfs.webnfs.enable option

Default: Off

Description: Turns WebNFS On and Off.

The nfs.webnfs.rootdir option

Default: None

Description: The specified directory becomes the root or public directory for WebNFS. When a request specifies a relative path, lookups for files are done with respect to this directory.

The nfs.webnfs .rootdir.set option

Default: FALSE

Description: When set to TRUE, sets the directory specified in the `nfs.webnfs.rootdir` option to be the WebNFS root or public directory.

The pcnfsd.enable option

Default: Off

Description: Specifies whether the filer runs the (PC)NFS server.

The pcnfsd.umask option

Default: 022

Description: Specifies that users can read and write their own files, but members of their group and others can only read the files. This option uses the standard UNIX permission mask.

NIS options

What the NIS options do

The NIS options control how the filer works with NIS.

For more information about NIS, see Chapter 4, “Network Administration.”

The nis.domainname option

Default: None

Description: Sets the NIS domain to the specified domain name.

The nis.enable option

Default: Off

Description: Enables the NIS client on the filer. You must set the NIS domain before you enable NIS.

RAID options

What the RAID options do

The RAID options control how the filer uses RAID.

For more information about RAID on the filer, see Chapter 3, "Disk and File System Management."

The raid.reconstruct_speed option

Default: 4

Description: Specifies the speed of RAID reconstruction. The speed ranges from 1 (slowest) to 10 (fastest). The filer uses the number to determine the percentage of CPU time used for RAID reconstruction.

The raid.scrub.enable option

Default: On

Description: Specifies whether the filer performs RAID scrubbing.

The raid.timeout option

Default: 24

Description: Sets the time, in hours (from 1 through 24), that the system runs in degraded mode before an automatic shutdown.

Timed options

What the timed options do

The `timed` options control whether and how the filer uses the `timed` daemon to synchronize time with a time server. For additional information about time synchronization and using the `timed` options, see "Filer system time synchronization" on page 60 and "Synchronizing filer system time" on page 64.

The timed.enable option

Default: Off

Description: Determines whether a time daemon (`timed`) runs on the filer and synchronizes time with a time server.

The timed.log option

Default: Off

Description: Determines whether to log to the console time changes initiated by the `timed` daemon.

The timed.max_skew option

Default: 30m

Description: Sets the maximum allowable discrepancy between filer time and server time. If there is a large discrepancy, it probably means that enough is wrong somewhere that the filer time should not be synchronized with the server time, no time synchronization takes place, and a message to that effect is sent to the console. The value is an integer followed by one of the following letters:

- s for seconds
- m for minutes
- h for hours

The default is 30 minutes.

The timed.proto option

Default: ntp

Description: Selects whether to use the protocol used by the `rdate` command or SNTP. The value can be one of the following:

- rdate for the protocol used by the `rdate` command
- ntp for SNTP

The timed.sched option

Default: hourly

Description: Schedules when to synchronize the time with a time server. The value can be one of the following:

- hourly to synchronize hourly
- multihourly to synchronize every six hours
- daily to synchronize every day at midnight
- a number followed by m to specify an interval of minutes or h to specify an interval of hours

To avoid overburdening the time server, the filer randomly selects the exact time of the synchronization within a 20-minute window of the specified schedule.

The timed.servers option

Default: None

Description: Specifies up to five servers in order of contact priority. The value can be a list of the host names or IP addresses of up to five time servers, separated by commas. The filer goes down the list until it finds a server that responds, then uses the time from that server. It starts with the first server in the list each time. An example list of three time servers is as follows:

sundial, sundial.netapp.com, 10.152.8.12

You can get a list of NTP (Network Time Protocol) time servers, which SNTP can use, from *http://www.eecis.udel.edu/~mills/ntp/servers.htm*.

Volume options

What the volume options do

The volume options control volume-level operations. You use these options only with the `vol options` command.

For more information about volume options, see "The vol options command" on page 54.

The minra option

Default: Off

Description: Configures the filer to perform minimal read ahead. By default, the option is disabled and the filer does aggressive read ahead.

The no_atime_update option

Default: Off

Description: Prevents the update of the access time (atime) on an inode when a file is read. This option prevents inode updates from contending with reads from other files. Use it only on a filer with extremely high read traffic (for example, on a news server used by an Internet access provider or on a filer used mainly as an HTTP server).

The nosnap option

Default: Off

Description: Temporarily disables automatic snapshots.

The nosnapdir option

Default: Off

Description: Makes invisible the snapshot directory that's usually present at the client mount point or at the root of the CIFS share. It also turns off access to the snapshot directory and all snapshot directories under the mount point or the root of the CIFS share.

After you toggle this option, you might not notice the effect immediately because the information about the snapshot directories might still be in the client's attribute cache. To force the change to take effect immediately, unmount and remount the file system.

The nvfail option

Default: Off

Description: Sets the filer to check for NVRAM errors during boot up. Change the value to On when you want the filer to send error messages to notify you of NVRAM errors that can effect the validity of database files.

The raidsize option

Default: None

Description: Sets the maximum size of a RAID group in *volume*. Must be an integer greater than one.

The root option

Default: None

Description: Makes *volume* the root volume.

The snapmirrored option

Default: None

Description: The filer automatically sets this option to On if the volume is a mirror for data replication. Otherwise, the option is set to Off. Change the value to Off if you want to convert a mirror to a regular volume. After you make the change, the volume is no longer read-only, and the filer stops making incremental changes to the volume for data replication.

Note

You cannot set this option to On. That is, you cannot use this option to convert a regular volume to a mirror. To use a volume as a mirror, follow the instructions in Chapter 19, "Data Replication Using SnapMirror," to start replicating data to the volume.

Miscellaneous options

What the miscellaneous options do

The miscellaneous options control additional aspects of filer operation.

The console.encoding option

Default: nfs

Description: Specifies how non-ASCII character information is presented. The value can be one of the following.

Value	Description
nfs	NFS character set. You can use both NFS extended (greater than 0x7F) and SGML characters for input.
sgml	SGML character format. You can use both NFS extended (greater than 0x7F) and SGML characters for input.
utf8	UTF-8 character sets. For input, any character greater than 0x7F is the beginning of a UTF-8 encoding.

The ip.match_any_ifaddr option

Default: On

Description: If the option is On (the default), the filer accepts any packet that is addressed to it even if that packet came in on the wrong interface.

Note

If you are concerned about security, you should turn this Off.

The ip.path_mtu_discovery.enable option

Default: On

Description: Enables or disables path MTU discovery, which is currently used only by TCP. When enabled, the filer can discover and use the largest packet size that the filer can send to another host without fragmenting a packet. This means that the filer doesn't have to limit itself to sending many small packets, which takes more time and resources than sending fewer large packets.

If you cannot establish a connection, set this option to Off.

The rsh.enable option

Default: On

Description: Enables the `rsh` server on the filer.

The snmp.enable option

Default: On

Description: Enables the SNMP server on the filer.

The telnet.enable option

Default: On

Description: Enables the Telnet server on the filer.

The telnet.hosts option

Default: All hosts

Description: Specifies a list of hosts that can log in to the filer using `telnet`. You can limit `telnet` access to up to five specified hosts. The hosts should be listed in a comma-separated list. You can disable `telnet` for all hosts by specifying a hyphen (-).

The vol.copy.throttle option

Default: 10

Description: Specifies the default speed of `vol copy` operations. The speed ranges from 10 (full-speed) to 1 (one-tenth of full-speed).

The wafl.convert_ucode option

Default: Off

Description: Setting this option to On forces conversion of all directories to Unicode format when accessed from both NFS and CIFS. By default, conversion to Unicode format occurs as follows:

- Access from CIFS causes conversion of pre-4.0 and 4.0 format directories.
- Access from NFS causes conversion of 4.0 format directories.

The wafl.create_ucode option

Default: Off

Description: Setting this option to On forces Unicode format directories to be created by default, both from NFS and CIFS. By default, all directories are created in pre-4.0 format and the first CIFS access converts a directory to Unicode format.

The wafl.default_nt_user option

Default: None

Description: Specifies the Windows NT user account to use when a UNIX user accesses a file with Windows NT security (has an ACL), and that UNIX user would not otherwise be mapped. If this option is set blank, such accesses are denied.

The wafl.default_unix_user option

Default: None

Description: Specifies the UNIX user account to use when a Windows NT user attempts to log in and that Windows NT user would not otherwise be mapped. If this option is set blank, such accesses are denied.

The wafl.maxdirsize option

Default: 10240

Description: Sets the maximum size, in kilobytes, of a directory file. A directory file with a size of 10,240 kilobytes can hold about 300,000 files or subdirectories.

The wafl.root_only_chown option

Default: On

Description: Enables only the root user to change the owner of a file. When you disable the option, the owner of a file can change its ownership without being root. By default, this option is enabled.

When a non-root user changes the owner of a file, the set-user-id and set-group-id bits are cleared. If a non-root user tries to change the owner of a file but the change would cause the file’s recipient to exceed his or her quota, the attempt fails.

The wafl.wcc_minutes_valid option

Default: 20

Description: Specifies the number of minutes a WAFL credential cache entry is valid. The value can range from 1 through 20160.

Glossary

ACL

Access control list. A list that contains the users' or groups' access rights to each share.

adapter card

A SCSI card, network card, hot swap adapter card, serial adapter card, or VGA adapter that plugs into a filer expansion slot.

Address Resolution

The procedure for determining a Media Access Control (MAC) address corresponding to the address of a LAN or WAN destination.

administration host

The client you specify during filer setup for managing the filer. The `setup` program automatically configures the filer to accept `telnet` and `rsh` connections from this client, to give permission to this client for mounting the */* and */home* directories, and to use this client as the mailhost for sending autosupport email messages. At any time after you run the `setup` program, you can configure the filer to work with other clients in the same way as it does with the administration host.

appliance

A device that performs a well-defined function and is easy to install and operate.

ATM

Asynchronous Transfer Mode is a network technology that combines the features of cell-switching and multiplexing to offer reliable and efficient network services. ATM provides an interface between devices such as workstations and routers, and the network.

authentication

A security step performed by a domain controller for the filer's domain, or by the filer itself, using its */etc/passwd* file.

autosupport

A filer daemon that triggers email messages from the customer site to Network Appliance or another specified email recipient when there is a potential filer problem.

big-endian A binary data format for storage and transmission in which the most significant bit or byte comes first.

CIFS Common Internet File System. A protocol for networking PCs.

client A computer that shares files on a filer.

console A terminal that is attached to a filer's serial port and is used to monitor and manage filer operation.

copy-on-write The technique for creating snapshots without consuming excess disk space.

degraded mode The operating mode of a filer when a disk is missing from the RAID array or the batteries on the NVRAM card are low.

disk ID number A number assigned by the filer to each disk when it probes the disks at boot time.

disk shelf A shelf that contains disk drives and is attached to the filer.

Ethernet adapter An Ethernet interface card.

expansion card A SCSI card, NVRAM card, network card, hot swap card, or console card that plugs into a filer expansion slot.

expansion slot The slots on the system board in which you insert expansion cards.

FDDI adapter A Fiber Distributed Data Interface (FDDI) interface card.

FDDI-fiber An FDDI adapter that supports a fiber-optic cable.

FDDI-TP An FDDI adapter that supports a twisted-pair cable.

GID Group identification number.

group A group of users defined in the filer's */etc/group* file.

hot spare disk A disk installed in the filer that can be used to substitute for a failed disk. Before the disk failure, the hot spare disk is not part of the RAID disk array.

hot swap The process of adding, removing, or replacing a disk while the filer is running.

hot swap adapter An expansion card that makes it possible to add or remove a hard disk with minimal interruption to file system activity.

inode A data structure containing information about files on a filer and in a UNIX file system.

interrupt switch A switch on some filer front panels used for debugging purposes.

LAN Emulation (LANE) The architecture, protocols, and services that create an Emulated LAN using ATM as an underlying network topology. LAN Emulation enables ATM-connected end systems to communicate with other LAN-based systems.

magic directory A directory that can be accessed by name but does not show up in a directory listing. The *.snapshot* directories, except for the one at the mount point or at the root of the share, are magic directories.

mail host The client host responsible for sending automatic email to Network Appliance when certain filer events occur.

Maintenance mode An option when booting a filer from a system boot disk. Maintenance mode provides special commands for troubleshooting hardware and configuration.

MIME Multipurpose Internet Mail Extensions. A specification that defines the mechanisms for specifying and describing the format of Internet message bodies. An HTTP response containing the MIME Content-Type header allows the HTTP client to invoke the application that is appropriate for the data received.

NDMP Network Data Management Protocol. A protocol that allows Network Appliance filers to communicate with backup applications, and provides capabilities for controlling the robotics of multiple tape backup devices.

network adapter An Ethernet, FDDI, or Asynchronous Transfer Mode (ATM) adapter.

NVRAM cache Nonvolatile RAM in the filer, used for logging incoming write data and NFS requests. Improves system performance and prevents loss of data in case of a filer or power failure.

NVRAM card Adapter card that contains the filer’s NVRAM cache.

panic A serious error condition causing the filer to halt. Similar to a software crash in the Windows system environment.

parity disk Disk on which parity information is stored for the RAID-4 disk drive array. Used to reconstruct data in failed disk blocks or on a failed disk.

PCI Peripheral Component Interconnect. The bus architecture used in newer filer models.

pcnfsd A filer daemon that permits PCs to mount filer file systems. The corresponding PC client software is called (PC)NFS.

PDC Primary Domain Controller. The domain controller that has negotiated to be, or has been assigned as, the primary authentication server for the domain.

POST Power-on self-tests. The tests run by the filer after the power is turned on.

PVC Permanent Virtual Circuit. A link with a static route defined in advance, usually by manual setup.

qtree A directory on which you can impose tree quotas, created by the `quota qtree` command. Formerly known as quota trees.

RAID Redundant Array of Independent Disks. A technique that protects against disk failure by computing parity information based on the contents of all the disks in the array. The filer uses RAID Level 4, which stores all parity information on a single disk.

RAID disk scrubbing The process in which the system reads each disk in the RAID group and tries to fix media errors by rewriting the data to another disk area.

SCSI adapter An expansion card that supports the SCSI disk drives and tape drives.

SCSI ID The number of a disk drive on the SCSI chain (0-6).

SCSI address The full address of a disk, consisting of the disk's SCSI adapter number and the disk's SCSI ID; for example, 9a.1.

serial adapter An expansion card for attaching a terminal as the console on some filer models.

serial console An ASCII or ANSI terminal attached to a filer's serial port. Used to monitor and manage filer operations.

share A directory or directory structure on the filer that has been made available to network users and can be mapped to a drive letter on a CIFS client.

SID Security identifier.

snapshot An on-line, read-only copy of the entire file system that protects against accidental deletions or modifications of files without duplicating file contents. Snapshots enable users to restore files and enable you to back up the filer to tape while the filer is in use.

SVC Switched Virtual Circuit. A connection established through signaling. The user defines the endpoints when the call is initiated.

system board A printed circuit board that contains the filer's CPU, expansion bus slots, and system memory.

tree quota A type of disk quota that restricts the disk usage of a directory created by the `quota qtree` command. Different from user and group quotas that restrict disk usage by files with a given UID or GID.

UID User identification number.

Unicode A 16-bit character set standard. It was designed and is maintained by the nonprofit consortium Unicode Inc.

VCI Virtual Channel Identifier. A unique numerical tag defined by a 16-bit field in the ATM cell header that identifies a virtual channel over which the cell is to travel.

VGA adapter Expansion card for attaching a VGA terminal as the console.

VPI Virtual Path Identifier. An eight-bit field in the ATM cell header that indicates the virtual path over which the cell should be routed.

WAFL Write Anywhere File Layout. The WAFL™ file system was designed for the Network Appliance filer to optimize write performance.

WINS Windows Internet Name Service.

workgroup A collection of computers running Windows NT or Windows for Workgroups that is grouped for browsing and sharing.

Index

Symbols

/etc/crash directory 41, 47
/etc/dgateways file 123, 125
/etc/exports file 41
/etc/fstab file
    balancing traffic among interfaces 561
    editing 561
/etc/hosts file 507
    accessing problems in 583
    contents of 118
    for IP address 40
/etc/hosts.equiv file 22, 507
/etc/httpd.access file 342
/etc/httpd.group file 343
/etc/httpd.hostprefixes file
    for virtual hosting 346
/etc/httpd.passwd file 343
/etc/messages file
    and files saved 51
    and messages from savecore 47
    and reporting to Network Appliance 566
    and warnings for disk failures 77
    configuring messages written to the 49–51
    for disk failures 577
/etc/netgroup file 247–248
/etc/nsswitch.conf file 36, 117
/etc/nvfail_rename file 154
/etc/passwd file, copying for (PC)NFS 254
/etc/quotas file, configuring 426–430
/etc/rc
    vif commands 142
/etc/rc file
    commands in 39–41
    damaged or missing 584
/etc/resolv.conf file 120
/etc/serialnum file 36
/etc/shadow file 36, 254
/etc/symlink.translations file 193
/etc/syslog.conf file 49–51
/etc/zoneinfo directory 41

Numerics

8.3 format for file names 198

A

access events, CIFS, auditing 313–325
access logging, CIFS
    disabling 319
    enabling 318
access rights
    assigning rights to users 302–306
    CIFS shares 302–306
active event log, CIFS
    default 313
    specifying 320
adding a foreign volume 91
adding disks to a volume 76, 89
adding ELAN to the adapter 174
adding multiple ELANs to the adapter 175
address resolution 162
adinfo
    command 167
    interpreting the output 168
    sample output 167
    syntax 167
administration host
    accessibility problems from 583
    modifying configuration files 35
    requirements for 17
administration tasks, periodic 13
administrative users
    changing password 29
    creation of 25
adstat
    command 168
    interpreting the output 169
    sample output 169
    syntax 168
applications, LAN 161, 162
ATM
    address 171
    cells 160

differences between LANs and ATM 161
FORE OC3 network interface 164
introducing 160
network mask for 40
primary address 170
switch 163, 166, 172
Switched Virtual Circuit 162
topology 160
ATM adapter
adding ELAN 174
adding multiple ELANs 175
adinfo command 167
adstat command 168
checking proper configuration on an ELAN 172
checking working physical connection 168
configuring additional 182
configuring the adapter 165
configuring the Ethernet interface 175
configuring the LECS address 171
correcting ELAN addition mistakes 176
correcting ELAN mistakes after IP address is assigned 176
passing self-test procedure 168
saving the configuration information 181
uniconfig command 169
unit numbers 168
ATM adapter on an ELAN, configuration 172
ATM address 171, 173
ATM Forum LANE 164
ATM LANE
bridging between ATM networks and LANs 161
Ethernet services 164
saving the configuration 181
services 161
SNMP/ILMI MIB 164
standards supported 164
AuthName directive 342
automatic email to Network Appliance
configuring 57–58
contents of the messages 56
designating the contact person 42
using administration host for 16
autosupport daemon
configuring with the options command 57–58
how it works 55
autosupport options 618–619
autosupport.doit option 58, 618
autosupport.enable option 57, 618
autosupport.from option 58, 618
autosupport.mailhost option 58, 618
autosupport.noteto option 619
autosupport.to option 58, 619

B

backup
amount of data 451
data format 447
data not in qtrees 462
devices 444
different methods 441
excluding files 462
incremental 445
metadata 444
name in the /etc/dumpdates file 462
qtrees 451
reasons for 441
specifying files and directories 461
use of nonqualified tape drives 455
Windows NT ACL information 448
blocking factor
definition 449
how to specify 461
specifying when restoring data 473
when backing up to remote tape drives 455
booting
/etc/rc script 38–41
and NVRAM 31
from diskette 567–569
with halt and reboot commands 31
with nvfail enabled 154, 155
without /etc/rc file 568
bridging between ATM networks and LANs 161
Broadcast and Unknown Server (BUS)
configuring 166
ELAN 161
Emulated LAN (ELAN) 163
explained 163

LAN Emulation Client (LEC) 163
LAN Emulation Service 163
broadcast filer address, setting using ifconfig 127
broadcast MAC address 163
broadcast, LAN 161, 163

C

C$ share 21
cables, checking 587
cf disable command 377, 382
cf enable command 377, 382
cf forcegiveback command 382, 387
cf forcetakeover command 378, 382
cf giveback command 372, 382, 386
cf giveback -f command 386
cf partner command 375
cf status command 375
cf takeover command 369, 377, 382
cf.timed.enable option 376
CGI requests, redirecting 350
changing size of RAID groups 85
character set types, supported 200
checking the ATM LANE configuration 177
CIFS
access event displays 314
access event log 313
access logging
disabling 319
enabling 318
active event log
default 313
specifying 320
adding users 275–277
assigning and changing access rights 302–306
C$ share 21
creating shares 287
deleting a share 297
displaying session information 331–332
displaying share information 294–296
displaying shares 284–286
displaying statistics 330
effect of filer reboot on clients 32
event auditing configuration 317
events, viewing 325–327
file name case 216
file names, preserving case 218
filer command-line only operations 267
generic account
creation by default 282
users 283
guest access 282
home directory shares, creating 299–301
local groups, adding to filer 278–280
login tracing, toggling 232–233
lost record event detail display 316
open files, limits 271
oplocks 328–329
options 620–623
restoring files 472
rsh use to enter filer commands 281
session information, displaying 331–332
sessions, limits 271
shares
changing description 295
creating and changing 287–293
deleting 297–298
displaying 284–286
displaying information 294–296
limits 271
renaming volume, effect on 269
statistics, displaying 330
symbolic links 192–194
system ACL (SACL), setting 321–324
UNIX file access detail displays 315
Unsuccessful file access detail display 315
Windows file access detail displays 314–315
cifs access command 302, 305
cifs access -delete command 306
CIFS guest access 282
CIFS guest account 282
CIFS login tracing, toggling 232–233
cifs restart command 336
CIFS sessions
starting 333–336
stopping 333–336
cifs sessions command 331–332
cifs shares -add command 290
cifs shares -change command 292
cifs shares command 284, 285, 287, 295–296, 333

cifs shares -delete command 298
cifs stat command 330
cifs terminate command 334, 335
cifs.access_logging.enable option 318, 319, 320
cifs.access_logging.filename option 620
cifs.access_logging.filernamе option 320
cifs.bypass_traverse_checking option 620
cifs.cifs.show_snapshot option 622
cifs.guest_account option 620
cifs.home_dir option 300, 621
cifs.netbios_aliases option 621
cifs.oplocks.enable option 621
cifs.perm_check_use_gid option 621
cifs.save_case option 217, 218
cifs.scopeid option 622
cifs.search_domains option 622
cifs.symlinks.cycleguard option 194, 622
cifs.symlinks.enable option 193, 623
clients
    accessing snapshots from 408–410
    NFS statistics in the custom MIB 113
cloned_tapes file, format 488
cluster
    commands that are different in partner mode 383–384
    commands unavailable in partner mode 382
    configuration information 375–376
    console messages 602
    defined 355
    dump on virtual filer 384
    giveback 372–373, 386–387
    hardware components 357
    interconnect 358
    LCD messages 597–598
    management tasks 374
    monitor 358
    monitoring process 360
    normal mode 375–376
    partner mode 381–384
    partner name 375
    statistics in takeover mode 379–380
    status 375
    syslog messages 599–601
    takeover
        initiating 377–378
        managing 377
        process 369
        result 371
        understanding 369–370
    takeover mode 379–380
    terminology 358–359
    time synchronization 360, 376
    unlicensing 376
collisions 588
commands
    aborted by Ctrl-C 23
    accessible through rsh 23
    adinfo 167
    adstat 168
    elconfig 168, 171
    elconfig add 174
    elconfig delete 176
    elconfig show 172, 178
    ifconfig 175, 176
    ping 177
    unavailable in partner mode 382
    uniconfig 169
    useradmin 24
communication link, checking 177
configuration
    of volumes 81
    planning for multiple volumes 82
configuration files
    /etc/dgateways 125
    /etc/dgateways file 123
    /etc/hosts 118–119
    /etc/hosts.equiv file 22
    /etc/netgroup file 247
    /etc/nvfail_rename 154, 157
    /etc/rc 39–41
    /etc/resolv.conf file 120
    /etc/syslog.conf 49–51
    default 33–41
    editing 19
    errors in 584
configuration files, accessing 3
configuration problems
    booting with diskette for 584
    filer accessibility 583
    lost passwords 585

with /etc/rc file 583–584
configuration settings, checking 178, 181
configuring additional ATM adapters 182
configuring BUS 166
configuring LECS 166
configuring LES 166
configuring the /etc/hosts file 182
configuring the ATM adapter 165
configuring the Ethernet interface 175
configuring the LECS address 171
connectionless service 161
console encoding, setting 206
console messages, cluster 602
console session 22
console.encoding option 206, 635
copying a volume
changing the speed of 515
how to stop 515
possible errors during 512
recommendation for 508
requirements for 506
copy-on-write technique 390, 392
core files 47
cpio, copying files with 595
crash files 41
creating
administrative users 25
nvfail_rename file 157
creating an ELAN 166
creating and changing shares 287–293
creating qtrees 420
creating volumes 89

D

data access management 11
Data ONTAP
displaying version 544
initial volume configuration 80
overview 8–14
data rebuild, on the hot spare disk 578
data reconstruction
speed 78
when filer is shut down (degraded mode) 78
database file protection 154
enabling/disabling nvfail 156
date command 60
date date command, unavailable on virtual filer 382
default
configuration 33–37
route in routing table 123
router 40, 123
degraded mode 77
meaning of 87
reasons for 78
timeout period for automatic shutdown 78
when a hot spare disk is available 78
when a hot spare disk is not available 77
deprecated MIB objects 113
destroying a volume 92
df command 399–401, 438–439, 596
diagnostic messages 49–51, 566
differences between LANs and ATM 161
directories
conversion time 219
conversion to Unicode 211
created by snapshots 408–409
default format 210
exporting 41
reversion to pre-Data ONTAP 4.0 default creation 213
sizes 210
Unicode conversion upon CIFS access, reverting to 215
Unicode creation forcing 212
Unicode format conversion upon any access 214, 220
Directory directive 342
disk does not exist message 579
disk fail command 87
disk in use message 579
disk information, displaying 545, 547
disk remove command 87
disk scrub command 74
disk shelves 7
disks
adding 76
addressing, use of 73
changing the size of RAID groups 85
concepts 72

data 9
degraded mode 77, 78, 87
different sizes in same filer 75
different types in RAID group 72
failures, effects of 78
failures, handling 77
free space, accessing information through SNMP 113
free space, displaying 438–439
freeing space by deleting snapshots 406
hot swapping 75
information in the custom MIB 113
installing new 86
management tasks 85
maximum number of files 437
parity 9, 72, 438
problems 576–581
quotas 66, 114, 402, 426–435
removing 87
replacing 76
restricting usage 426–436
right-sizing 76
SCSI ID number 73
setting size of RAID group 85
snapshots, space used 393–407
spare 9
swapping 87
usable space 76
DNS
disabling 120
domain name in /etc/rc 40
enabling 120
options 624
querying the name server 119
resolving names with 117
dns.enable option 624
DOS attributes, changing from Windows NT 308
DOS file names
forcing to lowercase 217
double disk failures 578
dump command
data format 447
devices 444
different passes 447
effect on mirror update schedule 526
examples 463–469
exclude list 458, 462
excluding certain types of data 444
how it works 443
how it works with SnapMirror 526
incremental backup 445
multiple tape files 450
options 461
syntax 460
using snapshot 390
where to enter 445
Windows NT ACL information 448, 461
dump command virtual filer 384
dump level 461
dumpdates file, updating 462

E

editing configuration files 19
ELAN
adding multiple ELANs to the adapter 175
adding to the adapter 174
address resolution 162
checking elements 178
correcting addition mistakes 176
correcting mistakes after IP address is assigned 176
creating 166
explained 161
LE Service 161
multiple 161, 166
services 161
elconfig
command 168
example 171
syntax 171
elconfig add
command 174
example 174
interpreting the output 175
syntax 174
elconfig delete
command 176
example 176
syntax 176

elconfig show
example 172, 178, 179, 181
interpreting the output 172, 178
syntax 172, 178, 179, 180
environmental adapter, filer main unit 6
error messages, serious 603
errors
caused by copying a volume 512
caused by exceeding disk quotas 435
displayed by netstat 588
displayed by nvfail 155
Ethernet
configuring logical interface 175
interface names and numbers 44
network mask for 40
setting media type on 128
event auditing, CIFS, configuration 317
event log, CIFS
access 313
default 313
events, CIFS, viewing 325–327
explicit routes in routing table 123
exportfs command 41

F

failed filer
defined 358
FAX support 566
FC-AL disk shelf 358
FDDI
interface names and numbers 44
network mask for 40
setting the MTU 127–128
file locking, differences between NFS and CIFS 191
file names
CIFS, preserving case 218
conversion 197
DOS, forcing to lowercase 217
legal characters 198
maximum length 197
used by NFS and CIFS clients 198
file space, incorrect display 596
file system
inconsistent 581
maximum number of files 437
protection through RAID scrubbing 74
filer
characteristics in a cluster 355
description, changing and viewing 273–274
halting without takeover 377
restarting 570
virtual 381–384
accessing through a remote shell 382
accessing through Telnet 382
dump 384
dump command 384
restoring files 384
filer information, overall, displaying 546
filer main unit components 5
filer system load
systat command 79
FilerView, use for filer administration 4
files
copying with cpio 595
large 110
maximum size 110
ownership
changeable by root only 66
same file criteria 391
setting maximum number of 437
working with large 110
filestats command 93
firewall, virtual 345
FORE OC3 network interface 164
FORE/IP PVCs
incoming on the filer 185
outgoing, establishing on the filer 184
tearing down 186

G

generic account
creation by default 282
users 283
giveback
defined 358
performing 386–387
understanding 372–373

group quotas, creating 434
guest
    access 282
guest account, CIFS 282

H

halt command 31, 382, 571
halt -f command 377
halting a filer without takeover 377
heartbeat
    defined 358
    transmission 360
home directory shares, CIFS, creating 299–301
host key, SecureAdmin 607
host name resolution 117
host names
    for network interfaces 43
    resolving 118–122
hostname command 39
hosts file 182
hot spare disk 86
    availability
        sysconfig command 78
    overview 75
    removing 87
    replacement activity (/etc/messages) 78
hot swapping a disk 75
hourly snapshots 396, 397
HTTP
    displaying connection information 353
    displaying statistics 354
    options 625–626
    password protection 342–343
    root directory 341
    starting service 340
    virtual hosting, enabling 346
httpd.admin.enable option 625
httpd.enable option 625
httpd.log.max_file_size option 625
httpd.rootdir option 625
httpd.timeout option 625
httpd.timewait.enable option 626
httpstat statistics in takeover mode 379–380

I

ICMP redirect messages 123
identifying disks
    SCSI ID 73
IEEE 802.3 162
IERRS (input errors) displayed by netstat 588
ifconfig command 39, 40, 118, 127–131, 175, 176, 383, 587
ifconfig mtusize command 382
ifconfig -partner command 382
ifconfig partner command 382
ifstat command 137, 379–380, 551
illegal volume name message 575
ILMI address registration 164
inconsistent file system 581
inodes, effects of maximum number of files on 437
installing new disks 86
interfaces
    errors on 588
    how packets are sent and received 125
    using ifconfig to configure 127–131
invalid volume name message 575
IP addresses
    setting in /etc/rc file 40
    setting using ifconfig 127
ip.match_any_ifaddr option 635
ip.path_mtu_discovery.enable option 635

L

LAN
    address resolution 162
    applications 161, 162
    bridging between ATM networks and LANs 161
    broadcast 161, 163
    broadcast MAC address 163
    connectionless service 161
    differences between ATM and 161
    MAC address 161, 162
    multicast 161, 163
    shared medium characteristics 163
    Token Ring 161
    Token Ring emulation 164
    unresolved unicast frames 163

LAN Emulation Address Resolution Protocol (LE_ARP) 162
LAN Emulation Client (LEC)
    ELAN 162
    explained 162
    joining an ELAN 163, 171
    MAC layer interface 162
LAN Emulation Configuration Server (LECS) 163
    address 166, 172
    checking communication with the adapter 172
    configuring 166
    configuring the address for an ATM adapter 171
    ELAN 161
    explained 163
    LAN Emulation Client (LEC) 163
    LAN Emulation Server (LES) 163
    well-known address 166
LAN Emulation Server (LES) 163
    ATM address 173
    configuring 166
    explained 163
    registering MAC addresses 163
    resolving a route descriptor 163
LAN Emulation Service
    BUS 161, 163
    explained 161, 163
    LECS 161, 163
    LES 161, 163
LAN Emulation, explained 161
languages, supported
    displaying list 205
    list 202
large files 110
LCD messages, cluster 597–598
LCD, filer main unit 5
legal characters in file names 198
license command 67, 69, 382
licenses 67–69
local filer 358
local groups, CIFS, adding to filer 278–280
localhost 507
logging in to the filer 22
logging out of the filer 22
lost data from disk failures 577
lost passwords 585
lost record event detail display, CIFS 316
ls command, listing snapshot files 409, 410

M

MAC address 161, 162
mail host
    automatic email 16, 55
    designating other mail hosts 58
mailbox disk 359
Makefile, NIS 119, 248
making a volume inactive 90
management tasks
    for disks 85
    for volumes 89
maxfiles command 437
maximum number of files 437
media type for an Ethernet interface 128
messages
    configuring the syslogd daemon 49–51
    console, cluster 602
    file of 49–51
    for disk failures 577
    severity of 49
    syslog, cluster 599–601
MIB objects
    deprecated 113
    multivolume, locations 113
MIB, Network Appliance custom 112
MIB-II 111
MIME Content-Type, specifying 348
minra option 562
minra volume option 633
modes
    normal 375–376
    normal, clusters 359
    partner 381–384
    takeover 379–380
monitoring status
    of volumes 90
mount_rootonly option 66
mounting files
    if there are qtrees 439
    problems with 589

restricting mount privilege to root 66
mounting volumes 81
mt command 383, 492
multicast 161, 163
multiple ELANs 161, 166
multiple RAID groups 72
multiple volumes
configuration planning 82
limitations of 82
multiple-mode trunks
creating 146
trunks
multiple-mode 135
multivolume MIB objects, locations 113

N

name services, specifying the order in which contacted 117
names
resolving 117–122
volume naming conventions 80
NDMP
advantages of 495
log files for 500
tape devices used 496
ndmp command 498
netstat
statistics different in takeover mode 379–380
netstat command 123, 126, 383, 550, 588
network interfaces
balancing traffic among 561
configuring 127–131
naming conventions 43
network statistics, displaying 550
networks
connections, checking 587
how filer sends and receives traffic 125
management services, using SNMP 111–114
network mask
configuring using ifconfig 127
in /etc/rc file 40
problems with 587
statistics 587
using ifconfig to configure 127
NFS
how interfaces respond to packets 125
problems with 589
statistics in custom MIB 112
statistics, displaying (nfsstat command) 262
turning on 41
NFS and (PC)NFS options 627–628
nfs command 41
NFS guest access 282
NFS over UDP requests 125
nfs.big_endianize_fileid option 562
nfs.mount_rootonly option 627
nfs.per_client_stats.enable option 627
nfs.tcp.enable option 627
nfs.v2.df.2gb.lim option 627
nfs.v2.df_2gb_lim option 596, 627
nfs.v3.enable option 628
nfs.webnfs.enable option 628
nfs.webnfs.rootdir option 628
nfs.webnfs.rootdirset option 628
nfsstat command 262–265, 379–380
nightly snapshots 395–397
NIS
changing domain name 122
disabling 122
domain name
set in /etc/rc 41
specifying with option 629
enabled in /etc/rc 41
enabling
during setup 121
without using setup 122
maps supported 121
options 629
propagating changes
/etc/hosts on filer 119
/etc/netgroup on filer 248
/etc/passwd on filer 255
/etc/shadow for (PC)NFS on filer 255
nis.domainname option 629
nis.enable option 629
no_atime_update option 561
no_atime_update volume option 633
nonexistent disks 579
nonqualified tape drives 455, 488

nonvolatile RAM (NVRAM)
batteries 78, 571
failures in 566
inconsistent contents 571
logging requests 31
mirroring, cluster 359, 360
when to increase 548
normal mode, clusters 359, 375–376
nosnap volume option 633
nosnapdir option 394, 395
nosnapdir volume option 633
numbering network interfaces 43
nvfail
bootup process 154, 155
database file protection 154
enabling and disabling 156
error message 155
volume option 154
nvfail option 634
nvfail_rename file, creation of 157
NVRAM 5

O

oplocks 328–329
options
autosupport.doit 58, 618
autosupport.enable 57, 618
autosupport.from 58, 618
autosupport.mailhost 58, 618
autosupport.noteto 619
autosupport.to 58, 619
cf.timed.enable 376
cifs.access_logging.enable 318, 319, 320
cifs.access_logging.filename 320
cifs.guest_account 620
cifs.home_dir 300, 621
cifs.netbios_aliases 621
cifs.oplocks.enable 329, 621
cifs.save_case 217, 218
cifs.scopeid 622
cifs.show_snapshot 622
cifs.symlinks.cycleguard 194, 622
cifs.symlinks.enable 193, 623
console.encoding 206
dns.domainname 624
dns.enable 624
httpd.admin.enable 625
httpd.enable 340, 625
httpd.log.max_file_size 340, 625
httpd.rootdir 340, 346, 625
httpd.timeout 625
httpd.timewait.enable 626
ip.path_mtu_discovery 635
minra 562
mount_rootonly 66
nfs.big_endianize_fileid 562
nfs.mount_rootonly 627
nfs.per_client_stats.enable 627
nfs.tcp.enable 627
nfs.v2.df_2gb_lim 596, 627
nfs.v3.enable 628
nfs.webnfs.enable 628
nfs.webnfs.rootdir 628
nfs.webnfs.rootdirset 628
nis.domainname 629
nis.enable 629
no_atime_update 561
nosnapdir 394, 395
pcnfsd.enable 628
pcnfsd.umask 628
raid.reconstruct_speed 78, 630
raid.scrub.enable 74, 630
raidtimeout 77, 78, 630
root_only_chown 66
server 63
telnet.hosts 66, 636
timed.enable 63, 64
timed.log 63, 64
timed.max_skew 64
timed.proto 63, 65
timed.sched 63, 65
timed.servers 65
timed.skew 63
vol.copy.throttle 516
volume options
minra 633
no_atime_update 633
nosnap 633
nosnapdir 633

raidsize 634
root 634
vol.copy.throttle 636
wafl.convert_ucode 214, 215
wafl.create_ucode 212, 213
wafl.maxdirsize 561, 638
wafl.root_only_chown 638
wafl.wcc_minutes_valid 222, 225
options command
making changes permanent 53
using 52

P

packets, responses to 125
panic messages 603
parity disks 72
role 73
partner 359
partner command 381
partner mode 359, 381–384
commands that are different in 383–384
unavailable commands 382
passwd command 23
password
HTTP pages 342–343
maximum length when used with MD5 authentication 501
passwords
changing administrative user 29
changing with passwd command 23
lost 585
used for shell session 22
(PC)NFS 251–255
pcnfsd daemon 253–254
pcnfsd.enable option 628
pcnfsd.umask option 628
performance, improving 548
permissions
changing from Windows NT 308
for snapshots 391
on exported directories 240
ping command 177, 587
primary ATM address 170
protocol licenses 67–69
public-key encryption used by SecureAdmin 606
PVCs, FORE/IP
establishing on a remote ATM host 185, 186
establishing on interconnecting ATM switches 185
incoming, establishing on the filer 185
outgoing, establishing on the filer 184

Q

qtrees
administrative actions 10
backup of 451
creating 420
description 10
displaying information about 424
moving files in and out of 414
oplocks settings 422
parameters 411
restoring 457, 475, 478
security style 415–417
use for backups 413
use in projects 413
what they are 411
quota
command 430
disk, setting up 426–435
displaying report 433
effects of snapshots on 402
effects on clients when exceeded 435
how to impose on a former mirror 525
resizing 432
quotas, disk
information available through SNMP 114
information available through the custom MIB 114
when changing file ownership 66

R

RAID (Redundant Array of Independent Disks)
accessing information through SNMP 114
adding disks 76
data reconstruction speed 78
disk scrubbing 74
displaying information about 545

group size characteristics 72
groups 8, 72, 80
information in the custom MIB 114
options 630
spare disk use 72
support for multiple groups 72
RAID (Redundant Array of Independent Disks), displaying information about 547
raid.reconstruct_speed option 78, 630
raid.scrub.enable option 74, 630
raidsize volume option 634
raidtimeout command 78
raidtimeout option 77, 630
rc file 181
rdate command 382
/hosts/equiv file required line 62
accuracy 60
use with cron and crontab 61
read-ahead, minimal 562
read-only bit 195–196
reboot command 31, 382
rebooting
and NVRAM 31
automatic email to Network Appliance 48
from diskette 567–582
scripts for 39
with halt and reboot commands 31
recovering data, different methods 472
registering MAC addresses 163
remote shell (rsh)
accessing the filer from the administration host 22
accessing virtual filer with 382
for entering filer commands 23
removing
a hot spare disk 87
volumes from a filer 82
renamed volume not exported 574
renaming volumes 92
replacing
disks 76
failed disks 76
replicating a volume 517–542
requests, replies to 125
require group directive 343
require user directive 343
resolving a route descriptor 163
resolving host names 118
restarting a filer 570
restore command
data that cannot be restored 472
examples 479–483
function keys 473
how it differs from ufsrestore 475
incremental 475
options 473
purposes of 471
space required 475
syntax 473
when not to use 471
where to enter 476
Windows NT ACLs 474
restore_symboltable file, removing 476, 479
restoring files from virtual filer 384
reverting a volume 101–107
right-sizing disks 76
role of parity disks 73
root volume 18, 80
option 634
root_only_chown option 66
route command 123
routed command 124
routed daemon
/etc/dgateways file 125
purpose of 124
router
default 40, 123
problems with 588
Routing Information Protocol (RIP) 124
routing table, filer 123
rsh.enable option 636

S

savecore command 41, 47
SCSI ID, identifying disks 73
second-level virtual interfaces
configurations 138
creating 147–149
SecureAdmin

determining whether it is running 614
disabling 612
displaying version 615
keys 607–608
protocol supported 606
public-key encryption 606
reinitializing 610
restarting 613
setting up 609
sshd daemon 606
stopping 612
secureadmin disable ssh command 610, 612
secureadmin enable ssh command 609, 610, 613
secureadmin -f setup command 610
secureadmin setup command 609
secureadmin version command 615
security
options for 66
system password 23
serial ports, filer main unit 6
server key, SecureAdmin 607
session key, SecureAdmin 607
setting up the filer 33
shadow file 36
shared medium characteristics 163
shares
creating and changing 287–293
deleting 297–298
displaying 284–286
displaying information about 294–296
renaming volume, effect on 269
shelfchk command in partner mode 383
single-mode trunk
creating 143
single-mode trunks 134
preferred links 144–145
slots, filer main unit 6
snap command 394–398
SnapMirror
checking status 538
converting a mirror to a regular volume 540
disabling data replication for entire filer 534
disabling data replication for one volume 536
how it works 520
overview 517
procedure for mirroring a volume 530–533
purposes of 518
resuming 535
when used with quotas 525, 540
when used with the dump command 526
snapmirror.allow file format 527
snapmirror.conf file format 528
snapmirrored option 634
SnapRestore
effects of reverting a root volume 103
effects on backup and recovery 104
effects on snapshots 102
how it works 101
interaction with SnapMirror 103
procedure for reverting a volume 105–107
typical applications 102
when to use 102
snapshots
accessing 408–409
automatic 395–397
commands for 394
created by SnapMirror 522
definition 390
deleting to free space 406
directory name displayed on CIFS clients 409
disk consumption by 399–407
effects of SnapRestore 102
effects on disk quotas 402
information in the custom MIB 113
ls command 409, 410
magic directories 408
making snapshot directory invisible 394
operation of 391–392
options 630–634
reserving space for 400–401, 404–405, 406
reverting a volume 105
scheduling 403
snapshot_for_backup file 443
types 395
SNMP
commands
examples 111–112
in /etc/rc 41
configuring the agent 111
custom MIB 112

takeover mode 380
snmp command 111–112
snmp.enable option 636
SNMP/ILMI MIB 164
software licenses 67–69
spare disks in RAID groups 72
SSCOP operational state 170
sshd daemon for SecureAdmin 606
standards
ATM Forum LANE Version 1.0 164
ILMI address registration 164
SNMP/ILMI MIB 164
User Network Interface (UNI) 164
sticky bit 37
subnets, exporting to 249–250
svtx bit 37
swapping out disks 87
symbolic links, CIFS 192–194
sysconfig command
cluster configuration 375–376
different in partner mode 383
filer configuration 545–546
hot spare disk availability 78
usable disk space 76
sysconfig -d command 545
sysconfig -m command 490, 496
sysconfig -r command 545
sysconfig -t command 545
sysconfig -v command 566
syslog messages, cluster 599–601
syslogd daemon 49
sysstat command 548
for filer system load 79
statistics in takeover mode 379–380
system ACL (SACL), setting 321–324
System board, filer main unit 5
system crashes 41
system memory, filer main unit 5
system message files 49–51
system panics 47, 603

T

takeover
defined 359
disabling 377
enabling 377
initiating 377–378
managing 377
process 369
result 371
understanding 369–370
takeover filer 358
takeover mode 379–380
defined 359
prompt 379
statistics 379–380
tape block
definition 449
specifying the size for dump 461
specifying the size when restoring files 473
tape drives
backing up to remote 455
controlling 492–493
displaying information about 545
inaccessible on virtual filer 382
used by NDMP 496
tape file
definition 449
size used by remote host 455
specifying name 461
specifying size for dump 461
specifying when restoring data 474
types of 450
tape stackers, displaying information about 489
tapes, estimating number needed 453
TCP/IP, how interfaces respond to packets 126
Technical Support
address for receiving autosupport email 55
how to contact 566
URL 566
telnet connection to the filer
administration host 17
limiting host access 66, 636
virtual filer 382
telnet.enable option 636
telnet.hosts option 66, 636
The telnet.enable option 636
The wafl.convert_ucode option 636
The wafl.create_ucode option 638

time
date command 60
rdate command 60, 61–62
setting on the filer 60–65
setting the timed options to synchronize filer time 64–65
SNTP (Simple Network Time Protocol) 60
SNTP synchronization 63
synchronization, cluster 360, 376
synchronizing filer with another system 61
time zone, setting 41
timed.enable option 63, 64, 631
timed.log option 63, 64, 631
timed.max_skew option 63, 64, 631
timed.proto option 63, 65, 631
timed.sched option 63, 65, 632
timed.servers option 63, 65, 632
timezone command 41
timezone name command 382
Token Ring 161
topology, ATM 160
troubleshooting
configuration problems 583–585
cpio problems 595
df problems 596
disk problems 576–578
network problems 587
NFS problems 589
NVRAM problems 571–572
UNIX cpio problems 595
UNIX df problems 596
volume problems 573–580
trunks
defined 132
destroying 153
hardware requirements 136
physical interfaces, adding 150
single-mode 134
statistics, displaying 152
status, displaying 151
trusted host 23, 507

U

Unicode format
directory conversion to 211
forcing default directory creation 212
uniconfig
command 169
interpreting the output 170
sample output 169
UNIX access to NTFS files 221
UNIX file access detail displays, CIFS 315
unrecognized volume name message 575
unresolved unicast frames 163
Unsuccessful file access detail display, CIFS 315
uptime command 383, 549
URL, how filer translates 350
usable space on disks 76
user names, defining administrative 24
User Network Interface (UNI)
checking whether operational 167
standards supported 164
user quotas, creating 433
useradmin command 24
users, CIFS, adding to filer 275–277

V

version command 544
version of Data ONTAP, displaying 544
vif command 142
Virtual Channel Connection (VCC) 162
virtual filer
accessing through a remote shell 382
accessing through Telnet 382
defined 359
dump command 384
managing 381–384
restoring files 384
unavailable commands 382
virtual firewall 345
virtual host addresses, mapping 347
virtual hosting
enabling 346
setting up 346
virtual interfaces
defined 137
names 44, 137
second-level 138–141

trunking 138
vol command
add 89
create 85, 89, 208
destroy 92
lang 205, 207
offline 90
online 91
options 90
options raidsize 85
rename 92
status 90, 209
vol copy start command 509–513
vol copy throttle command 516
vol create -l command 208
vol options command 52
vol status command 375–376, 384, 547
vol status -d command 547
vol status -l command 209
vol status -r command 547
vol status -v command 547
vol.copy.throttle option 516, 636
volume
creating with specified language 208
setting language 207
volume copy operation numbers 514, 516
volume copying
checking status of 514
how to stop 515
possible errors during 512
recommendation for 508
requirements for 506
volume options
minra 633
nosnap 633
nosnapdir 633
nvfail 154
raidsize 634
root 634
snapmirrored 90
vol.copy.throttle 636
volumes
adding disks to 89
adding foreign 91
administrative actions 10
and management tasks 89
changing the speed of copying 515
concepts 80
configuration of 81
converting from mirror to regular 540
creating 89
defined 9
destroying 92
differences between mirror and regular 521
displaying file statistics 93
displaying information about 547
displaying language use 209
displaying state information 547
error message about invalid names 575
handling failures 92
making inactive 90
monitoring status 90
mounting 81
naming conventions 18, 80
problems with 573–580
removing 82
renaming 92
renaming and effect on shares 269
replicating 517–542
reverting 101–107
setting options 90
snapmirrored status 521

W

WAFL credential cache
entry, adding 226
entry, deleting 227
managing 222–233
setting how long entry is valid 225
statistics display 228
UNIX name mapping result 230
Windows name mapping result 231
wafl.convert_ucode option 214, 215
wafl.create_ucode option 212, 213
wafl.default_nt_user option 638
wafl.default_unix_user option 638
wafl.maxdirsize option 561, 638
wafl.root_only_chown option 638
wafl.wcc_minutes_valid option 222, 225, 639

warnings for disk failures 77
wcc command 222–224
WebNFS 258–260
weekly snapshots 395
Windows file access detail displays, CIFS event 314–315
Windows NT commands, performing WIndows filer tasks 3

Y

ypwhich command 122