Data ONTAP™ 6.5
System Administration
Storage Management Guide

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

Part number 210-00472
December 2003

Copyright and trademark information

Copyright information

Copyright © 1994–2003 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 by Carnegie Mellon University.

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

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

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

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

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

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

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

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

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

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

Network Appliance reserves the right to change any products described herein at any time, and without notice. Network Appliance assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by Network Appliance. The use 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

NetApp and the Network Appliance design are registered trademarks of Network Appliance, Inc. in the United States, Canada, and the European Union. Network Appliance is a registered trademark of Network Appliance, Inc. in Monaco and a trademark of Network Appliance, Inc. in the United States and Canada. FAServer is a registered trademark of Network Appliance, Inc. in the United States and the European Union. NetCache is a registered trademark of Network Appliance, Inc. in the United States, Canada, the European Union, and Japan. SnapCopy is a registered trademark of Network Appliance, Inc. in the European Union and a trademark of Network Appliance, Inc. in the United States. WAFL is a registered trademark of Network Appliance, Inc. in the United States, the European Union, and Canada. DataFabric, FilerView, NearStore, SecureShare, SnapManager, SnapMirror, and SnapRestore are registered trademarks of Network Appliance, Inc. in the United States. Data ONTAP is a trademark of Network Appliance, Inc. in the United States and Canada. Snapshot is a trademark of Network Appliance, Inc. in the United States and the European Union. NetApp–the Network Appliance Company is a registered trademark of Network Appliance, Inc. in the United States and other countries. ApplianceWatch, BareMetal, Center-to-Edge, ContentDirector, gFiler, MultiStore, SecureAdmin, Serving Data by Design, Smart SAN, SnapCache, SnapDrive, SnapVault, vFiler, and Web Filer are trademarks of Network Appliance, Inc. in the United States.

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

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

Table of Contents

Preface xiii

Chapter 1 **Introduction to Network Appliance Filers** 1

About filers 2

How you administer a filer 6

About filer main unit components 8

About disk shelves 10

About Data ONTAP 11

Filer administration with Data ONTAP 15

Getting command information in Data ONTAP 18

Documentation for Data ONTAP 19

Chapter 2 **Introduction to Filer Administration** 21

Managing access from administration hosts 22

Managing local user accounts 25

Accessing a filer to administer it 32
- Accessing a filer from the console or a Telnet session 33
- Accessing a filer with a Remote Shell connection 36
- Accessing a filer with FilerView 39

Editing command-line commands 43

Using Data ONTAP commands at different privilege levels 45

Receiving automatic technical support 47
- Understanding the Autosupport feature 48
- Events that trigger Autosupport messages 49
- Contents of Autosupport messages 50
- Configuring Autosupport 53

Understanding the filer's default volume 58
- Default directories in the default volume 59
- Contents of the /etc directory 60
- Editing configuration files 62

Configuring filer startup 64

Understanding the options and vol options commands 67

About the options command . . . 68
About the vol options command . . . 69

Performing common tasks . . . 70
Halting the filer . . . 71
Booting and rebooting the filer . . . 73
Editing the filer's boot configuration file . . . 74
Accessing the default directories on the filer . . . 75
Recovering from /etc/rc errors . . . 77
Synchronizing filer system time . . . 78
Setting the filer date and time . . . 83
Managing licenses . . . 85
Understanding core files . . . 87
Understanding large files . . . 89
Configuring message and event logging . . . 90
Configuring audit logging . . . 94
Configuring Telnet sessions . . . 95
Maintaining filer security through options . . . 96

Filer configuration backup and cloning . . . 102

Chapter 3 **Data Organization Using Volumes and Qtrees** . . . 105

Understanding data management . . . 106
Understanding volumes and qtrees . . . 107
Understanding a volume's RAID structure . . . 109
Volume structure . . . 110

Creating volumes . . . 112

Choosing a language for a volume . . . 117

Changing the language of a volume . . . 120

Changing the state of a volume . . . 122

Physically moving volumes . . . 125

Renaming volumes . . . 128

Changing the root volume . . . 130

Destroying volumes . . . 131

SnapLock storage solution . . . 133

Understanding qtree creation . . . 136

Creating qtrees . . . 138

Understanding qtree and volume security styles . . . 139

Understanding file access and security style 141
Changing the security style of a volume or a qtree 143
Changing the CIFS oplocks setting 145
Displaying qtree status 147
Making a directory into a qtree or a qtree into a directory 148
Deleting qtrees 151
Increasing the maximum number of files in a volume 152
Setting space reservation for files 154

Chapter 4 **RAID Group Management** 159
Understanding RAID groups 160
Configuring RAID type and group size while creating a volume 165
How RAID handles disk failures 168
Disk failure with a hot spare disk 169
Disk failure without a hot spare disk 171
Modifying the RAID type for a volume 172
Modifying the size of RAID groups 176
Controlling the speed of RAID operations 179
Controlling the speed of RAID data reconstruction 181
Controlling the speed of disk scrubbing 182
Controlling the speed of plex resynchronization 184
Controlling the speed of mirror verification 185
Scheduling or manually running disk scrubbing 186
Checking and correcting parity 190
Minimizing media error disruption of RAID reconstructions 193
Handling of media errors during RAID reconstruction 194
Continuous media scrub 195
Disk media error failure thresholds 198
Viewing RAID status 199

Chapter 5 **Disk Management** 201
Understanding disks on the filer 202
Multipath I/O for Fibre Channel disks 205

Storage subsystem management 208
Changing the state of a host adapter 209
Viewing storage subsystem information 211
Adding disks 221
Removing disks 227
Ascertaining free disk space 231
Disk sanitization 233
Disk ownership assignment on an FAS270/FAS270c appliance 247

Chapter 6 **User Information Maintenance** 253
About maintaining user information 254
Using configuration files to maintain user information 255
Understanding configuration files 256
Accessing the configuration files 259
Editing the configuration files 261
Using NIS to maintain user information 266
Changing the user name search order 268
Using LDAP to maintain user information 270

Chapter 7 **Cluster Management** 277
About clusters 278
How a clustered system works 281
Cluster operation modes 282
Normal operation mode 283
Takeover operation mode 285
Giveback operation mode 289
Understanding interface configurations for takeover 291
Configuring dedicated and standby interfaces 294
Configuring mirrored, dual-MAC interfaces 297
Enabling cluster failover in a DAFS configuration 300
Designating failover addresses in etc/hosts 302
Enabling failover paths on the filer 305
Managing a cluster 307
Managing a cluster in normal mode 308

Enabling negotiated failover between filers 316
Managing a cluster in takeover mode 318
Managing an emulated filer 320
Backing up and restoring from the emulated filer 325
Performing a giveback 326
Enabling automatic giveback. 329
Configuring automatic giveback 330

Checking cluster configuration 332

Chapter 8 **Disk Space Management Using Quotas** 333

Understanding quotas. 334

When quotas take effect 337

Understanding default quotas. 338

Understanding derived quotas 339

How Data ONTAP identifies users. 342

Understanding the /etc/quotas file 345
Overview of the /etc/quotas file 346
Fields of the /etc/quotas file 349
Sample quota entries 354
Special entries for mapping users 356
How disk space owned by default users is counted 360

Turning quotas on or off 361

Turning quota message logging on or off 363

Creating or resizing quotas 365

Deleting quotas 369

Effects of qtree changes on quotas 370

Understanding quota reports 372
Types of quota reports 373
Overview of the quota report format. 374
Detailed information about the ID and Quota Specifier fields. 376
How to display disk space thresholds 379
How to display soft limits 380
Displaying a quota report 381

Effects of hard quotas on clients 382

Chapter 9 **System Information and Performance** 385
Filer configuration information 386
Storage information 389
Volume information 391
File statistics for volumes 392
Options for the filestats command 396
Checksum upgrade status 399
Filer statistics 401
Environmental status information 404
Fibre Channel information 407
NFS and RPC statistics 408
CIFS statistics 414
DAFS statistics 423
Improving filer performance 424

Appendix A **Filer Problem Messages** 427
Serious error messages 428
Network error codes 429
Disk and FC-AL messages from the Storage Health Monitor 434
Cluster messages 439
LCD messages 440
Syslog messages 442
Console messages 445

Appendix B **Booting the Filer** 447
Booting from the system boot diskette set 448
Booting the F87 filer 451
Booting the FAS200 series appliances 453
Using filers as netboot servers 458

Glossary 461

Index 471

Preface

Introduction

This guide describes how to configure, operate, and manage Network Appliance™ filers that run Data ONTAP™ 6.5 software. It covers all filer models.

Audience

This guide is for system administrators who are familiar with operating systems, such as the UNIX®, Windows 95™, Windows NT®, and Windows 2000® operating systems, that run on the filer's clients. It also assumes that you are familiar with how to configure the filer and how Network File System (NFS), Common Internet File System (CIFS), and Hypertext Transport Protocol (HTTP) 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 NetApp® filer.

Terminology

When the term "filer" is used, the information applies to all NetApp filer models, including the FAS900 and R100 series of storage appliances.

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.

Command conventions

You can enter filer commands either on the system console or from any client computer that can access the filer through a Telnet session.

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.

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 font	Words or characters that require special attention. 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. 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.

Introduction to Network Appliance Filers 1

About this chapter

This chapter describes what a filer is, what it consists of, and how you can use the filer software, Data ONTAP, to configure and manage filers.

Topics in this chapter

This chapter discusses the following topics:

- "About filers" on page 2
- "How you administer a filer" on page 6
- "About filer main unit components" on page 8
- "About disk shelves" on page 10
- "About Data ONTAP" on page 11
- "Filer administration with Data ONTAP" on page 15
- "Getting command information in Data ONTAP" on page 18
- "Documentation for Data ONTAP" on page 19

About filers

What a filer is

A filer is a hardware and software storage system that serves data. 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.

Component	Function	To learn more, read...
Filer main unit, or chassis	The piece of hardware that receives and sends data	"About filer main unit components" on page 8
Disk shelves	Hardware that holds disks and is connected to the filer	"About disk shelves" on page 10
Data ONTAP	Software that is the operating system for the filer	"About Data ONTAP" on page 11

Note

Model FAS250, FAS270 and FAS 270c filers combine the filer head or heads and disk shelf into one unit.

What a filer does

A filer provides the following features.

Feature	Description
Network file service	The filer enables users on client workstations (or hosts) to create, delete, modify, and access files or blocks stored on it. Client workstations are connected to the filer through TCP/IP, direct-attached, or fabric-attached network connections.

Feature	Description
Multiprotocol file and block sharing	Clients can use the following protocols to access data on the filer: ◆ CIFS (Common Internet File System)—used by Windows® clients ◆ FCP (Fibre Channel Protocol)—used for block access in storage area networks ◆ DAFS (Direct Access File System)—used by DAFS clients ◆ HTTP (HyperText Transfer Protocol)—used by the World Wide Web ◆ iSCSI (Internet small computer system interface)—used for block access in storage area networks ◆ NFS (Network File System)—used by UNIX systems ◆ (PC)NFS (Personal Computer NFS)—used by PCs to access NFS Files written through 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. Blocks written through one protocol can also be accessed by clients using the other protocol.

Feature	Description
Data protection	The filer protects disk data in the following ways: Snapshot™ same-volume backup enables you to manually or automatically create, schedule, and maintain multiple backups (or snapshots) of volume and qtree images, using only a minimal amount of additional volume space at no performance cost. SnapMirror® (license required) filer-to-filer snapshot mirroring enables you to mirror snapshots on one filer to a partner filer. Should the original filer be disabled, this ensures quick restoration of data from the point of the last snapshot. SnapVault™ (license required) multiple-sources-to-central-system snapshot backup backs up snapshots from qtrees on multiple volumes and filers to a single Network Appliance SnapVault secondary storage system. This system is specialized for quick backup and restore of its filer sources. The SnapRestore® (license required) software feature performs fast restores of backed-up data on request from snapshots on an entire volume. Tape backup dump and restore commands enable you to back up filer or SnapVault snapshots to tape. Because the snapshot rather than the active file system is backed up to tape, the filer can continue its normal functions while the tape backup is still in operation. Vol copy (license required) fast block copy of data enables you to quickly copy blocks of stored filer data from one volume to another.

Feature	Description
Data protection (continued)	The SyncMirror™ feature performs real-time mirroring of data to two separate filer volumes physically connected to the same filer head. In case of multiple unrecoverable disk errors on one volume, the filer automatically switches access to the mirrored volume. NDMP (Network Data Management Protocol) support enables third-party applications that use NDMP to manage tape backup operations of filer data. The NDMP `copy` command carries out NDMP-compliant backups and restores. Security login restricts access to NDMP operations. The `nvfail` option provides protection against data corruption by nonvolatile RAM (NVRAM) failures.
Autosupport	Mail notifications about filer problems are automatically sent to Network Appliance Technical Support.

How you administer a filer

Methods of administering a filer

You can administer a filer that uses the Data ONTAP operating system by using the following methods:

- Command execution through the filer's command line
- Command execution through Windows NT and Windows 2000
- Configuration file editing
- Command execution through the FilerView® interface
- Command execution through DataFabric™ Manager

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 and Windows 2000

You use Windows NT and Windows 2000 commands to perform filer administrative tasks related to Windows NT and Windows 2000.

You can execute Windows NT commands that affect the filer by 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 for performing 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, and then editing the file from the client.

Command execution through FilerView

Another method of performing many common filer tasks is to use FilerView, the graphical management interface for viewing and managing a filer from a Web browser. FilerView is easy to use, and it includes Help that explains FilerView features and how to use them.

About filer main unit components

Two kinds of components

The main unit of the filer 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 RAM)	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 (not available on all models)	In an emergency and for major upgrades, Data ONTAP uses the diskette drive to boot from diskettes.
CompactFlash™ card (not available on all models)	The filer automatically boots from a Data ONTAP release stored on the CompactFlash card. The CompactFlash card also stores a backup version of Data ONTAP from which to boot the filer in an emergency.

Component	Description
LCD and LEDs	The filer displays status information on the LCD and LEDs.
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 the filer to 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.

Note

Model FAS250, FAS270, and FAS270c filers combine the filer head or heads and the disk shelf into one unit.

About Data ONTAP

Data ONTAP overview

You administer the filer by using Data ONTAP commands either directly or through the FilerView graphical user interface.

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 handles four types of data management:

- Data storage
- Data organization
- Data access
- Data migration

Data storage management

Data ONTAP stores data on disks in disk shelves. When filer volumes are created and these disks are assigned to them, Data ONTAP groups these disks into RAID groups within that volume. The RAID groups are organized to provide parity protection against data loss within each group due to disk or sector failure. RAID and disk-related terms 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 ensures data preservation and file access even if one disk or two disks in a RAID group are damaged.	You can specify the level of RAID protection (RAID 4 or RAID DP) for each volume on the filer.

Data storage concept	Explanation	Storage administrative actions
RAID group	A set of data disks and one or two parity disks, grouped to provide parity protection against data loss within that group due to disk or sector failure. Data ONTAP supports up to 400 RAID groups per filer or cluster.	You can control the size of a RAID group and the level of RAID protection. This enables you to customize backups and disk failure recovery.
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. Depending on the level of RAID protection enabled, each RAID group uses one parity disk (for single disk failure protection under RAID 4) or two parity disks (for double disk failure protection under RAID DP).	
Hot spare disks	Replace failed data disks automatically with reconstructed data.	
Logical unit numbers (LUNs)	Logical units of storage, each of which is identified by a number, used in a storage area network (SAN) environment.	You can create LUNs within volumes and specify the size of a LUN.

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 particular to filers are explained in the following table.

Data organization concept	Explanation	Data organization administrative actions
Volume	One or more independent file systems contained in a filer.	You can specify the following features for a volume: ◆ A security style, which determines whether a volume can contain files that use UNIX security, files that use NT file system (NTFS) security, or both types 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 to 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:

- Manage network connections
- Manage protocols
- Configure pairs of filers into clusters for takeover
- Dump data to tape and restore it to the filer
- Copy volumes
- Mirror volumes (synchronously and asynchronously)
- Create vFiler™ units

Data migration management

Data ONTAP manages data migration by means of

- Snapshots
- Asynchronous mirroring
- Synchronous mirroring
- Backup to tape
- Volume copy

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, the Setup Wizard, the `setup` program, or both if necessary, and then editing configuration files if necessary.

Monitoring and maintaining client access

You monitor client access to data on the filer by using 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 CIFS access so that you can identify potential security problems

You maintain client access by using 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
	Review DAFS utilization, if licensed for DAFS	`dafs stats` 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 Network Information Service (NIS) server	`ypwhich` command

Getting command information in Data ONTAP

Where information about Data ONTAP commands is

Basic information about each Data ONTAP command is stored in manual pages in the etc/man directory in the filer. You can access this information from the command line by using the `man` command.

How to get information about Data ONTAP commands

To get information about a Data ONTAP command, complete the following step.

Step	Action
1	Enter the following command: **`man [-s section] command_name`** *section* is one of the four sections of Data ONTAP man page information. Section 1 is Commands, Section 4 is Special Files, Section 5 is File Formats and Conventions, and Section 8 is System Management and Services. **Result:** The manual page information is displayed.

Documentation for Data ONTAP

Description and location of Data ONTAP documents

Documentation for Data ONTAP includes a release description, setup and upgrade guides, a number of system administration guides for various aspects of administering NetApp storage appliances, and a comprehensive command reference.

The documents are available on the NetApp on the Web (NOW) site. Click Documentation on the home page, then click Data ONTAP on the Access Documentation page to see links to the latest version of the documentation.

The following table describes the information in each document.

Document name	Information included
Release Description	Describes the latest release of Data ONTAP, including new features, requirements, and factors to consider before you upgrade to it.
Software Setup Guide	Describes how to set up Data ONTAP software on NetApp storage appliances.
Upgrade Guide	Describes how to upgrade to the latest version of Data ONTAP software and presents information you need to consider before upgrading your software or reverting it to an earlier version.
Storage Management Guide	Describes how to configure, operate, and manage Network Appliance storage appliances with Data ONTAP.
File Access Management Guide	Describes filer operations and how to manage file access for the NFS, CIFS, HTTP, FTP, WebDAV, and DAFS protocols.

Document name	Information included
Block Access Management Guide for FCP	Describes how to configure and manage the FCP protocol, and how to create and manage LUNs and initiator groups, with the optional FCP license.
Block Access Management Guide for iSCSI	Describes how to configure and manage the iSCSI protocol, and how to create and manage LUNs and initiator groups, with the optional iSCSI license.
Network Management Guide	Describes how to configure and manage networks associated with NetApp storage appliances.
Data Protection: Online Backup and Recovery	Describes how to back up and recover data using Data ONTAP online backup and recovery features.
Data Protection: Tape Backup and Recovery	Describes how to back up and recover data using Data ONTAP tape backup and recovery features.
MultiStore™ Management Guide	Describes how to administer virtual filers (vFilers) with the optional MultiStore software available by license with Data ONTAP.
Commands: Manual Page Reference	Contains the manual (man) pages for Data ONTAP commands.

Documentation for FilerView and DataFabric Manager

FilerView, the graphical user interface you can use instead of the command-line interface to manage your filer, has Help built in. DataFabric Manager, an optional software program for managing multiple NetApp appliances, comes with an administrator's guide and a command reference as well as built-in Help.

Introduction to Filer Administration 2

About this chapter

This chapter introduces you to filer administration.

Topics in this chapter

This chapter discusses the following topics:

- "Managing access from administration hosts" on page 22
- "Managing local user accounts" on page 25
- "Accessing a filer to administer it" on page 32
- "Editing command-line commands" on page 43
- "Using Data ONTAP commands at different privilege levels" on page 45
- "Receiving automatic technical support" on page 47
- "Understanding the filer's default volume" on page 58
- "Configuring filer startup" on page 64
- "Understanding the options and vol options commands" on page 67
- "Performing common tasks" on page 70
- "Filer configuration backup and cloning" on page 102

Managing access from administration hosts

Why administration hosts are used

Administration hosts are used for two reasons:

- Some administration tasks require that the administrator edit configuration files, but the Data ONTAP software does not include a file editor. The only way you can edit files is by accessing them from an administration host.
- It is often more convenient to administer a filer remotely than to go to the filer's physical location.

What an administration host is

The administration host is a remote workstation from which you can administer the filer. The remote workstation can be an NFS client or a CIFS client. When you designate an administration host, the filer's root file system (/vol/vol0 by default) is accessible to the host through the NFS protocol and as a CIFS share.

When to add administration hosts

Initially, you make a client an administration host when you set up the filer. After booting the filer, you can add as many clients as you want to be administration hosts.

Administration host privileges

The filer grants root permissions to the administration host after the `setup` procedure is 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 connection.
A CIFS client	Edit configuration files from any CIFS client as long as you connect to the filer as root or "Administrator."

Requirements for using an NFS client

If you plan to use an NFS client to manage the filer, the NFS 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 by using the NFS protocol

Requirements for using a CIFS client

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

Where and how administration hosts are specified

Administration hosts are specified in the /etc/hosts.equiv file in the following format:

`hostname` [*`username`*]

`hostname [ “user name” ]`

Example:

```
client1 peter
client1 “joe smith”
```

Rules for /etc/hosts.equiv file entries

The following rules apply to entries in the /etc/hosts.equiv file:

- If multiple users on the same host require access to the filer through a Remote Shell, you must specify each user’s entry as in the following example:

  ```
  client1 anna
  client1 carl
  client1 peter
  client2 leann
  client2 sue
  client2 tony
  ```

- If *hostname* specifies an NFS client, the user name is optional. If you do not specify a user name, you must be the root user on that NFS client to execute a filer command through the `rsh` command.
- If *hostname* specifies a CIFS client, you must enter the user name for that CIFS client.

The following is an example of the contents of an /etc/hosts.equiv file:

```
nfsclient1
client2 lena
client2 root
client3 fred
client3 root
```

Adding administration hosts

To designate additional NFS clients or CIFS clients as administration hosts, complete the following steps.

Step	Action
1	Open the /etc/hosts.equiv configuration file with an editor.
2	Add the host names and user names of the clients that you want designated as administration hosts.
3	Save the /etc/hosts.equiv file.

Removing administration hosts

To remove an NFS client or CIFS client from the administration hosts list, complete the following steps.

Step	Action
1	Open the /etc/hosts.equiv configuration file with an editor.
2	Locate and delete the entries for the host names and user names you want to remove.
3	Save the /etc/hosts.equiv file.

Managing local user accounts

Reasons for creating local user accounts

There are two primary reasons for creating local user accounts on the filer:

- By default, Data ONTAP supports a system account, which is root. You can, however, create optional local user accounts so that other administrators can manage the filer. The accounts you create for this purpose are also called administrative user accounts.

  An administrator using an administrative user account can perform the same tasks that the administrator can perform by using the system account. Administrative user accounts add security by allowing accounts to be established for individuals on one or more filers, unlike system accounts, which are created on all filers. This enables the system account password to remain secret.
- If a filer running CIFS is a member of a domain or a Windows NT workgroup, you can create local user accounts for authenticating users. For more information about authenticating users using local user accounts, see the section on local user accounts in the CIFS chapter of the *Block Access Management Guide*.

Tasks for managing local user accounts

The following are tasks you can perform to manage passwords and local user accounts:

- "Changing the system password" on page 26
- "Creating a local user account" on page 27
- "Changing a local user account password" on page 29
- "Changing a local user account group or comment" on page 30
- "Deleting a local user account" on page 30

Changing the system password

To change the system password, which is also the password for the root user account, complete the following step.

Step	Action	
1	**If you are using a...**	**Then...**

	Telnet session or the console to administer the filer	1. Enter the following command:
		`passwd`
		2. Enter the system account name:
		`root`
		3. Enter the existing system password.
		4. Enter a new password, and then enter it a second time to confirm it.

	Remote Shell connection to administer the filer	Enter the following command: `rsh` *`filer_name`* `passwd` *`old_password new_password`*

Listing local user accounts

To list one or more local users, complete the following step.

Step	Action	
1	**If you are using a...**	**Then...**
	Telnet session or the console to administer the filer	Enter the following command: **`useradmin userlist [`*`username list`*`]`** *username list* is a space-separated list of the login names you want to display.
	Remote Shell connection to administer the filer	Enter the following command: **`rsh useradmin userlist [`*`username list`*`]`** *username list* is a space-separated list of the login names you want to display.
	Note If you do not specify any user names, all local users, except root, are listed.	

Creating a local user account

You must be logged in as a local user to create a user account.

User account passwords and user names have the following restrictions:

- The password must contain at least six characters, including at least two alphabetical characters and one nonalphabetical character.
- Passwords must contain no more than 16 characters.
- The login name is case-sensitive.
- The login name cannot contain any of the following characters:

  * \ : , = / > [< | + ?] ; space

To create a local user account, complete the following step.

Step	Action	
1	**If you are using a...**	**Then...**
	Telnet session or the console to administer the filer	1. Enter the following command: **useradmin useradd [-c "*comments*"] *login_name*** *comments* is an optional opportunity to enter comments about the account, such as a user's full name. *login_name* is the new local user login name. 2. When Data ONTAP prompts you for a password, enter the password associated with the new local user login name. 3. When Data ONTAP prompts you to reenter the password, enter the password again.
	Remote Shell connection to administer the filer	Enter the following command: **rsh *filer_name* useradmin useradd [-c "*comments*"] *login_name password*** *comments* is an optional opportunity to enter comments about the account, such as a user's full name. *login_name* is the new local user login name. *password* is the password associated with the new user name.

Changing a local user account password

To change a local user account password, complete the following step.

Step	Action	
1	**If you are using a...**	**Then...**
	Telnet session or the console to administer the filer	**1.** Enter the following command: `passwd` **2.** When Data ONTAP prompts you, enter the name of the local user whose password you want to change. **3.** When Data ONTAP prompts you, enter the new password. **4.** Enter the new password again.
	Remote Shell connection to administer the filer	Enter the following command: `rsh` ***`filer_name`*** `passwd` ***`new_password username`***

Changing a local user account group or comment

To change a local user account group membership or comment, complete the following step.

Step	Action	
1	**If you are using a...**	**Then...**
	Telnet session or the console to administer the filer	Enter the following command: **`useradmin usermodify [-c "`*`comments`*`"] login_name`** *comments* is the new comment that you want to specify for the affected user. *login_name* is the local user whose login group membership or comment is being modified.
	Remote Shell connection to administer the filer	Enter the following command: **`rsh `*`filer_name`*` useradmin usermodify [-c "`*`comments`*`"] login_name`**

Deleting a local user account

To delete a local user account, complete the following step.

Caution

Data ONTAP deletes a local user account immediately, without asking you to confirm the deletion.

Step	Action	
1	**If you are using a...**	**Then...**
	Telnet session or the console to administer the filer	Enter the following command: **`useradmin userdel`** ***`login_name`*** *login_name* is the login name of the local user account you want to delete.
	Remote Shell connection to administer the filer	Enter the following command: **`rsh`** ***`filer_name`*** **`useradmin userdel`** ***`login_name`*** *login_name* is the login name of the local user account you want to delete.

Accessing a filer to administer it

Ways to access a filer

There are four ways you can directly access a filer to administer it:

- From a console that is attached by a cable to the filer's serial port
- From an NFS client or CIFS client by using a Telnet session
- From an NFS client or CIFS client by using a Remote Shell connection
- From an NFS client or CIFS client by using a Web browser and the FilerView interface

Detailed information

The following sections discuss the ways to access a filer:

- "Accessing a filer from the console or a Telnet session" on page 33
- "Accessing a filer with a Remote Shell connection" on page 36
- "Accessing a filer with FilerView" on page 39

Accessing a filer from the console or a Telnet session

Background information

You need to understand the following information before you access the filer through the console or a Telnet session:

- The console and Telnet share a single session. Commands typed at either location are echoed to the other location. Messages are displayed at both locations.
- You can abort commands entered at the console or through a Telnet session by pressing Ctrl-C.

Note

Because the console and a Telnet session share a single session, pressing Ctrl-C aborts the current command regardless of where the command was entered. You can prevent commands from being aborted at the console or through a Telnet session by using the `rsh` command to initiate commands from an administration host.

- Only one Telnet session can be active at a time.

Accessing a filer from the console

To access a filer from a console that is attached by a cable to the filer's serial port, complete the following steps.

Step	Action
1	Press Enter on the console. **Result:** The filer responds with the login or password prompt.

Step	Action
2	If the filer displays the login prompt, do one of the following: ◆ To access the filer with the system account, enter the following account name: `root` ◆ To access the filer with an alternative administrative user account, enter the following: ***`username`*** *username* is the administrative user account. **Result:** The filer responds with the password prompt.
3	Enter the password for the root or administrative user account. **Note** If no password is defined for the account, press Enter.

Accessing a filer with Telnet

To access a filer from an NFS client or a CIFS client through a Telnet session, complete the following steps.

Step	Action
1	Open a Telnet session on a client. **Note** The client from which you run Telnet must be listed in the filer /etc/hosts file.
2	Connect to the filer using its host name.

Step	Action
3	If the filer displays the login prompt, do one of the following: ◆ To access the filer with the system account, enter the following account name: `root` ◆ To access the filer with an alternative administrative user account, enter the following: ***`username`*** *username* is the administrative user account. **Result:** The filer responds with the password prompt.
4	Enter the password for the root or administrative user account. **Note** If no password is defined for the account, press Enter.

Terminating a Telnet session

To terminate a Telnet session with a filer, complete the following steps.

Step	Action
1	Press Ctrl-] or enter `logout telnet` to log out of the filer, or press Ctrl-D to close the Telnet session. **Note** You can use the `logout telnet` command from either the filer console or a Remote Shell connection.
2	Close the Telnet client.

Accessing a filer with a Remote Shell connection

Remote Shell with trusted remote hosts supported

The filer supports a Remote Shell with trusted remote hosts—those remote hosts listed in /etc/hosts.equiv on the root volume.

Use of rsh command with user names and passwords

Under some circumstances, you might need to supply a user name and a password when using the `rsh` protocol to run a command on the filer. The following table describes the circumstances under which a user name and password are required.

If you are logged in as...	Then...
Root on a UNIX host that is listed in the filer's /etc/hosts.equiv file	You do not need to supply a user name or a password.
A user other than root on a UNIX host that is listed in the filer's /etc/hosts.equiv file	You need to supply a user name and a password. The user name can be root or the name of an administrative user that has been defined on the filer.

Note

To issue commands from a Remote Shell on a PC, you must always supply a user name for the PC in the filer's /etc/hosts.equiv file. For more information, see the na_hosts.equiv(5) man page.

Format for rsh commands with user name and password

The format for `rsh` commands that include a user name and a password is as follows:

`rsh` *`username`*`:`*`password command`*

Example: The following example illustrates an `rsh` command entered with a user name and a password:

```
rsh -l root:figby cifs shares
```

Caution

Passing a password from an NFS client in this manner is a security risk.

Commands not accepted from rsh

You cannot execute the following commands by using the `rsh` command:

```
?
arp
cdacmd
license
netdiag
orouted
ping
routed
savecore
setup
shelfchk
traceroute
```

Accessing a filer from an NFS client by using an rsh command

To access a filer from an NFS client by using the `rsh` command, complete the following step.

Step	Action
1	Enter the `rsh` command, replacing *command* and *arguments* with the desired command and arguments: **`rsh filername command arguments`** **Example:** `rsh filer12 sysstat 1`

Accessing a filer from a CIFS client by using a Remote Shell application

To access a filer from a CIFS client by using a Remote Shell application, complete the following steps.

Step	Action
1	Run the Remote Shell application on the CIFS client.
2	Enter the client in the /etc/hosts.equiv file.

Step	Action
3	Enter the `rsh` command, replacing *command* and *arguments* with the desired command and arguments: **`rsh filername command arguments`** **Example:** `rsh filer12 sysstat 1`

Accessing a filer with FilerView

What FilerView is

FilerView is a Web-based graphical management interface that enables you to manage most filer functions from a Web browser rather than from the console, a Telnet session, or the `rsh` command.

FilerView requires a Web browser

FilerView requires a Netscape Navigator® 4.51 or later or Microsoft® Internet Explorer™ 4.0 or later browser.

Note

The browser must have Java and JavaScript enabled. Other browsers that support Java and JavaScript might be compatible with FilerView.

Client platforms that support FilerView

The following platforms support FilerView:

- Windows 95
- Windows 98
- Windows 2000
- Windows NT
- Windows XP
- UNIX
- Solaris
- Linux
- HPUX

Accessing a filer from a client by using FilerView

To access a filer from a client by using FilerView, complete the following steps.

Steps	Action
1	Start your Web browser.

Steps	Action
2	Enter the following URL, replacing *filername* with the name of your filer: `http://filername/na_admin` *filername* is either the fully qualified name or the short name of the filer or the IP address of the filer.
3	Click the FilerView icon. If FilerView prompts for them, supply an administrative user name and password. **Result:** FilerView starts up in a new browser window. **Note** If you are running SecureAdmin™ 2.1.1 software or later, click the Secure FilerView icon to start an encrypted browser session.
4	Select a management function.

Using the interface

The FilerView interface consists of three main frames: a left frame, a right frame, and a title frame.

Left frame: The left frame contains an expandable list of categories:

- The Real Time Status category contains choices that launch separate tools that monitor filer performance.
- The other categories represent management functions.

To view the real-time displays or the functions in a category, complete the following step.

Step	Action
1	Click Real Time Status or a category name.

To select a real-time display or a function, complete the following step.

Step	Action
1	Click the display or the function name.

Right frame: As you select functions from the left frame, the right frame changes to display forms that provide information about the filer configuration and enables you to change the filer configuration.

To change the filer configuration, complete the following steps.

Step	Action
1	Change information that is displayed in the right frame.
2	Save the information.

Title frame: The title frame contains the name of the function you select from the left frame, followed by the path to the function. For example, if you select Report in the Volumes category, the title frame shows the path as Volumes > Report.

Help buttons: Help buttons are situated next to the categories in the left frame and in the title frame. Help provides a description of the function, descriptions of the fields that the function uses, and procedures for tasks you can perform with the function.

To view Help, complete the following steps.

Step	Action
1	Click the Help button next to the category. **Result:** A two-frame Help screen appears. The left frame displays an expandable table of contents, with additional tabs you can select at the top labeled Index and Search. As you select topics from the left frame, the right frame displays a page with tabs you can select at the top labeled Concepts, Procedures, and More Information, which relate to the topic you selected.
2	To view tasks, click the Procedures tab in the left frame.

Step	Action
3	To view additional information about a topic, when available, click the More Information tab in the left frame.
4	To view descriptions of fields, click the Help icon next to the appropriate field in the left FilerView frame.

Editing command-line commands

Features for editing commands

Data ONTAP provides the following two features to help you when you enter commands on the command line:

- History
- Command-line editor

Using the history feature

The history feature enables you to scroll through recently entered commands.

To use the history feature, complete the following step.

Step	Action	
1	**If you are using...**	**Then...**
	The cursor keys	Press the Up arrow key to scroll back through commands and the Down arrow key to scroll forward through commands.
	Key combinations	Press Ctrl-P to scroll back through commands and Ctrl-N to scroll forward through commands.

Using the command-line editor

The command-line editor enables you to position the cursor anywhere in a partially typed command and insert characters at the cursor position. You can use various key combinations to move the cursor within the same line and edit the command.

To use the command-line editor on the command line, complete the following step.

Step	Action	
1	**If you want to...**	**Then...**
	Move the cursor right one position	Press Ctrl-F or the Right arrow key.
	Move the cursor left one position	Press Ctrl-B or the Left arrow key.
	Move the cursor to the end of the line	Press Ctrl-E.
	Move the cursor to the beginning of the line	Press Ctrl-A.
	Delete all characters from the cursor to the end of the line	Press Ctrl-K.
	Delete the character to the left of the cursor and move the cursor left one position	Press Ctrl-H.
	Delete the line	Press Ctrl-U.
	Delete a word	Press Ctrl-W.
	Reprint the line	Press Ctrl-R.

Using Data ONTAP commands at different privilege levels

About privilege levels

Data ONTAP provides two sets of commands, depending on the privilege level, which can be administrative or advanced. You can set the privilege level by using the `priv` command.

At the administrative level, you have access to commands that are sufficient for managing your filer. At the advanced level, in addition to all commands available at the administrative level, you have access to commands for troubleshooting. Commands accessible only at the advanced level should be used under the guidance of Network Appliance Technical Support.

Caution

Advanced commands can be dangerous. Using some advanced commands without consulting Network Appliance Technical Support might result in data loss.

How different privilege settings apply to different sessions

The console session and the Telnet session share the same privilege setting. For example, if you set the privilege level to advanced at the console, the advanced commands also become available to the administrator who is connected to the filer using Telnet.

You can, however, set a different privilege level for each invocation of `rsh`. Suppose your privilege level at the console is administrative and, through `rsh`, another administrator sets the privilege level to advanced. Your privilege level at the console remains unchanged.

Initial privilege level

The initial privilege level for the console and each `rsh` session is administrative.

Data ONTAP resets the privilege level to administrative for each `rsh` session. If you have a script invoking multiple `rsh` connections and you want to execute advanced commands in each connection, you must set the privilege level accordingly for each `rsh` session. If you set the privilege level for the first `rsh` session only, Data ONTAP fails to execute the advanced commands through the subsequent `rsh` sessions, because the privilege level for each subsequent session is administrative.

Setting the privilege level

To set the privilege level, complete the following step.

Step	Action
1	Enter the following command: **`priv set { admin | advanced }`** `admin` sets the privilege level to administrative. `advanced` sets the privilege level to advanced.

Receiving automatic technical support

How to get automatic technical support

You can receive automatic technical support for your filer by purchasing a support subscription and configuring the Autosupport feature.

Detailed information

The following sections discuss how the Autosupport feature works and how you can customize its configuration:

- "Understanding the Autosupport feature" on page 48
- "Events that trigger Autosupport messages" on page 49
- "Contents of Autosupport messages" on page 50
- "Configuring Autosupport" on page 53

Understanding the Autosupport feature

What Autosupport does

The `autosupport` daemon monitors the filer's operations and sends automatic messages to Network Appliance Technical Support to alert it to potential filer problems. If necessary, Network Appliance Technical Support contacts you at the e-mail address that you specify to help resolve a potential system problem. The `autosupport` daemon is enabled by default on the filer. The filer can send Autosupport messages by SMTP, HTTP, or HTTPS (the secure socket layer [SSL] Internet transport protocol). HTTPS is the default.

Autosupport daemon requires an external mail host if you use SMTP

The filer does not function as a mail host—it requires an external mail host at your site to send mail. The mail host is a host that runs a mail server that listens on the SMTP port (25).

Examples of mail servers include the following:

- A UNIX host running an SMTP server such as the sendmail program
- A Windows NT server running the Microsoft Exchange server

The filer uses the mail host's mail server to send periodic e-mail messages automatically to Network Appliance Technical Support about the filer's status. You can configure `autosupport` to use one or more mail hosts.

Autosupport messages using HTTP/HTTPS

To use HTTP or HTTPS to send Autosupport messages, you may need to configure an HTTP or HTTPS proxy. For more information, see "Configuring Autosupport" on page 53.

HTTP uses port 80; HTTPS uses port 443. If the network connection does not allow HTTPS or HTTP, you must configure Autosupport for SMTP.

Events that trigger Autosupport messages

Event and log-level descriptions

The filer sends Autosupport messages to Network Appliance Technical Support about your filer after any of several events. The Autosupport message includes a log level that indicates the priority assignment from Network Appliance Technical Support:

- CRITICAL—Priority 1
- ERROR—Priority 2
- WARNING—Priority 3
- NOTICE—Informational, no response expected
- INFO—Informational, no response expected
- DEBUG—Informational, no response expected

Where to get Autosupport message descriptions

If you are using Autosupport locally, you will see the previously mentioned log levels in the subject line of the Autosupport e-mail that you receive.

To read descriptions of the Autosupport messages that you might receive, complete the following steps.

Step	Action
1	Go to the Network Appliance NOW site, and click to Service and Support > Access Product Documentation > Autosupport Features and Functions > Autosupport Message Matrices.
2	On the Message Matrices page under Select a Release, select version 6.5 or higher for Data ONTAP and click View Matrix.
3	All Autosupport message descriptions are listed alphabetically by subject line.

Contents of Autosupport messages

Contents of Autosupport event messages and weekly reports

Autosupport messages contain various kinds of information, such as dates, version numbers, and serial numbers. The contents shown here are represented as follows:

- Items marked with an asterisk (*) are suppressed in the `Autosupport.content Minimal` format.
- Items marked with two asterisks (**) are filtered in the `Autosupport.content Minimal` format.
- Commands marked with a dagger (†) are advanced commands. It you want more information on the output of these commands, call Network Appliance Customer Service.
- Commands marked with a double dagger (‡) print as shown if the verbose option is *not* set for them.

  Commands that show both the verbose and regular options print as verbose if the verbose option is set. For example, the output of the command `dafs status -n` is printed unless `autosuppport.dafs.verbose` is set to on, in which case the output of `dafs status -nv` is printed to display the status of the NICs in verbose form.

Each Autosupport message contains the following types of information.

- Date and timestamp of the message
- NetApp software version
- Serial number of the filer
- Encrypted software licenses*
- Host name of the filer*
- SNMP contact name and location (if specified)*
- Console encoding type

- Output of the following commands (some are applicable only to the licensed protocols):

availtime	fcstat fcal_stats	storage show adapter -a
cf monitor all	fcstat link_stats	storage show disk -a
cf rsrctbl†	httpstat	storage show fabric
cf timers†	ic stats error -v	storage show hub -a
cifs domaininfo *	ic stats performance	storage show port
cifs sessions*‡	ifconfig -a*	storage show switch
cifs sessions -t‡	ifstat -a**	sysconfig -a**
cifs shares*‡	igroup show	sysconfig -c
cifs shares -t‡	iscsi config	sysconfig -d
cifs stat	iscsi show adapter	sysconfig -D
dafs status -n‡	iscsi show initiator	sysconfig -r
dafs status -nv‡	iscsi stats	vfiler status -a
dafs status -sr‡	iscsi status -v	viadmin list -v
dafs status -srv‡	license	vif status*
dafs status -v	lun show‡	vlan stat*
df	lun show -v‡	vol scrub status -v
df -i	lun stats	vol status
disk shm_stats ata	nbstat -c *	vol status -c
disk show -n	netstat -s	vol status -l*
ems event status**	nfsstat -c	vol status -v
ems log status*	nis info*	vscan scanners *
environment status all	options**	vscan options
fcp config	perf report -t	wafl swarmstats
fcp show adapter‡	qtree status -i -v	
fcp show adapter -v‡	snap sched	
fcp show initiator‡	snet status -v†	
fcp show initiator -v‡		
fcp status -v		
fcstat device_map		

- Checksum status
- ECC memory scrubber statistics
- The following information if clustering is licensed:**

  System ID of the cluster partner

  Host name of the cluster partner

  Cluster node status, including cluster monitor and cluster interconnect statistics
- Contents of the following /etc directory files

  /etc/messages (to last WEEKLY_LOG event)**

  /etc/log/ems files (to last WEEKLY_LOG event) (optional)**

/etc/serialnum file

/etc/rc file*

/etc/nsswitch.cong file*

/etc/exports file*

/etc/snapmirror.allow file (if the SnapMirror license is enabled)*

/etc/snapmirror.conf file (if the SnapMirror license is enabled)*

/etc/syslog.conf file*

/etc/hosts file*

- Registry information
- Usage information*
- Service statistics
- Boot time statistics*
- NVLOG statistics*
- WAFL check log
- Modified configurations
- X-header information

Using the `options` command, you can specify the value of `autosupport.content` as `complete` or `minimal` to control the detail level of event messages and weekly reports. Complete Autosupport messages are required for normal technical support. Minimal Autosupport messages omit sections and values that might be considered sensitive information and reduce the amount of information sent. Choosing `minimal` greatly affects the level of support you can receive. For more information about setting Autosupport options, see “Autosupport options” on page 53. You can also find information about Autosupport features and functions on the NOW site at http://now.netapp.com/autosupport.

Configuring Autosupport

Autosupport options

To configure Autosupport, you specify Autosupport options. The Autosupport options are shown in the following table.

Autosupport option	Description
`autosupport.cifs.verbose [ on | off ]`	Enables and disables inclusion of CIFS sessions and shares information in Autosupport messages. The default is Off.
`autosupport.content [ complete | minimal ]`	Indicates the type of content that Autosupport messages should contain. The default is Complete. **Note** If this setting is changed from complete to minimal, any complete content Autosupport message not yet sent is cleared from the outgoing message spool and a message to that effect appears on the console.
`autosupport.dafs.verbose [ on | off ]`	Enables and disables inclusion of DAFS information in Autosupport messages. The default is Off.

Autosupport option	Description
`autosupport.doit [ message ]`	Tells the `autosupport` daemon to send an Autosupport notification immediately. The message can be a single word or a string enclosed in single quotation marks. The message is included in the subject line of the Autosupport notification and should be used to explain the reason for the notification. You can verify that Autosupport is working by using the "Call Home Check" function, which sends an `autosupport.doit` message with a subject line containing any variation of the word TEST or TESTING. When such a message is sent to autosupport@netapp.com, the mail handler sends an automated response to those addresses, indicating that the test Autosupport message was received successfully.
`autosupport.enable [ on | off ]`	Enables and disables Autosupport notification. The default is On.
`autosupport.from sender`	Defines the user to be designated as the sender of the notification.
`autosupport.mailhost host1[, ..., host5]`	Defines the list of up to five mail host names. The host names should be entered as a comma-separated list with no spaces in between. The default an empty string.
`autosupport.minimal.subject.id [hostname | systemid]`	Defines how the system is identified in the Autosupport message title if `autosupport.content` is Minimal. The default is System ID.

Autosupport option	Description
`autosupport.noteto` *`address1`*`[, ..., `*`address5`*`]`	Defines the list of recipients for the Autosupport short note e-mail. Up to five mail addresses are allowed. Enter the addresses as a comma-separated list with no spaces in between. The default is a null list to disable short note e-mails.
`autosupport.retry.count` *`#retries`*	Defines the number of times the filer will try to resend the Autosupport notification before giving up. Retries can be between 5 and 4,294,967,295. The default is 15.
`autosupport.retry.interval` *`interval`*	Defines the time to delay before trying to resend the Autosupport notification. The values can end with s, m, or h to indicate seconds, minutes, or hours, respectively. If no units are specified, the value is assumed to be in seconds. Values can range from 30 seconds to 1 day. The default is 4 minutes.
`autosupport.support.enable [ on | off ]`	Enables and disables the Autosupport notification to Network Appliance. The default is On.
`autosupport.support.proxy`	Allows you to set an HTTP proxy if necessary. This is useful only if `autosupport.support.transport` is http or https.
`autosupport.support.to`	Indicates where Autosupport notifications to Network Appliance are sent if `autosupport.support.transport` is smtp. This option is read-only and is shown for informational purposes only.

Autosupport option	Description
`autosupport.support.transport [ http | https | smtp ]`	Defines the type of delivery for Autosupport notifications to Network Appliance. The default is https.
`autosupport.support.url`	Indicates where Autosupport notifications to Network Appliance are sent if `autosupport.support.transport` is http or https. This option is read-only and is shown for informational purposes only.
`autosupport.throttle [ on | off ]`	Drops additional messages when too many Autosupport messages of the same type are sent in too short a time. The default is On.
`autosupport.to` *`address1`*`[, ..., address5]`	Defines the list of recipients for the Autosupport e-mail notification. Up to five mail addresses are allowed, or the list can be left empty. Enter the addresses as a comma-separated list with no spaces in between. The default is no list.

Configuring the Autosupport feature

To configure Autosupport, complete the following step.

Note

All Autosupport options except the `autosupport.doit` option are persistent between reboots.

Step	Action
1	Issue any of the Autosupport option commands shown in the preceding table to the filer: **`options autosupport.option arguments`** *autosupport.option* is one of the options shown in the preceding table. *arguments* refers to any required or optional arguments shown in the preceding table.

Testing Autosupport

To test Autosupport, complete the following step.

Step	Action
1	Enter the following command: **`options autosupport.doit message`** *message* is the subject line for the test Autosupport e-mail.

If you use the keyword TEST in the message, you will receive a return message from Network Appliance, indicating that the Autosupport process is working correctly. For more details, see the information on the `autosupport.doit` command in “Autosupport options” on page 53.

Understanding the filer's default volume

What the default volume is

The filer contains a default volume that was created when the filer was set up. Unless the installer selected a unique volume name during setup, the default volume is named /vol/vol0.

Note

The /vol portion of the name indicates to the filer that the remainder of the name specifies a volume. The /vol prefix is required whenever a filer volume is referred to from a client or when a volume name is used in a command given to the filer.

Detailed information

The following sections discuss the filer's default volume:

- "Default directories in the default volume" on page 59
- "Contents of the /etc directory" on page 60
- "Editing configuration files" on page 62

Default directories in the default volume

What the default directories are

The default volume contains a root directory, which, in turn, contains the following directories that were created when the filer was set up:

- The /etc directory, which contains configuration files that the filer needs in order to operate
- The /home directory, which is a default location you can use to store data

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 /etc directory	NFS	◆ Full permissions for the root user on the administration host (-rwx) ◆ No permissions for any other user or host
	CIFS	◆ Read and write permissions to all files for the administrative user when logged in to the filer by use of the root password (Full Control) ◆ No permissions for other users
The /home directory	NFS	Permissions associated with individual users and with groups through a UNIX security database
	CIFS	Permissions for the HOME$ share are Full Control for Everyone

Contents of the /etc directory

The /etc directory containing configuration files

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.

Descriptions of configuration files

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

File name	Contents
Files used by CIFS, (PC)NFS, and NFS	
hosts	Known hosts and their IP addresses
hosts.equiv	Trusted hosts and users for `rsh`
nsswitch.conf	The order in which the filer contacts name services
quotas	Stored user and group quota specifications
rc	Script of commands to be executed when the filer is initializing
usermap.cfg	Map of Windows user names to UNIX user names
Files used only by CIFS	
.cifs.dat	Domain information (only if the filer is a member of a domain)
cifsconfig.cfg	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.

File name	Contents
group	CIFS group names, GIDs (group identification numbers), and members' names Not used if you use NIS to authenticate groups
Files used by CIFS, (PC)NFS, and FTP	
passwd	Users' names, UIDs (user identification numbers), and primary GIDs Not used if you use NIS to authenticate users
Files used only by NFS	
exports	NFS export points
netgroup	Network groups

Character coding of /etc/usermap.cfg and /etc/quotas

The /etc/usermap.cfg and /etc/quotas files support two types of encoding: Unicode and root volume UNIX encoding. As a result, you can edit the files from either a PC or a UNIX workstation. Data ONTAP can detect whether a file was edited and saved by a Unicode-capable editor, such as Notepad. If so, Data ONTAP considers all entries in the file to be in Unicode. Otherwise, Data ONTAP considers the entries to be in the root volume UNIX encoding. SGML entities are allowed in either encoding.

Editing configuration files

Configuration files must be edited from clients

Data ONTAP does not include an editor. You cannot edit files by using the system console or by establishing a Telnet session to the filer. You must use an editor from an NFS client or a CIFS client to edit filer configuration files.

Enabling an NFS client to edit configuration files

For an NFS client to edit configuration files, the client must be authorized to access the root file system.

If the NFS client was specified as the administration host during setup or added as an administration host after setup was completed, it is already authorized to access the root file system.

If the NFS client is not specified as an administration host, complete the following steps to authorize access to the root file system.

Step	Action
1	Mount the filer root volume on 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 on the client.

Editing configuration files from an NFS client

To edit configuration files from an NFS client, complete the following step.

Step	Action
1	From the client, use a text editor to edit the files in the /etc directory.

Editing configuration files from a CIFS client

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. 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.

To edit configuration files from a CIFS client, complete the following steps.

Step	Action
1	Connect from a CIFS client to the filer as Administrator.
2	Display the contents of the filer's C$ share, and select a file to edit.

Configuration files that should not be edited

When you configure Data ONTAP, it creates some files that you should not edit. The following configuration files should not be edited.

- cifsconfig.cfg
- cifssec.cfg
- lclgroups.cfg
- filesid.cfg
- sysconfigtab
- registry
- registry.bck
- registry.default
- registry.lastgood

Configuring filer startup

Startup commands are stored in /etc/rc file

The /etc/rc file contains commands that the filer executes at boot time to configure the filer.

What startup commands do

Startup commands are placed into the /etc/rc file automatically after you run the `setup` command or the Setup Wizard.

Commands in the /etc/rc file configure the filer to do the following actions:

- Communicate on your network
- Use the NIS and DNS services
- Set the filer's time zone
- Save the core dump that might exist if the filer crashed before it was booted
- Export directories to NFS, if NFS is licensed and configured
- Enable NFS, if NFS is licensed

Default /etc/rc file contents

To understand the commands used in the /etc/rc file on the root volume, examine the following sample /etc/rc file, which contains default startup commands:

```
#Auto-generated /etc/rc Tue May 30 14:51:36 PST 2000
hostname toaster
ifconfig e0 'hostname'-0
ifconfig e1 'hostname'-1
ifconfig f0 'hostname'-f0
ifconfig a5 'hostname'-a5
route add default MyRouterBox
routed on
timezone US/Pacific
savecore
exportfs -a
nfs on
```

Explanation of default /etc/rc contents

This section explains the sample /etc/rc file:

Description: `hostname toaster`

Sets the filer host name to “toaster.”

Description:

```
ifconfig e0 ‘hostname‘-0
ifconfig e1 ‘hostname‘-1
ifconfig f0 ‘hostname‘-f0
ifconfig a5 ‘hostname‘-a5
```

Sets the IP address for the filer network 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 filer root volume. If you prefer to have the actual IP addresses in the /etc/rc file, you can enter IP addresses directly in /etc/rc on the root volume. For more information about the `ifconfig` command, see the *Network Management Guide*.

Description: `route add default MyRouterBox`

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 MyRouterBox must be in /etc/hosts on the root volume.

Description: `routed on`

Starts the routing daemon. For more information about routing, see the *Network Management Guide*.

Description: `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.

Description: `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.

Description: `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.

Caution

Avoid running `exportfs -a` at a time when applications are likely to be accessing any file system listed in the etc/exports file. Because the `exportfs -a` operation momentarily unexports and then re-exports all the listed file systems, any application attempt to access a temporarily unexported file can result in serious application error.

Description: `nfs on`

Turns on NFS file service. By default, this command is included even if NFS is not licensed for your filer. When the filer runs this command from the /etc/rc file 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.
```

Understanding the options and vol options commands

Commands for setting filer options

The filer recognizes two commands, `options` and `vol options`, for setting options.

Detailed information

The following sections discuss the `options` and `vol options` commands:

- "About the options command" on page 68
- "About the vol options command" on page 69

About the options command

What the options command does

You use the `options` command to display option values or to change the behavior of a filer.

Syntax of the options command

The syntax of the `options` command is as follows. If you omit a value for an option, the command displays the current value of the option. If you omit an option and a value, the command displays the current value for all options.

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

option is the name of the option.

value is the value of the option.

Note

The values of all options except the `Autosupport.doit` option are persistent across reboots. You do not need to include an `options` command in the /etc/rc file to make an `options` setting permanent.

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

Example of the options command

The following command sets the DNS domain name to appliance.com:

```
options dns.domainname appliance.com
```

About the vol options command

What the vol options command does

You use the `vol options` command to configure volume-level behavior. You can use this command only with volume options, which are described in the na_vol(1) man page.

Syntax of the vol options command

The syntax of the `vol options` command is as follows. If you omit a value for an option, the command displays the current value of the option.

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

volname is the name of the volume to which the option applies.

option is the name of the option.

value is the value of the option.

Note

Changes made with the `vol options` command are persistent across reboots. You do not need to include a `vol options` command in the /etc/rc file to make that option setting permanent.

For more information about the `vol options` command, see the na_vol(1) man page.

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
```

Performing common tasks

Detailed information

The following sections discuss commonly performed tasks:

- "Halting the filer" on page 71
- "Booting and rebooting the filer" on page 73
- "Editing the filer's boot configuration file" on page 74
- "Accessing the default directories on the filer" on page 75
- "Recovering from /etc/rc errors" on page 77
- "Synchronizing filer system time" on page 78
- "Setting the filer date and time" on page 83
- "Managing licenses" on page 85
- "Understanding core files" on page 87
- "Understanding large files" on page 89
- "Configuring message and event logging" on page 90
- "Configuring audit logging" on page 94
- "Configuring Telnet sessions" on page 95
- "Maintaining filer security through options" on page 96

Halting the filer

Command to use

Use the `halt` command to perform an orderly shutdown that flushes file system updates to disk and cleans the NVRAM.

Reasons to use the halt command

The filer stores requests it receives in nonvolatile random-access memory (NVRAM). For the following reasons, you should always execute the `halt` command before turning the filer off:

- The `halt` command flushes all data from memory 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.

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 ability to display messages from the filer is built into Windows NT.

Halting the filer

To halt the filer, complete the following step.

Step	Action
1	Enter the following command: **`halt [-d] [-t interval] [-f]`** `-d` causes the filer to perform a core dump before halting. `-t` *interval* causes the filer to halt after the number of minutes specified by *interval*. `-f` prevents one partner in a clustered filer pair from taking over the other after the filer halts. **Result:** The filer displays the following prompt: `ok`

Booting and rebooting the filer

Where the filer boots from

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

Booting the filer

You can boot the filer if it has been turned on or if the filer console is displaying the firmware (`ok`) prompt.

If the filer is turned on and is displaying the command prompt, see the next section, "Rebooting the filer."

To boot the filer, complete the following step.

Step	Action
1	Enter the following command: **`boot`** **Note** If you are booting the filer from a diskette, ensure that boot diskette 1 is in the diskette drive. Use of any other diskette results in an error message.

Rebooting the filer

You can reboot the filer if the filer console is displaying the command prompt. Rebooting the filer is equivalent to halting and booting the filer. During a reboot, the contents of the filer's NVRAM are flushed to disk, and the filer sends a warning message to CIFS clients.

To reboot the filer, complete the following step.

Step	Action
1	Enter the following command: **`reboot`**

Editing the filer's boot configuration file

What the filer's boot configuration file contains

The filer's boot configuration file contains commands that are run automatically whenever you boot the filer. The configuration file is named rc and is in the /etc directory of the filer's default volume (the default is /vol/vol0/etc/rc).

Editing the filer's boot configuration file

To edit the filer's boot configuration file, complete the following steps.

Step	Action
1	Make a backup copy of the /etc/rc file.
2	Edit the /etc/rc file. **Note** Do not add CIFS commands to /etc/rc.
3	Save the edited file.
4	Reboot the filer to test the new configuration. If the new configuration does not work as you want, repeat Steps 2 through 4.

Accessing the default directories on the filer

What the default directories are

The filer root directory in the filer default volume contains two subdirectories:

- The etc directory, which contains configuration files
- The home directory, which, by default, is empty

Accessing the /etc directory from an NFS client

To access the /etc directory from an NFS client, complete the following steps.

Step	Action
1	Mount the following path: **`filer:/vol/vol0`** *filer* is the name of your filer. **Result:** You have access to the filer root directory.
2	Change directories to the /etc directory, using the following command: **`cd etc`**

Accessing the /etc directory from a CIFS client

To access the root directory from a CIFS client, complete the following steps.

Step	Action
1	Map a drive to the following path: **`\\filer\C$`** *filer* is the name of your filer. **Result:** You have access to the filer root directory.
2	Double-click the /etc folder.

Accessing the home directory from an NFS client

To access the home directory from an NFS client, complete the following step.

Step	Action
1	Mount the following path: ***filer*:/vol/vol0/home** *filer* is the name of your filer.

Accessing the home directory from a CIFS client

To access the home directory from a CIFS client, complete the following step.

Step	Action
1	Map a drive to the following path: ***filer*\HOME** *filer* is the name of your filer. **Note** You can also browse the Network Neighborhood to locate the filer and the HOME directory.

Recovering from /etc/rc errors

Why the filer becomes inaccessible to the administration host

The filer can become inaccessible to the administration host if you introduce one of the following errors into the /etc/rc file when you edit it:

- You specify an incorrect network address, using the interface configuration command. The filer is inaccessible because it is not on the network.
- You improperly export filer directories to the NFS client that is the administration host. The filer is inaccessible because you cannot mount the filer root directory on the NFS client.

Recovering from /etc/rc errors

To recover from the error, complete the following steps.

Step	Action	
1	Enter commands on the console to configure the interface with the correct address.	
2	**If you are in...**	**Then...**
	An NFS environment	Enter the `exportfs` command to export the filer root directory to the administration host.
	A CIFS environment	Add a share to the filer root directory.
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.	

Synchronizing filer system time

About synchronizing filer system time

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

Note

To ensure proper interaction between your storage appliance and external entities, such as Windows domain controllers, you must synchronize your appliance's time with the same external source as the domain controllers in the same domain. The lack of time synchronization might make CIFS clients unable to connect to the storage appliance.

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.

Synchronizing filer time

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

Step	Action
1	Specify, by entering the following command, the maximum discrepancy allowed before time is no longer synchronized: **`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
2	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.
3	Specify a synchronization schedule by entering the following command: **`options timed.sched schedule`** *schedule* is one of the following: ◆ hourly—Synchronizes hourly ◆ multihourly—Synchronizes every six hours ◆ daily—Synchronizes every day at midnight ◆ an integer followed by m to specify an interval of minutes or h to specify an interval of hours The default is hourly.
4	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 without spaces.
5	Turn on the `timed` daemon by entering the following command: **`options timed.enable on`**
6	If desired, turn on console logging of time changes initiated by the `timed` daemon by entering the following command: **`options timed.log on`**

Synchronizing the time with the rdate command

The `rdate` command synchronizes the filer date and time to an accuracy within one second of the date and time of the remote host.

How the rdate command communicates with UNIX computers: The `rdate` command uses the UDP protocol (typically, UDP port 37). This service is available on most UNIX computers. For detailed information about the protocol that the `rdate` command uses, see RFC 868.

Note

There is currently no Windows 95, Windows NT, or Windows 2000 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 network.

Use UNIX scripts to synchronize the time periodically: To keep the filer clock accurate, regularly run the `rdate` command on the filer that specifies the time server. A typical scheme is to have a UNIX computer run a periodic `cron` job that executes the appropriate `rdate` command on your filer through a Remote Shell. The computer and the user (if not root) running the `cron` job must be in the filer /etc/hosts.equiv file on the root volume.

To synchronize the date and time with a time server using the `rdate` command, complete the following steps.

Step	Action
1	Check /etc/services and /etc/inetd.conf on the target computer—that is, the computer you choose as the time server—to see whether the UDP time service is supported. If the service is not supported, locate a host that supports the service before continuing to Step 2.
2	Access the filer command line through the console or through a Telnet session.
3	Enter the following command: **`rdate`** ***`host`*** *host* is the name of the time server.

Synchronizing the time with SNTP

SNTP provides the highest degree of time synchronization accuracy. For detailed information about SNTP, see RFC 2030.

The following table describes the options used by the `timed` daemon. See the na_options(1) man page for a full description of the options and their possible values.

Timed option	Function	Default
`timed.log`	Determines whether to log time changes initiated by the `timed` daemon to the console. **Example:** `options timed.log on`	Off
`timed.max_skew`	Sets the maximum allowable discrepancy between filer time and server time. **Example:** `options timed.max_skew 10m`	30m
`timed.proto`	Selects whether to use the protocol used by the `rdate` command or the SNTP protocol (ntp). **Example:** `options timed.proto rdate`	ntp
`timed.sched`	Schedules when to synchronize the time with a time server. **Example:** `options timed.sched daily`	hourly
`timed.servers`	Specifies up to five time servers in order of contact priority. **Example:** `options timed.servers timeser1, timeser2.ntap.com,10.15.46.92`	None

To synchronize the date and time with a time server, complete the following steps.

Step	Action
1	Access the filer command line through the console or through a Telnet session.
2	Set any of the `timed` daemon options as described in the preceding table.
3	Enter the following command to enable the `timed` daemon: **`options timed.enable on`**

Setting the filer date and time

Setting the date and time with the date command

To set the date and time using the `date` command, complete the following steps.

Step	Action
1	Access the filer command line through the console or through a Telnet session.
2	Enter the following command, substituting the current date and time for the number string: **`date [-u] [[[`*`CC`*`]`*`yy`*`]`*`mmddhhmm`*`[.`*`ss`*`]]`** `-u` sets the date and time to Greenwich Mean Time instead of the local time. *CC* is the first two digits of the current year. *yy* is the second two digits of the current year. *mm* is the current month. If the month is omitted, the default is the current month. *dd* is the current day. If the day is omitted, the default is the current day. *hh* is the current hour, using a 24-hour clock. *mm* is the current minute. *ss* is the current second. If the seconds are omitted, the default is 0. **Example:** The following command sets the date and time to 22 May 2002 at 9:25 a.m. `date 200205220925` **Note** If the first two digits of the year are omitted, they default to 20; if all four digits are omitted, they default to the current year. Time changes for daylight saving and standard time, and for leap seconds and years, are handled automatically.

Setting the date while running SnapMirror

Using the `date` or `rdate` command to make the date on a filer go backward when SnapMirror is running can cause snapshots to appear out of sequence. When this occurs, SnapMirror assumes that the older snapshot was created before the newer one, causing SnapMirror to ask for a new, complete transfer before proceeding with any incremental transfers. You can avoid this problem in the following ways:

- Turn SnapMirror off until the filer completes the changes.
- Change the date prior to the next scheduled SnapMirror transfer.

Managing licenses

About license codes

A license code is a string of characters that is unique to a particular filer (such as ABCDEFG). You receive license codes for every protocol and option that you purchase. Not all purchased license codes are installed on a filer before it is shipped from the factory; some must be installed after the system is set up. You can purchase license codes to enable additional protocols and options at any time.

License code replacement

If you misplace a license code, you can contact Network Appliance Technical Support or log in to the NOW site at http://now.netapp.com to obtain a copy.

License management tasks

You can perform the following tasks to manage licenses:

- Add licenses
- Display current license codes
- Delete licenses

Enabling a license

To add the license for a protocol or option, complete the following step.

Step	Action
1	Enter the following command: **`license add <code1> <code2> ...`** *code* is the license code provided to you by Network Appliance.

Displaying current license codes

To display licensing information for all your protocols and options, complete the following step.

Step	Action
1	Enter the following command without parameters: `license` **Result:** Data ONTAP displays a list of the licenses that are enabled and their codes.

Disabling a license

To disable a license for a protocol or option, complete the following step.

Note
You cannot disable licenses for the disk sanitization or SnapLock features after you enable them.

Step	Action

1	Enter the following command:
	`license delete` ***`service`***
	service is one of the following:
	`cifs`
	`cluster`
	`cluster_remote`
	`dafs`
	`fcp`
	`http`
	`iscsi`
	`multistore`
	`nfs`
	`smdomino`
	`smsql`
	`snaplock_enterprise`
	`snapmirror`
	`snapmirror_sync`
	`snapmover`
	`snaprestore`
	`sv_linux_pri`
	`sv_ontap_pri`
	`sv_ontap_sec`
	`sv_unix_pri`
	`sv_windows_pri`
	`syncmirror_local`
	`vld`

Understanding 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. 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, does the following things:

- 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 writing

A core dump file contains the contents of memory and NVRAM. Core dumps are written over reserved sections any nonbroken disk owned by the local filer. When a core dump is created, it is stored in uncompressed format if sufficient space is available; otherwise, it is stored in compressed format. If there is insufficient space to store a complete core dump in compressed format, the core dump is canceled.

Note

If the failed filer belongs to a cluster and `options cf.takeover.on_panic` is enabled, a core dump file is written to a spare disk on that filer.

Core dump files are not compatible between Data ONTAP releases because where the core starts on disks depends on the release. Because of this incompatibility, Data ONTAP might fail to find a core dump file dumped by another release.

Automatic technical support notification

Your filer sends e-mail automatically to Network Appliance Technical Support upon each system reboot if the Autosupport feature is enabled and configured correctly. Network Appliance Technical Support uses the Autosupport message and the core file to troubleshoot the problem.

Manual technical support notification

If you have disabled automatic e-mail, you should contact Network Appliance Technical Support when your system creates a core file.

Understanding large files

Maximum size of 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

When a displayed file size is incorrect

A file size is displayed incorrectly 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 file size on your server. In this case, the maximum file size appears as 2 GB, even though the file on the client is larger than 2 GB.

Software requirements

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

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

Configuring message and event logging

What message logging is

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

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

What event logging is

The filer maintains a log of system events in the /etc/log/ems file on the filer's root volume. For more information about configuring event logging and the filer's Event Management System (EMS) mechanisms, see the na_ems(1) man page.

Where messages are sent

Message logging is done by a `syslogd` daemon. The /etc/syslog.conf configuration file on the filer's root volume determines how system messages are logged. Depending on their severity and origin, messages can be sent to

- The console
- A file
- A remote system

By default, all system messages (except those with debug-level severity) are sent to the console and logged in the /etc/messages file.

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; therefore, you can have up to seven message files (/etc/messages.0 through /etc/messages.5 and the current /etc/messages file).

Checking the /etc/messages file daily

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

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 most to least severe.

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

Level	Description
*	All levels of errors

The action parameter

The `action` parameter specifies where to send messages. 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 *

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 of levels emerg, alert, crit,
# and err to /etc/messages.
kern.err                        /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
```

Configuring message logging

To configure message logging, complete the following steps.

Step	Action
1	Open the /etc/syslog.conf file with an editor from a client.
2	Add one or more lines using the following format: `facility.level <tab> destination`
3	Save and close the /etc/syslog.conf file.

For more information

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

Configuring audit logging

About the audit log

An audit log is a record of commands executed at the console, through a Telnet shell, or by using the `rsh` command. The audit log data is logged in the /etc/log directory in a file called auditlog. Administrative HTTP operations, such as those resulting from the use of FilerView, are also logged. The maximum size of the auditlog file is specified by the `auditlog.max_file_size` option. By default, Data ONTAP is configured to save an audit log.

The /etc/log/auditlog file restart schedule

Every Saturday at midnight, the /etc/log/auditlog file is copied to /etc/log/auditlog.0, /etc/log/auditlog.0 is copied to /etc/log/auditlog.1, and so on. This also occurs if the auditlog file reaches the maximum size specified by `auditlog.max_file_size`.

Assuming they do not get full, the system saves auditlog files for six weeks.

Configuring audit logging

To configure audit logging, complete the following steps.

Step	Action
1	If audit logging is turned off, enter the following command to turn audit logging on: **`options auditlog.enable on`** Otherwise, go to Step 2.
2	To change the maximum size of the audit log file, enter the following command: **`options auditlog.max_file_size`** ***`value`*** *value* is the maximum size in bytes. The default value is 10000000 (10 MB).

Configuring Telnet sessions

What you can configure

You can configure whether Telnet sessions have a timeout period and how long that timeout period is. By default, Telnet sessions have timeout periods of 60 minutes.

Changing the timeout period

To change the timeout period for Telnet sessions, complete the following step.

Step	Action
1	Enter the following command: **`options autologout.telnet.timeout`** ***`minutes`*** *minutes* is the length of the timeout period.

Disabling the timeout period

To disable the timeout period for Telnet sessions, complete the following step.

Step	Action
1	Enter the following command: **`options autologout.telnet.enable off`**

When options command values take effect

A console or Telnet session must be reestablished before any of the following `options` command values take effect:

- `autologout.console.enable`
- `autologout.console.timeout`
- `outlogout.telnet.enable`
- `outlogout.telnet.timeout`

Maintaining filer security through options

Options that help maintain security

The following table describes options that help you to maintain filer security.

Option	Description
`trusted.hosts`	Controls which hosts can access the filer through a Telnet session for administrative purposes. You can disable Telnet access completely, allow Telnet access from up to five hosts, or allow Telnet access from all hosts.
`telnet.access`	Controls which hosts can access the filer through a Telnet session for administrative purposes. You can restrict Telnet access to the filer by specifying host names, IP addresses, or network interface names.
`rsh.access`	Controls which hosts can access the filer through a Remote Shell session for administrative purposes. You can restrict Remote Shell access to the filer by specifying host names, IP addresses, or network interface names.
`nfs.mount_rootonly`	Controls whether the filer's volumes can be mounted from NFS clients only by the root user on privileged ports (ports 1 through 1,023) or by all users on all ports. This option is applicable only if the NFS protocol is licensed.
`wafl.root_only_chown`	Controls whether all users or only the root user can change directory and file ownership. This option is applicable only if the NFS protocol is licensed.

Option	Description
`cifs.restrict_anonymous.enable`	Controls whether anonymous CIFS users can look up CIFS shares, users, or groups on a filer. This option is applicable only if the CIFS protocol is licensed.

Note

For more information about the options in this table, see the na_protocolaccess(8) man page.

Controlling Telnet access using host names

To limit Telnet access to specific hosts, complete the following steps.

Step	Action	
1	Access the filer command line through the console or through a Telnet session.	
2	**If...**	**Then...**
	You want to disable Telnet access for all hosts	Enter the following command: **`options trusted.hosts -`**
	You want to restrict Telnet access to up to five hosts	Enter the following command: **`options trusted.hosts host1[, ..., host5]`**
	You want to allow Telnet access for all hosts	Enter the following command: **`options trusted.hosts *`**

Controlling Telnet access using host names, IP addresses, and network interface names

To limit Telnet access to specific host names, IP addresses, and network interface names, complete the following steps.

Step	Action
1	Access the filer command line through the console or through a Telnet session.
2	Enter the following command: **`options telnet.access host=[ hostname | IP_address | interface_name ]`** *hostname* is the name of the host to which you want to allow Telnet access. *IP_address* is the IP address of the host to which you want to allow Telnet access. *interface_name* is the network interface name of the host to which you want to allow Telnet access. **Note** If this option is set, the trusted.hosts option is ignored for Telnet.

Note

For more information on controlling Telnet access to a filer using multiple host names, IP addresses, and network interfaces, see the na_protocolaccess(8) man page.

Controlling Remote Shell access

To allow Remote Shell access to a filer from a single host, complete the following steps.

Step	Action
1	Access the filer command line through the console or through a Telnet session.

Step	Action
2	Enter the following command: **`options rsh.access host=`*`hostname`*** *hostname* is the name of the host to which you want to allow Remote Shell access.

Note

For information on controlling Remote Shell access to a filer from multiple hosts using the host name, IP address or network interface name, see the na_protocolaccess(8) man page.

Controlling mount privilege

To control the NFS mount privilege for the filer's volumes, complete the following steps.

Note

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

Step	Action	
1	Access the filer command line through the console or through a Telnet session.	
2	**If...**	**Then...**
	You want to restrict the mount privilege to only the root user using privileged ports (ports 1 through 1,024)	Enter the following command: **`options nfs.mount_rootonly on`**
	You want to allow the mount privilege for all users on all ports	Enter the following command: **`options nfs.mount_rootonly off`**

Controlling file ownership change privileges

To control who has privileges to change directory and file ownership, complete the following steps.

Note

The following behaviors apply to ownership changes:

- When a nonroot user changes the owner of a file, the set-user-id and set-group-id bits are cleared.
- If a nonroot user tries to change the owner of a file but the change causes the file's recipient to exceed the quota, the attempt fails.

Step	Action	
1	Access the filer command line through the console or through a Telnet session.	
2	**If...**	**Then...**
	You want to restrict the privilege of changing directory and file ownership to the root user	Enter the following command: `options wafl.root_only_chown on`
	You want to allow the privilege of changing directory and file ownership to all users	Enter the following command: `options wafl.root_only_chown off`

Controlling anonymous CIFS share lookups

To control whether anonymous CIFS users can look up CIFS shares, users, or groups on a filer, complete the following steps.

Step	Action
1	Access the filer command line through the console or through a Telnet session.

Step	Action	
2	**If...**	**Then...**
	You want to restrict anonymous share lookups	Enter the following command: `options cifs.restrict_anonymous.enable on`
	You want to allow anonymous share lookups	Enter the following command: `options cifs.restrict_anonymous.enable off`

Filer configuration backup and cloning

About filer configuration backup and cloning

Filer configuration backup and cloning is performed using the `config` command. The configuration backup operation stores the filer's configuration information in a single file with a name you specify. The configuration backup file enables you to restore the filer configuration in case of disasters or emergencies. Configuration cloning enables you to clone the configuration of an existing filer to a new filer.

Backing up a filer configuration

When you use the `config` command to back up a filer configuration, the filer configuration is saved in a single file with a file name that you specify. By default, backup configuration files are created in the /etc/configs directory.

To back up a filer configuration, complete the following step.

Step	Action

1	Enter the following command:
	`config dump [-f] [-v] config_file`
	`-f` forces the new file to override an existing backup.
	`-v` causes Data ONTAP to also back up a volume-specific configuration.
	config_file is the name or the path and name of the backup file you are creating.
	Example 1: The following is an example of the `config dump` command using the default directory to back up a filer-specific configuration to the file /etc/configs/08_02_2001.
	`config dump 08_02_2001`
	Example 2: The following is an example of the `config dump` command with a directory that you specify.
	`config dump /home/users/08_02_2001`

Cloning a filer configuration

To clone the configuration of one filer to another, complete the following step.

Step	Action
1	Enter the following command: **`config clone filer username:password`** *filer* is the name of the remote filer from which you want to clone the configuration. *username* is the remote filer user name. *password* is the remote user password. **Example:** The following is an example of the `config clone` command cloning the tpubs-dot configuration to the filer toaster. `toaster> config clone tpubs-dot root:hello`

About restoring a filer configuration

You can restore a filer configuration using the backup configuration files you created with the `config dump` command.

Restoring a filer configuration

To restore filer configuration information from a backup configuration file, complete the following steps.

Step	Action
1	Enter the following command: **`config restore [-v] config_file`** `-v` enables you to restore volume-specific configuration files, as well as filer specific configuration files. **Example:** The following is an example of the `config restore` command restoring the backup configuration file from the default /etc/configs directory. `config restore 08_02_2001`
2	Reboot the filer to run commands in the /etc/rc file.

Comparing filer configurations and backup configuration files

To compare a filer's current configuration with a backup configuration file, or to compare differences between two backup configuration files, complete the following step.

Step	Action

1	Enter the following command:
	`config diff [-o output_file] {config_file1} {config_file2}`
	`-o` is the option you use to create the output file.
	output_file is the name of the file containing the differences.
	config_file1 is the name of the first file you want to compare.
	config_file2 is the name of the second file you want to compare.
	Example 1: The following example compares the filer's current configuration with the configuration information in the backup file.
	`config diff 11_15_2001`
	Example 2: The following example compares the configuration information in two backup files.
	`config diff 11_05_2001 11_15_2001`

Data Organization Using Volumes and Qtrees 3

About this chapter

This chapter describes how to use volumes and qtrees to manage user data. Read this chapter if you plan to organize user data into large units (volumes) for reliability, or smaller units (qtrees) for flexibility.

Topics in this chapter

This chapter discusses the following topics:

- "Understanding data management" on page 106
- "Creating volumes" on page 112
- "Choosing a language for a volume" on page 117
- "Changing the state of a volume" on page 122
- "Physically moving volumes" on page 125
- "Renaming volumes" on page 128
- "Changing the root volume" on page 130
- "Destroying volumes" on page 131
- "SnapLock storage solution" on page 133
- "Displaying qtree status" on page 147
- "Making a directory into a qtree or a qtree into a directory" on page 148
- "Deleting qtrees" on page 151
- "Increasing the maximum number of files in a volume" on page 152
- "Setting space reservation for files" on page 154

Understanding data management

Data access management with volumes and qtrees

You manage data access by using both volumes (physical file systems) and qtrees (logical file systems within a volume) to segregate files that use UNIX security from files that use Windows NT security. You can also enable access to both kinds of files from both UNIX and Windows NT, depending on the needs of your users.

Data storage management

You manage data storage by managing volumes. You can have more than one volume. You can create additional volumes to ensure data integrity and to manage backups and restores. You can also create additional volumes to accommodate your users if they use more than one language. For example, you can have one volume for files in French and another for files in German.

Data resource management with quotas

You can use quotas to limit the number of files and the data space available to users. Using quotas is described in Chapter 8, “Disk Space Management Using Quotas,” on page 333.

Detailed information

The following topics discuss the major activities involved in managing data on a filer:

- “Understanding volumes and qtrees” on page 107
- “Volume structure” on page 110
- “Creating volumes” on page 112
- “Choosing a language for a volume” on page 117
- “Changing the language of a volume” on page 120

Understanding volumes and qtrees

What volumes and qtrees are

A volume is a file system on a filer. It is physically defined by the disks that are assigned to it.

Root volumes: Each filer has a special volume known as a root volume. The root volume contains the filer's configuration files, and a file known as the /etc/rc file, which contains startup commands. There can be only one root volume on a filer.

Data volumes: In addition to the root volume, you can create one or more data volumes, which hold the user visible directories and files that you want to store.

qtrees: A qtree is a logically defined file system that can exist as a special subdirectory of the root directory of a volume.

When creating volumes is appropriate

You can create up to 199 additional volumes when you want to treat large amounts of data as a unit. You create one or more additional volumes under the following circumstances:

- A volume has more disks or storage than it is convenient to manage as one unit.
- You want to accommodate users with different security styles or oplocks requirements that current volumes do not meet and you do not want to create a qtree for that purpose.
- You have a multilingual environment and want to provide one volume for each language and character set.

If your filer has a large number of disks, creating volumes and having a multivolume configuration can also aid in administration. Having multiple volumes enables you to do the following:

- Perform administrative and maintenance tasks, such as backup and restore, on individual volumes rather than on a single, large file system.
- Set `vol option` command values, such as `snap sched`, `raidsize`, `minra`, and `no_atime_update`, differently for individual volumes.
- Take individual volumes offline—to perform administrative tasks on them, for example—without interrupting data availability of the other volumes.
- Have more flexible quota and snapshot configurations.

- Have backups and restores of individual volumes, which take less time.

However, there is a trade-off: You must use additional disks as parity disks and as hot spares rather than for data storage.

When creating qtrees is appropriate

You create qtrees for amounts of data that do not expand the size of volumes to the point that the volume requires additional disks. You can set up qtrees and remove them easily. However, there is a trade-off. Because qtrees are part of a volume, they inherit many volume-level settings.

Detailed comparison

Qtrees and volumes have other differences and similarities that can play a role in your decision, as shown in the following table.

<table>
<tr><th>Qtrees</th><th>Volumes</th></tr>
<tr><td colspan="2">

Both enable you to organize user data.
Both enable you to group users with similar needs.
Both have a security style and an oplocks setting.

The security style determines whether files use UNIX or Windows NT permissions.
The oplocks setting determines whether files and directories use CIFS opportunistic locks.

Both can be used as units in a backup and a restore.
Neither can be crossed by hard links.

</td></tr>
<tr><td>You can add and delete qtrees on the fly; this provides flexibility.</td><td>You can add volumes by adding disks; this makes more storage available on the filer.</td></tr>
<tr><td>You can expand and shrink qtrees.</td><td>You can expand but not shrink volumes.</td></tr>
<tr><td>You cannot create individual qtree snapshots unless you use SnapMirror. You are dependent on volume snapshots.</td><td>You can create volume snapshots, which enable you to recover and mirror data.</td></tr>
</table>

Understanding a volume's RAID structure

About RAID groups in a volume

A volume consists of disks assembled into one or multiple RAID groups. A RAID group consists of one or more data disks, across which client data is striped and stored, plus one or two parity disks that ensure data recoverability even if as many as two disks within the RAID group fail. The purpose of a RAID group is to provide parity protection of data across its included disks from loss due to single or double disk failure within that group.

Types of RAID protection

Data ONTAP supports two types of RAID protection, which you can assign on a per-volume basis: RAID 4 and RAID DP.

RAID 4: Before Data ONTAP 6.5, RAID 4 was the only RAID protection for Data ONTAP volumes. Within its RAID groups, it allots a single disk for holding parity data, which ensures against data loss due to a single disk failure within a group.

RAID DP (double-parity): RAID DP provides a higher level of RAID protection for Data ONTAP volumes. Within its RAID groups, it allots two disks for holding parity and double-parity data. Double-parity protection ensures against data loss due to a double disk failure within a group.

For more information on RAID 4 and RAID DP, see "Types of RAID protection" on page 160.

Size of RAID groups

When you create a volume, you can also control the size of a RAID group. Generally, larger RAID groups maximize your data storage space (by providing a greater ratio of data disks to parity disks) and maximize system performance (data striping across a larger number of disks). Smaller RAID groups minimize the risk of multiple disk failure and the performance impact during the subsequent reconstruction. For information on RAID size considerations, see "Larger versus smaller RAID groups" on page 164.

Volume structure

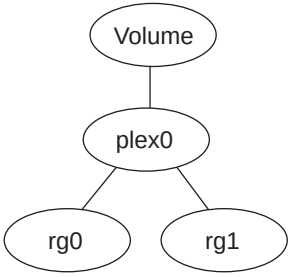
Unmirrored and mirrored volumes

A volume can be unmirrored or mirrored. An unmirrored volume has one plex; a mirrored volume has two plexes that are updated simultaneously. A plex contains all of the data in a volume. Therefore, a mirrored volume contains two copies of the data. This chapter discusses unmirrored volumes unless mirroring is explicitly stated.

You must have a SyncMirror license to use mirror capabilities. See the *Data Protection Guide* for more information about mirrored volumes.

What a plex is

A plex is a physical copy of a file system and consists of one or more RAID groups. The following diagram illustrates a plex.



How volume plexes and RAID groups are named

The syntax that Data ONTAP uses to name plexes is as follows:

vol-name/plex*plexid*

vol-name is the name of the volume.

plexid is the number of the plex, which is assigned by Data ONTAP.

The syntax that Data ONTAP uses to name RAID groups is as follows:

vol-name/plex*plexid*/rg*raid-group-id*

raid-group-id is the number of the RAID group.

Examples: vol0 specifies all the plexes and RAID groups of volume vol0.

vol0/plex1 specifies all the RAID groups of plex 1 of volume vol0.

vol0/plex1/rg2 specifies RAID group 2 of plex 1 of volume vol0.

Creating volumes

About creating volumes

When you create a volume, you must provide the following information:

- A name for the volume.
- Data disks to include in the volume. You can specify these either by disk ID or by the number of disks of a specified size.

  Disks are formatted as either block or zoned checksum type. Spare disks are listed according to checksum types. Rules for using checksum types are in the next section, "Checksum type usage."

You can accept the defaults for the following items or specify other values:

- RAID group size—The default for volumes using RAID 4 protection is 6 on R150 platforms, 7 on FAS270 and FAS270c platforms and 8 on all other platforms. The Default for RAID DP protection is 12 on R200 and R150 platforms and 14 on all other platforms.
- RAID type—The default on NearStore appliances is RAID DP. The default on all other filers is RAID 4.
- The language used for file names—The default is the language of the root volume.

Checksum type usage

A checksum type applies to an entire volume. A volume can have only one checksum type.

Checksum type for root volume on new filer: When you run the `setup` command on a new filer, the checksum type for the root volume is chosen for you.

Checksum type for new volumes: When you create a volume, the filer determines the checksum type of the volume.

- If enough block checksum disks are available, the volume uses block checksums.
- Otherwise, the volume uses zoned checksums.

To use block checksums when you create a new volume, you must have at least the same number of block checksum spare disks available as you specify in the `vol create` command.

No checksum type conversion: You cannot convert a disk from one checksum type to the other, and you cannot convert a volume from one checksum type to another.

Checksum types for existing volumes: The following table shows the types of disks that you can add to an existing volume of each type.

Disk type	Block checksum volume	Zoned checksum volume
Block checksum	OK to add	OK to add
Zoned checksum	Not OK to add	OK to add

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

Creating a volume

To create a volume, complete the following steps.

Step	Action
1	View a list of the spare disks on your NetApp appliance. The spare disks are the disks that are available for you to assign to the volume that you want to create. In the filer console, enter the following command: `vol status -s` **Result:** The output of `vol status -s` lists all the spare disks that you can select for inclusion in the volume and their capacities. **Note** If you are setting up volumes on an FAS270c appliance with two internal filer heads, you might have to assign the FAS270c disks to one or the other of the two filers before creating volumes on those filers. See “Disk ownership assignment on an FAS270/FAS270c appliance” on page 247.

Step	Action
2	At the system prompt, enter the following command: **`vol create vol-name [-r raid-group-size] [-t {raid4 | raid_dp}][-l language_code] disk-list`** *vol-name* is the name for the new volume (without the /vol/ prefix). *raid-group-size* is the RAID group size you want. {raid4|raid_dp} is the type of RAID parity protection that you want to apply to this volume. **Note** For a complete description of RAID related options for `vol create` see "Specifying the RAID type at volume creation" on page 165 *language_code* specifies a language other than that of the root volume. See "Language codes" on page 117. *disk-list* is one of the following: ◆ *ndisks*[*@disk-size*] *ndisks* is the number of disks to use. It must be at least 2 (3 if RAID DP is configured). *disk-size* is the disk size to use, in gigabytes. You must have at least *ndisk* available disks of the size you specify. ◆ `-d` *disk_name1 disk_name2... disk_nameN* *disk_name1*, *disk_name2*, and *disk_nameN* are disk IDs of one or more available disks; use a space to separate multiple disks.
3	Enter the following command to verify that the volume exists as you specified: **`vol status vol_name -r`** *vol_name* is the name of the volume whose existence you want to verify. **Result:** The system displays the RAID groups and disks of the specified volume on your filer.

Step	Action	
4	**If you have...**	**Then...**
	An NFS-only filer	Go to Step 6.
	A CIFS-only or multiprotocol filer	Create shares that refer to the new volume to enable clients to access it, then go to Step 5.
5	**If you have...**	**Then...**
	A CIFS-only filer	You are done.
	A multiprotocol filer	Go to Step 6.
6	Update the system /etc/exports file with the new volume.	
7	Run `exportfs`.	
8	Add the appropriate mount point information to the /etc/fstab or /etc/vfstab file on clients that mount volumes from the filer.	

Volume creation example

The following command creates a volume called newvol, with no more than eight disks in a RAID group, using the French character set, consisting of the disks with disk IDs 8.1, 8.2, 8.3, and 8.4.

```
vol create newvol -r 8 -l fr -d 8.1 8.2 8.3 8.4.
```

Parameters to accept or change after volume creation

After you create a volume, you can accept the defaults for CIFS oplocks and security style settings or change the values. You should decide what to do as soon as possible after creating the volume. If you change the parameters after files are in the volume, they might become inaccessible to users because of conflicts between the old and new values. For example, UNIX files available under mixed security might not be available after you change to NTFS security.

CIFS oplocks setting: The CIFS oplocks setting determines whether the volume uses CIFS oplocks. The default is to use CIFS oplocks.

Security style: The security style determines whether the files in a volume use NTFS security, UNIX security, or both.

When you have a new filer, the default depends on what protocols you licensed, as shown in the following table.

Protocol licenses	Default volume security style
CIFS only	NTFS
NFS only	UNIX
CIFS and NFS	UNIX

When you upgrade to Data ONTAP 6.1 from Data ONTAP 5.0 or earlier, the default security style is multiprotocol.

When you change the configuration of a filer from one protocol to another, the default security style for new volumes changes as shown in the following table.

From	To	Default for new volumes	Note
NTFS	Multiprotocol	UNIX	The security styles of volumes are not changed.
Multiprotocol	NTFS	NTFS	The security style of all volumes is changed to NTFS.

Choosing a language for a volume

About volumes and languages

Every volume has a language. The filer uses a character set appropriate to the language. You can specify a language for each volume if you do not want to use the language of the root volume. You can specify the language of a volume when you create the volume or change it later.

The language you specify determines which character set the filer uses for the following names:

- File names
- User names
- Share names
- System and domain names

Note

Names of the following objects must be in ASCII characters:

- Qtrees
- Snapshots
- Volumes

Language codes

You use a language code to specify a language, as described in "Creating a volume" on page 113.

The filer supports the languages shown in the following table. The language code for each language appears next to the language. If you want to override the normal UNIX character set so that a UNIX system uses UTF-8 as the NFS character set, add UTF-8 to the language code.

Language	Language code
Chinese (Simplified)	zh
Chinese (Simplified - GBK)	zh.GBK
Chinese (Traditional - euc-tw)	zh_TW
Chinese (Traditional - Big 5)	zh_TW.BIG5

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
Korean	ko
Norwegian	no
Portuguese	pt
POSIX (ASCII only for DOS and ISO-Latin1 for NFS. Equivalent to the Sun locale.)	C
Spanish	es
Swedish	sv

UTF-8 Examples: The following table shows examples of language codes for UTF-8.

Language	Language code	Language code for UTF-8
Danish	da	da.UTF-8
Japanese PCK (sjis)	ja_JP.PCK	ja_JP.PCK.UTF-8

Viewing the language list online

It might be useful to view the list of languages before you choose one for a volume. To view the list of languages, complete the following step.

Step	Action
1	Enter the following command: `vol lang`

Choosing a language for a volume

To choose a language for a volume, complete the following step.

Step	Action	
1	**If the volume is used for...**	**Then...**
	NFS only	Do nothing; the language does not matter.
	CIFS only	Set the language of the volume to the language of its clients.
	Both NFS and CIFS	Set the language of the volume to the locale used by NFS.

Changing the language of a volume

When to change the language of a volume

You change the language of a volume to make it available to users of a language other than the default language. You should do this before any files are created in the volume so that all file names use the same language.

Caution

Changing the language after the volume contains files can cause some NFS encodings to be invalid. The file names are then unreadable and can't be found, making the files inaccessible.

Displaying volume language use

You can display a list of volumes with the language each volume is configured to use. This is useful for the following kinds of decisions:

- How to match the language of a volume to the language of clients.
- Whether to create a volume to accommodate clients that use a language for which you don't have a volume.
- Whether to change the language of a volume (usually from the default language).

To display which language a volume is configured to use, complete the following step.

Step	Action

1	Enter the following command:
	`vol status [`*`vol-name`*`] -l`
	vol-name is the name of the volume about which you want information. Leave out *vol-name* to get information about every volume on the filer.
	Result: Each row of the list displays the name of the volume, the language code, and the language, as shown in the following sample output.
	`Volume Language`
	`vol0   ja (Japanese euc-j)`

Changing the language

To change the language that a volume uses to store file names, complete the following steps.

Step	Action
1	Enter the following command: **`vol lang vol-name language`** *vol-name* is the name of the volume about which you want information. *language* is the code for the language you want the volume to use.
2	Enter the following command to verify that the change has successfully taken place: **`vol status vol-name -l`** *vol-name* is the name of the volume whose language you changed.

Changing the state of a volume

Volume states

A volume can be in one of the following three states:

- Offline—Read or write access is not allowed.
- Restricted—Some operations, such as copying volumes and parity reconstruction, are allowed, but data access is not allowed.
- Online—Read and write access is allowed. An online volume can be further described as follows:
  - Normal—The volume is unmirrored and all of its RAID groups are functional.
  - Mirrored—The volume is mirrored and all of its RAID groups are functional.
  - Degraded—The volume contains at least one degraded RAID group that is not being reconstructed after single disk failure.
  - Double degraded—The volume contains at least one RAID group with double disk failure that is not being reconstructed (this state is possible if RAID DP protection is enabled for the affected volume).
  - Mirror degraded—The volume is a mirrored volume and one of its plexes is offline or resynchronizing.
  - Foreign—Disks that the volume contains were moved to the current filer from another filer.
  - Reconstruction xx% complete—At least one RAID group in the volume is being reconstructed after experiencing single disk failure.
  - Double reconstruction xx% complete—At least one RAID group in the volume is being reconstructed after experiencing double disk failure (this state is possible if RAID DP protection is enabled for the affected volume).
  - Partial—At least one disk was found for the volume, but two or more disks are missing.
  - Growing—Disks are in the process of being added to the volume.
  - Initializing—The volume is in the process of being initialized.
  - Needs check—WAFL consistency check needs to be performed on the volume.
  - Resyncing—The volume contains two plexes; one plex is resynchronizing with the volume.

- Ironing—A WAFL consistency check is being performed on the volume.
- Copying—The volume is currently the target volume of active `vol copy` or `snapmirror` command operations.
- Snapmirrored—The volume is a snapmirrored replica of another volume.
- Verifying—Mirror verification operation is currently running on the volume.
- Wafl inconsistent—The volume has been marked corrupted; contact Network Appliance Technical Support.
- Invalid—The volume does not contain a valid file system.

Determining the state of volumes

To determine what state a volume is in, complete the following step.

Step	Action
1	Enter the following command: **`vol status`**

When to take a volume offline

You can take a volume offline and make it unavailable to the filer. You do this for the following reasons:

- To perform maintenance on the volume
- To move a volume to another filer
- To destroy a volume

Taking a volume offline

To take a volume offline, complete the following step.

Step	Action
1	Enter the following command: **`vol offline`** ***`vol-name`*** *vol-name* is the name of the volume to be taken offline.

When to make a volume restricted

When you make a volume restricted, it is available for only a few operations. You do this for the following reasons:

- To copy a volume to another volume (see the *Data Protection Guide* for volume copy details)
- To perform a level-0 SnapMirror operation (see the *Data Protection Guide* for SnapMirror operation details)

Restricting a volume

To restrict a volume, complete the following step.

Step	Action
1	Enter the following command: `vol restrict vol-name` *vol-name* is the name of the volume to be made restricted.

Bringing a volume online

You bring a volume online to make it available to the filer after you deactivated that volume. To bring a volume back online, complete the following step.

Step	Action
1	Enter the following command: `vol online vol-name` *vol-name* is the name of the volume to reactivate. If the volume is inconsistent, the command prompts you for confirmation. **Note** If you bring an inconsistent volume online, it might suffer further file system corruption.

Physically moving volumes

About physically moving volumes

You can physically move volumes from one filer to another. You might want to move a volume to a different filer to perform one of the following tasks:

- Replace a disk shelf with one that has a greater storage capacity
- Replace current disks with larger disks
- Gain access to the files on disks on a malfunctioning filer

You can physically move disks, disk shelves, or loops to move a volume from one filer to another. You need the manual for your disk shelf to move a volume.

The following terms are used:

- The *source filer* is the filer from which you are moving the volume.
- The *destination filer* is the filer to which you are moving the volume.
- The volume you are moving is a *foreign volume* to the destination filer.

Note

If MultiStore and SnapMover licenses are installed, you might be able to move volumes without moving the drives on which they are located. See the MultiStore guide for more information.

Moving a volume

To physically move a volume, complete the following steps.

Step	Action
1	Enter the following command at the source filer to locate the disks that contain the volume: `vol status vol_name -r` **Result:** The locations of the data and parity disks in the volume appear under the volume name on the same line as the labels Data and Parity.
2	Place the volume offline and follow the instructions in the disk shelf hardware guide to remove the disks from the source filer.
3	Halt and turn off the destination filer.

Step	Action
4	Install the disks in a disk shelf connected to the destination filer.
5	Turn on and boot the destination filer. **Result:** When the destination filer boots, it places the foreign volume offline. 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 causes them to become hot spare disks.
6	If the filer renamed the foreign volume because of a name conflict, enter the following command to rename the volume: **`vol rename vol-name new-name`** *vol-name* is the name of the volume you want to rename. *new-name* is the new name of the volume. **Example:** The following command renames the users(1) volume as newusers: `vol rename users(1) newusers`
7	Enter the following command to bring the volume online in the new filer: **`vol online vol-name`** *vol-name* is the name of the volume. **Result:** The following events occur: ◆ The volume is online in its new location. ◆ Entries in the /etc/exports file and shares 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.

Step	Action	
8	Enter the following command to confirm that the added volume came online: **`vol status`** ***`vol-name`*** *vol-name* is the name of the volume.	
9	**If you have...**	**Then...**
	An NFS-only filer	Go to Step 11.
	A CIFS-only or multiprotocol filer	Create shares that refer to the new volume to enable clients to access it, then go to Step 10.
10	**If you have...**	**Then...**
	A CIFS-only filer	You are done.
	A multiprotocol filer	Go to Step 11.
11	Update the system /etc/exports file with the new volume.	
12	Run `exportfs`.	
13	Add the appropriate mount point information to the /etc/fstab or /etc/vfstab file on clients that mount volumes from the filer.	

Renaming volumes

Reasons for renaming volumes

You usually rename volumes to give them descriptive names. In particular, you might want to change the root volume which has the undescriptive default name of /vol0.

Renaming a volume

To rename a volume, complete the following steps.

Step	Action	
1	Enter the following command: **`vol rename vol-name new-name`** *vol-name* is the name of the volume you want to rename. *new-name* is the new name of the volume. **Result:** The following events occur: ◆ The volume is renamed. ◆ Entries in the /etc/exports file and shares that refer to the volume become incorrect and must be updated. ◆ The in-memory information about active exports gets updated automatically, and clients continue to access the exports without problems.	
2	**If you have...**	**Then...**
	An NFS-only filer	Go to Step 4.
	A CIFS-only or multiprotocol filer	Create shares that refer to the new volume to enable clients to access it, then go to Step 3.
3	**If you have...**	**Then...**
	A CIFS-only filer	You are done.
	A multiprotocol filer	Go to Step 4.
4	Update the system /etc/exports file with the new volume name.	

Step	Action
5	Run `exportfs`.
6	Add the appropriate mount point information to the /etc/fstab or /etc/vfstab file on clients that mount volumes from the filer.

Changing the root volume

About changing the root volume

You can move a filer root volume to another filer or move the data in a root volume and use the disks as spares. But every filer must have a root volume; therefore, you must designate another volume to be the new root volume.

Changing the root volume of a filer

To change the root volume of a filer, complete the following steps.

Step	Action
1	Create the new root volume using the `vol create` command.
2	Move the /etc directory from the current root volume to the new root volume.
3	Enter the following command: **`vol options`** ***`vol-name`*** **`root`** *vol-name* is the name of the new root volume.
4	Enter the following command to reboot the filer: **`reboot`**

Destroying volumes

About destroying volumes

There are two reasons to destroy a volume:

- You no longer need the data it contains.
- You no longer need the volume because you copied its data to reside elsewhere.

When you destroy a volume, you convert its parity disk and all its data disks back into hot spares. You can then use them in other volumes and other filers.

Caution

If you destroy a volume, all the data in the volume is destroyed and no longer accessible.

Note

SnapLock volumes cannot be destroyed.

Destroying a volume

To destroy a volume, complete the following steps.

Step	Action
1	Take the volume offline by entering the following command: `vol offline vol-name` *vol-name* is the name of the volume that you intend to destroy and whose disks you are converting to hot spares.

<table>
<tr><th>Step</th><th colspan="2">Action</th></tr>
<tr><td>2</td><td colspan="2">Enter the following command to destroy the volume:
<tt>vol destroy vol-name</tt>
vol-name is the name of the volume that you intend to destroy and whose disks you are converting to hot spares.
Result: The following events occur:
◆ The volume is destroyed.
◆ Entries in the /etc/exports file and shares 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.</td></tr>
<tr><td rowspan="2">3</td><td>If you have...</td><td>Then...</td></tr>
<tr><td>An NFS-only filer</td><td>Go to Step 5.</td></tr>
<tr><td>A CIFS-only or multiprotocol filer</td><td>Create shares that refer to the new volume to enable clients to access it.</td></tr>
<tr><td rowspan="3">4</td><td>If you have...</td><td>Then...</td></tr>
<tr><td>A CIFS-only filer</td><td>You are done.</td></tr>
<tr><td>A multiprotocol filer</td><td>Go to Step 5.</td></tr>
<tr><td>5</td><td colspan="2">Update the system /etc/exports file with the new volume name.</td></tr>
<tr><td>6</td><td colspan="2">Run <tt>exportfs</tt>.</td></tr>
<tr><td>7</td><td colspan="2">Add the appropriate mount point information to the /etc/fstab or /etc/vfstab file on clients that mount volumes from the filer.</td></tr>
</table>

SnapLock storage solution

What is SnapLock?

SnapLock is an advanced storage solution that provides an alternative to traditional optical WORM (write-once-read-many) storage systems for non-rewritable data. SnapLock is licensed, open-protocol functionality that works with application software to administer nonrewritable storage of reference data.

How SnapLock works

WORM data resides on SnapLock storage volumes that are administered much like regular storage volumes. SnapLock volumes operate in WORM mode and support standard file system semantics. Data on a SnapLock volume can be created and committed to WORM state by transitioning the data from a writable state to a read-only state.

Marking a currently writable file as read-only or marking a directory as read-only on a SnapLock volume commits the data as WORM. This commit process prevents the file from being altered or deleted by applications, users, or administrators.

Data that is committed to WORM state on a SnapLock volume is immutable and cannot be deleted. The only exceptions are empty directories and files that are not committed to a WORM state. Additionally, once directories are created, they cannot be renamed.

Licensing SnapLock functionality

SnapLock volumes can be licensed as SnapLock Compliance or SnapLock Enterprise. These two licenses are mutually exclusive and cannot be enabled at the same time.

- SnapLock Compliance

  A SnapLock Compliance volume is recommended for strictly regulated environments. This license enables basic functionality and restricts administrative access to files.

- SnapLock Enterprise

  A SnapLock Enterprise volume is recommended for less regulated environments. This license enables general functionality, and allows you to store and administer secure data.

You can add or remove the license for SnapLock functionality by using the `license` command. To add the license, enter

```
sys1> license add <code>
```

The following message is displayed:

```
A snaplock site license has been installed. Snaplock enabled.
```

To remove the license, enter

```
sys1> license delete snaplock
```

The following message is displayed:

```
Unlicensed snaplock. Snaplock disabled.
```

Creating a SnapLock volume

A SnapLock volume is created in the same way a standard filer volume is created. To create a SnapLock volume, you must have at least two new or data-free disks: a SnapLock volume and a parity disk to store data reconstruction information.

To specify a volume as SnapLock, use the `-L` parameter with the `vol create` command. For example:

`vol create` ***`vol-name`*** **`-L [-r`** ***`raid-group-size`*****`] [-l`** ***`language_code`*****`]`**
`disk-list`

vol-name is the name of the new volume (without the /vol/ prefix).

-L designates the volume as SnapLock. The -L option does not take an argument. Confirmation is required when you are creating a SnapLock volume.

raid-group-size is the RAID group size you specify. The default RAID group size is eight disks.

language_code specifies a language other than that of the root volume. The default is the language of the root volume.

disk-list is one of the following:

- *ndisks*[*@disk-size*]

  ndisks is the number of disks to use. The number must be at least 2.

  disk-size is the disk size to use, in gigabytes. You must have at least *ndisks* available disks of the size you specify.
- `-d` *disk_name1 disk_name2... disk_nameN*

  disk_name1, *disk_name2*, and *disk_nameN* are disk IDs of one or more available disks. Use a space to separate multiple disks.

Volume status

You can use the `vol status` command to verify that the volume exists. The `vol status` command output displays the attribute of the SnapLock volume in the Options column. For example:

```
sys1> vol status

Volume State           Status         Options

vol0 online            normal         root

vol_name online        normal         snaplock
```

Autosupport

The filer sends Autosupport messages to Network Appliance Technical Support. These messages include event and log-level descriptions. SnapLock volume state and options are included in Autosupport output.

Replicating SnapLock volumes

You can replicate SnapLock volumes to another filer using the SnapMirror feature of Data ONTAP. If an original volume becomes disabled, SnapMirror ensures quick restoration of data. See the SnapMirror chapter in the *Data Protection, Online Backup and Recovery Guide* for more information.

Understanding qtree creation

When to create qtrees

You create qtrees when you want to group files without creating a volume. You can group files by any combination of the following criteria:

- Security style
- Oplocks setting
- Quota limit
- Backup unit

Using qtrees for projects

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, and yet another project can use Windows as well as UNIX files.
- If the project uses Windows, set CIFS oplocks (opportunistic locks) as appropriate to the project, without affecting other projects.

  For example, if one project uses a database that requires no CIFS oplocks, you can set CIFS oplocks to Off on that project qtree. If another project uses CIFS oplocks, it can be in another qtree that has oplocks set to On.
- Use quotas to limit the disk space and number of files available to a project qtree so that the project does not use up resources that other projects and users need. For instructions about managing disk space by using quotas, see Chapter 8, “Disk Space Management Using Quotas,” on page 333.
- Back up and restore all the project files as a unit.

Using qtrees for backups

You can back up individual qtrees to

- Add flexibility to backup schedules
- Modularize backups by backing up only one set of qtrees at a time
- Keep the size of each backup to one tape

Qtree limitations

Qtrees have the following limitations:

- There is a limit of 254 qtrees per volume.
- When you enter a `df` command on a UNIX client with a path name of a qtree, the command shows the smaller of the client's file system limit or the filer disk space. This makes the qtree look fuller than it is. For additional information, see "Ascertaining free disk space" on page 231.

Detailed information

Creating a qtree involves the activities described in the following topics:

- "Creating qtrees" on page 138
- "Understanding qtree and volume security styles" on page 139
- "Understanding file access and security style" on page 141

If you do not want to accept the default security style of a volume or a qtree, you can change it, as described in "Changing the security style of a volume or a qtree" on page 143.

If you do not want to accept the default CIFS oplocks setting of a volume or a qtree, you can change it, as described in "Changing the CIFS oplocks setting" on page 145.

Creating qtrees

Creating a qtree

To create a qtree, complete the following step.

Step	Action
1	Enter the following command: **`qtree create path`** *path* is the path name of the qtree. ◆ If you want to create the qtree in a volume other than the root volume, include the volume in the name. ◆ If *path* does not begin with a slash (/), the qtree is created in the root volume.

Examples: The following command creates the news qtree in the users volume:

```
qtree create /vol/users/news
```

The following command creates the news qtree in the root volume:

```
qtree create news
```

Understanding qtree and volume security styles

About security styles

Every qtree and every volume has a security style setting. This setting determines whether files in a qtree can use Windows NT or UNIX security.

Security styles

Three security styles apply to qtrees and volumes. They are described in the following table.

Security style	Description	Effect of changing to the style
NTFS	Exactly like Windows NT NTFS: Files and directories have Windows NT file-level permission settings. **Note** To use NTFS security, the filer must be licensed for CIFS.	If the change is from a mixed qtree, Windows NT permissions determine file access for a file that had Windows NT permissions. Otherwise, UNIX-style permission bits determine file access for files created before the change. **Note** If the change is from a CIFS filer to a multiprotocol filer, as described in “Parameters to accept or change after volume creation” on page 115, and the /etc directory is a qtree, its security style changes to NTFS.
UNIX	Exactly like UNIX: Files and directories have UNIX permissions.	The filer disregards any Windows NT permissions established previously.

Security style	Description	Effect of changing to the style
Mixed	Both NTFS and UNIX security are allowed: A file or directory can have either Windows NT permissions or UNIX permissions. The default security style of a file is the style most recently used to set permissions on that file.	If NTFS permissions on a file are changed, the filer recomputes UNIX permissions on that file. If UNIX permissions or ownership on a file are changed, the filer deletes any NTFS permissions on that file.

Understanding file access and security style

Types of file access

The following types of file access are possible in Data ONTAP:

- Native file access
- CIFS access to UNIX files
- NFS access to Windows files

About native file access

CIFS accesses to Windows files obey Windows security rules, and NFS accesses to UNIX files obey UNIX security rules.

About CIFS access to UNIX files

The following principles apply to accessing UNIX files from CIFS:

- The filer maps each CIFS user to a UNIX UID and GID.
- File accesses use UNIX security or PC security.
  - 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.
- Windows users can change UNIX permissions from Windows using SecureShare® Access, as described in the CIFS chapter of the *File Access Management Guide*.

About NFS access to Windows files

The following principles apply to accessing Windows files from NFS:

- The filer maps each UNIX user to a Windows NT user.
- The filer uses Windows permissions from the mapped Windows NT user in comparing the permissions with the Windows ACL.

- NFS users cannot change permissions in an NTFS qtree.

Note

When you create an NTFS qtree or change a qtree to NTFS, every Windows user is given full access to the qtree, by default. 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.

Changing the security style of a volume or a qtree

When to change the security style of a volume or qtree

There are many circumstances in which you might want to change volume or qtree security style. Two examples are as follows:

- You might want to change the security style of a volume or a qtree after creating it to match the needs of the users of the volume or the qtree.
- You might 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 from an NTFS qtree to a mixed qtree.

Effects of changing the security style on quotas

Changing the security style of a volume or a qtree requires quota reinitialization if quotas are in effect in the qtree or the volume. For information about how changing the security style affects quota calculation, see "Effects of qtree changes on quotas" on page 370.

Changing the security style of a volume or a qtree

To change the security style of a volume or a qtree, complete the following step.

Step	Action
1	Enter the following command: **`qtree security [`*`path`*` {unix | ntfs | mixed}]`** *path* is the path name of the volume or qtree. Use `unix` for a UNIX qtree. Use `ntfs` for an NTFS qtree. Use `mixed` for a qtree with both UNIX and NTFS files.

Step	Action
2	If you have user quotas in effect on the qtree whose security style you just changed, turn quotas off and then on again using the `quotas off` *volumename* and `quotas on` *volumename* commands on the volume containing this qtree. Reinitializing the quotas in this manner allows Data ONTAP to recalculate the quota usage for users who own files with ACL or UNIX security on this qtree.

Caution

There are two changes to the security style of a qtree that you cannot perform while CIFS is running and users are connected to shares on that qtree. You cannot change UNIX security style to mixed or NTFS. You also cannot change NTFS or mixed security style to UNIX.

Example with a qtree: To change the security style of /vol/users/docs to be the same as that of Windows NT, use the following command:

```
qtree security /vol/users/docs ntfs
```

Example with a volume: To change the security style of the root directory of the users volume to mixed, so that, outside a qtree in the volume, one file can have NTFS security and another file can have UNIX security, use the following command:

```
qtree security /vol/users/ mixed
```

Changing the CIFS oplocks setting

What CIFS oplocks do

CIFS 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 turn CIFS oplocks off

CIFS oplocks on the filer are on by default.

You might turn CIFS oplocks off on a volume or a qtree under either of the following circumstances:

- You are using a database application whose documentation recommends that CIFS oplocks be turned off.
- You are handling critical data and cannot afford even the slightest data loss.

Otherwise, you can leave CIFS oplocks on.

Effect of the cifs.oplocks.enable option

The `cifs.oplocks.enable` option enables and disables CIFS oplocks for the entire filer. The setting of the `cifs.oplocks.enable` option must be consistent with the volume or qtree oplocks settings.

Setting the `cifs.oplocks.enable` option has the following effects:

- If you set the `cifs.oplocks.enable` option to Off, all CIFS oplocks on all volumes and qtrees on the filer are turned off.
- If you set the `cifs.oplocks.enable` option back to On, the setting for each qtree and volume takes effect and CIFS oplocks are turned on for those volumes and qtrees in which CIFS oplocks are enabled.

Changing the CIFS oplocks setting

To change the CIFS oplocks setting of a volume or a qtree, complete the following steps.

Caution

If you disable the CIFS oplocks feature on a volume or a qtree, any existing CIFS oplocks in the qtree will be broken.

<table>
<tr><th>Step</th><th colspan="2">Action</th></tr>
<tr><td rowspan="3">1</td><td>If you want...</td><td>Then...</td></tr>
<tr><td>CIFS oplocks to be enabled for the entire filer</td><td>Enter the following command:
<code>options cifs.oplocks.enable on</code></td></tr>
<tr><td>CIFS oplocks to be disabled for the entire filer</td><td>Enter the following command:
<code>options cifs.oplocks.enable off</code></td></tr>
<tr><td>2</td><td colspan="2">Enter the following command to enable or disable CIFS oplocks on a specified volume or qtree:
<code>qtree oplocks path [enable | disable]</code>
path is the path name of the volume or the qtree.
<code>enable</code> enables CIFS oplocks.
<code>disable</code> disables CIFS oplocks.</td></tr>
</table>

Displaying qtree status

Determining the status of qtrees

To find the security style, oplocks attribute, and SnapMirror status for all volumes and qtrees on the filer or for a specified volume, complete the following step.

Step	Action
1	Enter the following command: **`qtree status [-i] [path]`** The `-i` option includes the qtree ID number in the display.

Example 1:

```
toaster> qtree status
  Volume       Tree        Style     Oplocks  Status
  --------     --------    -----     -------- ---------
  vol0                     unix      enabled  normal
  vol0         marketing   ntfs      enabled  normal
  vol1                     unix      enabled  normal
  vol1         engr        ntfs      disabled normal
  vol1         backup      unix      enabled  snapmirrored
```

Example 2:

```
toaster> qtree status vol1
Volume       Tree        Style     Oplocks  Status
--------     --------    -----     -------- ---------
vol1                     unix      enabled  normal
vol1         engr        ntfs      disabled normal
vol1         backup      unix      enabled  snapmirrored
```

Example 3:

```
toaster> qtree status -i vol1
Volume    Tree     Style    Oplocks    Status        ID
------    ----     -----    --------   ------------  ----
vol1               unix     enabled    normal        0
vol1      engr     ntfs     disabled   normal        1
vol1      backup   unix     enabled    snapmirrored  2
```

Making a directory into a qtree or a qtree into a directory

Making a rooted directory into a qtree

A rooted directory is a directory at the root of a volume. If you have a rooted directory that you want to make into a qtree, you must do it through your client application, Windows or UNIX. Whichever client application you use, you perform this task by following a four-step process:

- Rename the directory.
- Create a qtree with the original directory name.
- Use the client application to move the contents of the directory into the new qtree.
- Delete the now-empty directory.

Note

You cannot delete a directory if it is associated with an existing CIFS share.

Following are procedures showing how to complete this process on Windows clients and on UNIX clients.

Note

These procedures are not supported in the Windows command-line interface or DOS prompt.

Making a rooted directory into a qtree through a Windows client

To make a rooted directory into a qtree by using a Windows client, complete the following steps.

Step	Action
1	Open Windows Explorer.
2	Click on the folder representation of the directory you want to change.
3	From the File menu, select Rename to give this directory a different name.
4	On the filer, use the `qtree create` command to create a new qtree with the original name.

Step	Action
5	In Windows Explorer, open the renamed folder and select the files inside it.
6	Drag these files into the folder representation of the new qtree. **Note** The more subfolders contained in a folder that you are moving across qtrees, the longer the move operation for that folder will take.
7	From the File menu, select Delete to delete the renamed, now-empty directory folder.

Making a rooted directory into a qtree through a UNIX client

To make a rooted directory into a qtree through a UNIX client, complete the following steps.

Step	Action
1	Open a UNIX window.
2	Use the `mv` command to rename the directory. **Example:** `client: mv /n/joel/vol1/dir1 /n/joel/vol1/olddir`
3	From the filer, use the `qtree create` command to create a qtree with the original name. **Example:** `filer: qtree create /n/joel/vol1/dir1`
4	From the client, use the `mv` command to move the contents of the old directory into the qtree. **Example:** `client: mv /n/joel/vol1/olddir/* /n/joel/vol1/dir1` **Note** The more subdirectories contained in a directory that you are moving across qtrees, the longer move operation for that directory will take.

Step	Action
5	Use the `rm` command to delete the old, now-empty directory. **Example:** `client: rm -rf /n/joel/vol1/olddir`

Deleting qtrees

Before deleting a qtree

Before you delete a qtree, ensure that the following is true:

- The volume that contains the qtree you want to delete is mounted (for NFS) or mapped (for CIFS).
- The qtree you are deleting is not directly mounted or has a CIFS share directly associated with it.
- The qtree permissions allow you to delete the qtree.

Deleting a qtree

To delete a qtree, complete the following steps.

Step	Action
1	Find the qtree you want to delete. **Note** The qtree appears as a normal directory at the root of the volume.
2	Delete the qtree using the method appropriate for your client. **Example:** The following command on a UNIX host deletes a qtree that contains files and subdirectories: `rm -Rf` *`directory`* **Note** On a Windows host, delete a qtree by using Windows Explorer.

Increasing the maximum number of files in a volume

About increasing the maximum number of files

The filer automatically sets the maximum number of files for a newly created volume based on the amount of disk space in the volume. The filer increases the maximum number of files when you add a disk to a volume. The number set by the filer never exceeds 33,554,432. This prevents a filer with terabytes of storage from creating a very large inode file, such as a 64-GB inode file, that might be unnecessary.

If you need more files on your filer, use the `maxfiles` command to increase the number beyond 33,554,432.

Caution

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.

Increasing the maximum number of files allowed on a volume

To increase the maximum number of files allowed on a volume, complete the following step.

Step	Action
1	Enter the following command: **`maxfiles` *`vol_name max`*** *vol_name* is the volume whose maximum number of files you are increasing. *max* is the maximum number of files.

Displaying the number of files in a volume

To see how many files are in a volume, complete the following step.

Step	Action
1	Enter the following command: **`maxfiles`** ***`vol_name`*** *vol_name* is the volume whose maximum number of files you are increasing. **Result:** A display like the following appears: `Volume home: maximum number of files is currently 120962 (2872 used)`

Setting space reservation for files

What space reservation is

When space reservation is enabled for one or more files, Data ONTAP reserves enough space in the volume so that writes to a those files do not fail because of a lack of disk space (ENOSPC error). Other operations such as snapshots or the creation of new files can occur only if there is enough available unreserved space; other operations are restricted from using reserved space.

Writes to new or existing unreserved space in a volume will fail when the total amount of free space in the volume is less than the amount needed according to the current file reserve values. Once free space in a volume goes below this value, only writes to files with reserved space are guaranteed to succeed.

File space reservation is persistent across filer reboots and across takeovers and givebacks, but does not persist through reversions to an earlier Data ONTAP software version.

Note

This manual does not discuss enabling space reservation for LUNs on NetApp appliances configured for block storage. For details on enabling space reservation for LUNs, see the *Block Access Management Guide for FCP* or the *Block Access Management Guide for iSCSI*.

File space reservation disk space requirements

In volumes that keep snapshot data, files with space reservation require disk space in addition to their file size:

- Two times their file size (space claimed by the actual file and the equivalent amount of space reserved for any necessary rewrites)
- Plus any of their blocks that have been overwritten through file edits and are currently saved as delta data in volume snapshots

Note

In volumes that do not keep snapshots, files with space reservation require no extra disk space.

Enabling space reservation for a specific file

To enable space reservation for a file, complete the following step.

Step	Action
1	Enter the following command: **`file reservation file_name [enable|disable]`** *file_name* is the file in which file space reservation is set. `enable` turns space reservation on for the file *file_name* `disable` turns space reservation off for the file *file_name*. **Example:** `file reservation myfile enable`

Turning on space reservation for a file will fail if there is not enough free space in the volume for the new reservation.

Querying space reservation for files: To find out the status of space reservation for files in a volume, complete the following step.

Step	Action
1	Enter the following command: **`file reservation file_name`** *file_name* is the file in which file reservations are set. **Example:** `file reservation myfile`

Enabling default file space reservation for volumes

You can use the `create_reserved` parameter on the `vol options` command to enable default file space reservation for specified volumes. When default file space reservation is enabled on a volume, all subsequent files created in that volume are, by default, enabled for space reservation.

Caution

Enabling default file space reservation for a volume, means that *all* subsequent files created for that volume will require the extra volume space described in "File space reservation disk space requirements" on page 154. You can disable files for space reservation only one at a time through the `file reservation` *file_name* `disable` command.

To set space reservation values for new files in a specified volume, complete the following step.

Step	Action
1	Enter the following command: **`vol options vol_name create_reserved {on | off}`** *vol_name* is the volume in which file reservations are set. `create_reserved on` enables space reservation for all new files that are created or stored in the volume *vol_name*. `create_reserved off` disables space reservation for all new files that are created or stored in the volume *vol_name*.

Querying space reservation for volumes: To find out the status of space reservation for a volume, complete the following step.

Step	Action
1	Enter the following command: **`vol options vol_name`** *vol_name* is the volume in which reservations are set. **Example:** `vol options vol2` **Result:** `root, diskroot, nosnap=off, nosnapdir=off, minra=off, no_atime_update=off, raidtype=raid4, raidsize=8, nvfail=off, snapmirrored=off, resyncsnaptime=60, create_ucode=off, convert_ucode=off, maxdirsize=10240, fs_size_fixed=off, create_reserved=on`

Enabling default file space reservation for qtrees

You can use the `qtree reservation` command to enable default file space reservation for specified qtrees. When default file space reservation is enabled for a qtree, all subsequent files created in or copied to that qtree are, by default, enabled for space reservation.

Caution

Enabling default file space reservation for a qtree means that *all* subsequent files created for that qtree will require the extra volume space described in "File space reservation disk space requirements" on page 154. You can disable files for space reservation only one at a time through the `file reservation` *filename* `disable` command.

To set space reservation values for new files in a specified qtree, complete the following step.

Step	Action
1	Enter the following command: **`qtree reservation` *`qtree_name`* `{enable | disable}`** *qtree_name* is the qtree in which file reservations are set. `enable` enables space reservation for all new files that are created or stored in the qtree *qtree_name*. `disable` disables space reservation for all new files that are created or stored in the qtree *qtree_name*.

If default file space reservation is enabled for the volume: Default file space reservation can be set for volumes as well as for qtrees.

- If default file space reservation is enabled in a volume (through the `vol options create_reserved on` command), default file space reservation is enabled for all new files created in that volume, even for files in qtrees with default file space reservation set to off.
- If default file space reservation is disabled for a volume (`vol options create_reserved off`), enabling default file space reservation for a qtree in that volume means that all new files created in that qtree will have the values turned on.

Querying space reservation for qtrees: To find out the status of space reservation for a qtree, complete the following step.

Step	Action
1	Enter the following command: **`qtree reservation` *`qtree_name`*** *vol_name* is the volume in which reservations are set.

RAID Group Management 4

About this chapter

This chapter describes how to manage RAID (Redundant Array of Independent Disks) on filer volumes. Data ONTAP uses RAID Level 4 or RAID DP (double-parity) to ensure data integrity within a group of disks even if one or two of those disks fail.

- If a filer volume is configured for RAID 4 protection, Data ONTAP can reconstruct the data from a single failed disk within a RAID group and transfer that reconstructed data to a spare disk.
- If a filer volume is configured for RAID DP protection, Data ONTAP can reconstruct the data from one or two failed disks within a RAID group and transfer that reconstructed data to one or two spare disks as necessary.

Note

The RAID principles and management operations described in this chapter do not apply to gFiler gateways, which are supported by RAID systems specific to the storage subsystems that they manage.

Topics in this chapter

This chapter discusses the following topics:

- "Understanding RAID groups" on page 160
- "Configuring RAID type and group size while creating a volume" on page 165
- "Modifying the RAID type for a volume" on page 172
- "Modifying the size of RAID groups" on page 176
- "Controlling the speed of RAID operations" on page 179
- "Scheduling or manually running disk scrubbing" on page 186
- "Checking and correcting parity" on page 190
- "Minimizing media error disruption of RAID reconstructions" on page 193
- "Viewing RAID status" on page 199

Understanding RAID groups

About RAID groups in Data ONTAP

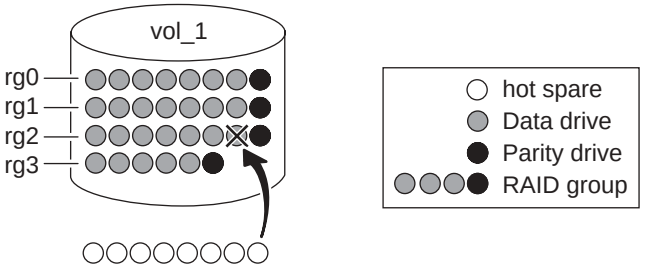
A volume consists of disks assembled into one or multiple RAID groups. A RAID group consists of one or more data disks, across which client data is striped and stored, plus one or two parity disks that ensure data recoverability even if as many as two disks within the RAID group fail. The purpose of a RAID group is to provide parity protection of data across its included disks from loss due to single or double disk failure within that group.

Types of RAID protection

Data ONTAP supports two types of RAID protection, RAID 4 or RAID DP, which you can assign on a per-volume basis.

RAID 4 protection: RAID 4 provides single-parity disk protection against single disk failure within a RAID group. The minimum number of disks in a RAID 4 group is two: at least one data disk and one parity disk. In case of single disk failure, Data ONTAP replaces the failed disk in the RAID group with a spare disk and uses the parity data to reconstruct the failed disk's data on the replacement disk. The following figure diagrams a volume configured for RAID 4 protection.

RAID-4 protection

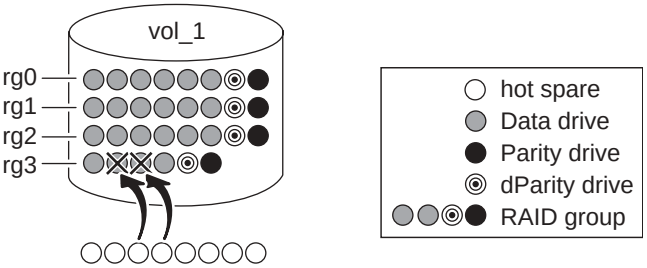


RAID DP protection: RAID DP provides double-parity disk protection against single or double disk failure within a RAID group. The minimum number of disks in a RAID DP group is three: at least one data disk, one regular parity disk, and one double-parity (or dParity) disk. (A three-disk root volume, for example, could be parity-protected even if one data and one parity disk failed).

In case of single disk failure, Data ONTAP replaces the failed disk in the RAID group with a spare disk and uses the parity data to reconstruct the data of the

failed disk on the replacement disk. In case of double disk failure, Data ONTAP replaces the failed disks in the RAID group with two spare disks and uses the double-parity data to reconstruct the data of the failed disks on the replacement disks. The following figure diagrams a volume configured for RAID DP protection.

RAID-DP protection



How Data ONTAP creates RAID groups automatically

When you create a filer volume or add disks to a filer volume, Data ONTAP creates new RAID groups as each RAID group is filled with its maximum number of disks. Within each volume, RAID groups are named rg0, rg1, rg2, and so on in order of their creation. The last RAID group formed might contain fewer disks than are specified for the volume's RAID group size. In that case, any disks added to the volume are also added to the last RAID group until the specified RAID group size is reached.

- If a volume is configured for RAID 4 protection, Data ONTAP assigns the role of parity disk to the largest disk in each RAID group.

  Note

  If an existing RAID 4 group is assigned an additional disk that is larger than the group's existing parity disk, then Data ONTAP reassigns the new disk as parity disk for that RAID group.

- If a volume is configured for RAID DP protection, Data ONTAP assigns the role of dParity disk and regular parity disk to the largest disk and the second-larger disk in the RAID group.

  Note

  If an existing RAID DP group is assigned an additional disk that is larger than the group's existing dParity disk, then Data ONTAP reassigns the new disk as the regular parity disk for that RAID group but restricts its capacity to be no greater than that of the existing dParity disk.

Hot spare disks

A hot spare disk is a filer disk that has not been assigned to a RAID group. It does not yet hold data but is ready for use. In event of disk failure within a RAID group, Data ONTAP automatically assigns hot spare disks to RAID groups to replace the failed disks. Hot spare disks do not have to be in the same disk shelf as other disks of a RAID group to be available to a RAID group.

Hot spare disk size recommendations: Network Appliance recommends keeping at least one hot spare disk for each disk size and disk type installed in your filer. This allows the filer to use a disk of the same size and type as a failed disk when reconstructing a failed disk. If a disk fails and a hot spare disk of the same size is not available, the filer uses a spare disk of the next available size up. See "Disk failure with a hot spare disk" on page 169 for more information.

Note

If no spare disks exist in a filer, Data ONTAP can continue to function in degraded mode. Data ONTAP supports degraded mode in case of single disk failure for volumes configured with RAID 4 protection and in case of single or double disk failure in volumes configured for RAID DP protection. For details see "Disk failure without a hot spare disk" on page 171.

Maximum number of RAID groups

Data ONTAP enables you to support up to 400 RAID groups per filer or cluster. When configuring your volumes, keep in mind that each volume requires at least one RAID group and that the total of all RAID groups in a filer cannot exceed 400.

RAID 4, RAID DP, and SyncMirror

RAID 4 or RAID DP can be used in combination with the Data ONTAP SyncMirror feature, which also offers protection against data loss due to disk or other hardware component failure. SyncMirror protects against data loss by maintaining two mirrored plexes of the same volume. Any data loss due to disk failure in one plex is repaired by the undamaged data in the opposite plex. The advantages and disadvantages of using RAID 4 or RAID DP, with and without the SyncMirror feature, are listed in the following table.

Factor affected by RAID type	RAID 4	RAID DP	RAID 4 with SyncMirror	RAID DP with SyncMirror
The kinds of failures that RAID and SyncMirror protect against	Single disk failure within one or multiple RAID groups	Single or double disk failure within one or multiple RAID groups	Single, double, or greater disk failure within one or multiple RAID groups across one sync-mirrored plex. (See "Advantages of RAID 4 with SyncMirror" on page 163.) Storage subsystem failures (HBA, cables, shelf) on the filer	Single, double, or greater disk failure within one or multiple RAID groups across both sync-mirrored plexes. (See "Advantages of RAID DP with SyncMirror" on page 164.) Storage subsystem failures (HBA, cables, shelf) on the filer
Required disk resources per RAID group	*n* data disks + 1 parity disk	*n* data disks + 2 parity disks	2 x (*n* data disks + 1 parity disk)	2 x (*n* data disks + 2 parity disks)
The performance cost	None	Almost none	Mirroring overhead	Mirroring overhead
Additional cost and complexity	None	None	SyncMirror license and configuration Possible cluster license and configuration	SyncMirror license and configuration Possible cluster license and configuration

Advantages of RAID 4 with SyncMirror: On SyncMirror-replicated volumes using RAID 4, any combination of multiple disk failures within single RAID groups in one plex is restorable, as long as multiple disk failures are not concurrently occurring in the same RAID groups in the opposite plex of the mirrored volume. Between two mirrored RAID groups, a concurrent failure of any three disks is tolerable.

Advantages of RAID DP with SyncMirror: On SyncMirror-replicated volumes using RAID DP, any combination of multiple disk failures within single RAID groups in one plex is restorable, as long as concurrent failures of greater than two disks are not occurring in the same RAID groups in the opposite plex of the mirrored volume. Between mirrored RAID groups, concurrent failure of any 5 disks is tolerable.

For more SyncMirror information: For more information on the Data ONTAP SyncMirror feature, see the *Data ONTAP System Administration Data Protection Guide*.

Larger versus smaller RAID groups

Configuring an optimum RAID group size for a volume requires a trade-off of factors. You must decide which factor—speed of recovery, assurance against data loss, or maximizing data storage space— is most important for the volume that you are configuring. For a listing of default and maximum RAID group sizes, see "Maximum and default RAID group sizes" on page 176.

Advantages of large RAID groups: Large RAID group configurations offer the following advantages:

- More data drives available. A volume configured into a few large RAID groups requires fewer drives reserved for parity than that same volume configured into many small RAID groups.
- Better system performance. Read/write operations are usually faster over large RAID groups than over smaller RAID groups.

Advantages of small RAID groups: Small RAID group configurations offer the following advantages:

- Shorter disk reconstruction times. In case of disk failure within a small RAID group, data reconstruction time is usually shorter than it would be within a large RAID group.
- Decreased risk of data loss due to multiple disk failures. The probability of data loss through double disk failure within a RAID 4 group or through triple disk failure within a RAID DP group is lower within a small RAID group than in a large RAID group.

Configuring RAID type and group size while creating a volume

About RAID group type

When you use the `vol create` command to create a volume, the volume's RAID type and group size can be set by default Data ONTAP values, or you can directly specify the size and type of a volume's RAID groups through the `vol create` command options.

Specifying the RAID type at volume creation

To specify the size and type of a volume's RAID groups when you create that volume, complete the following steps.

Step	Action
1	To view the filer disks that are "spare" and thus available to include in a new volume. Enter the following command: `vol status -s` **Result:** The device number, shelf number, and capacity of all spare disks on the filer are listed. The filer should contain spare disks of sufficient number and capacity to ◆ Support the volume that you want to create ◆ Provide enough spare disks to serve as replacement disks should disk failure occur in any filer volume. To fully support a volume's RAID disk failure protection, at least one spare disk is required for that volume. The size of the spare disks should be equal to or greater than the size of the volume disks that the spare disks might replace. Network Appliance recommends that spare disks include disks to match the size of any of the volume disks that might require replacement.

Step	Action
2	To specify RAID group size and RAID group type, use the `vol create` command's `-r` and `-t` options when you create the volume. Enter `vol create` *`volname`* `[-m] [-t {raid_4|raid_dp}] [-r` *`raid_group_size`*`]` *`disk_list`* *volname* is the name of the volume you want to create. `-m` specifies the optional creation of a SyncMirror-replicated volume if you want to supplement RAID protection with SyncMirror protection. A SyncMirror license is required for this feature. `-t {raid_4|raid_dp}` specifies the type of RAID protection (RAID 4 or RAID DP) that you want to provide for this volume. If no RAID type is specified, the default value (`raid_dp` on NearStore appliances, `raid_4` on all others) is applied. `-r` *raid-group-size* is the number of drives per RAID group that you want for this volume. If no RAID group size is specified, the default value for your appliance model is applied. For a listing of default and maximum RAID group sizes, see "Maximum and default RAID group sizes" on page 176. *disk_list* specifies the disks to include in the volume that you want to create. It can be expressed in the following formats: ◆ *ndisks*[*@disk-size*] *ndisks* is the number of disks to use. It must be at least 2. *disk-size* is the disk size to use, in gigabytes. You must have at least *ndisk* available disks of the size you specify. ◆ `-d` *disk_name1 disk_name2... disk_nameN* *disk_name1*, *disk_name2*, and *disk_nameN* are disk IDs of one or more available disks; use a space to separate multiple disks. **Example:** The following command creates the volume newvol with RAID DP protection. RAID group size is 16 disks. Since the volume consists of 32 disks, those disks will form 2 RAID groups, rg0 and rg1: `vol create newvol -t raid_dp -r 16 32@36`

Step	Action
3	To verify the RAID structure of the volume that you just created, enter the following command: `vol status volname -r` **Result:** The parity and data disks for each RAID group in the volume just created are listed. In volumes with RAID 4 protection, you will see “parity,” and “data” disks listed for each RAID group in a volume. In volumes with RAID DP protection you will see “parity,” “dParity,” and “data” disks listed for each RAID group
4	To verify that spare disks of sufficient number and size exist on the filer to serve as replacement disks in event of disk failure in one of the RAID groups in the volume that you just created, enter the following command: `vol status -s`

How RAID handles disk failures

About disk failures

Disks can fail. RAID enables data on a disk to be preserved even if that disk fails. This section describes how RAID handles disk failures in two situations.

Detailed information

For detailed information about how RAID handles disk failures, see the following topics:

- "Disk failure with a hot spare disk" on page 169
- "Disk failure without a hot spare disk" on page 171

Disk failure with a hot spare disk

About this section

This section describes how the filer reacts to a single or double disk failure when a hot spare disk is available.

Filer replaces disk with spare and reconstructs data

If a disk fails, the filer

- Replaces the failed disk with a hot spare disk (if RAID DP is enabled and double disk failure occurs in the RAID group, the filer replaces each failed disk with a separate spare disk). Data ONTAP first attempts to use a hot spare disk of the same size as the failed disk. If no disk of the same size is available, Data ONTAP replaces the failed disk with a spare disk of the next available size up.
- Reconstructs in the background the missing data onto the hot spare disk or disks
- Logs the activity in the /etc/messages file on the root volume
- Sends an Autosupport message

Note

If RAID DP is enabled, the above processes can be carried out even in the event of simultaneous failure on two disks in a RAID group.

During reconstruction, file service can slow down.

Caution

After the filer is finished reconstructing data, replace the failed disk or disks with new hot spare disks as soon as possible so that hot spare disks are always available in the system. For information about replacing a disk, see Chapter 5, "Disk Management," on page 201.

If a disk fails and no hot spare disk is available, contact Network Appliance Technical Support.

Network Appliance recommends keeping at least one matching hot spare disk for each disk size and disk type installed in your filer. This allows the filer to use a disk of the same size and type as a failed disk when reconstructing a failed disk. If a disk fails and a hot spare disk of the same size is not available, the filer uses a spare disk of the next available size up.

Information

In addition to disk failure and hot spare disk replacement, the /etc/messages file on the root volume logs any failure in a periodic check of the hot spare disk. For more information, see the na_sysconfig(1) and na_messages(5) man pages.

Disk failure without a hot spare disk

About this section

This section describes how the filer reacts to a disk failure when hot spare disks are not available.

Filer runs in degraded mode

If a disk fails, the filer continues to run without losing any data but performance is somewhat degraded. (If RAID DP is enabled, the filer can continue running with one or two disks failed in a RAID group.)

Caution

Replace the failed disks as soon as possible, because additional disk failure might 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 at the end of a period that you set with the `raid.timeout` option of 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.

Filer sends messages about failures

Check the /etc/messages file on the root volume once a day for important messages. You can automate checking of this file from a remote host with a script that periodically searches the file and then alerts you.

Alternatively, you can monitor Autosupport messages. Data ONTAP sends Autosupport messages when a disk fails.

Filer reconstructs data after disk is replaced

After you replace a disk, the filer detects the new disk immediately and uses it. The filer starts file service and reconstructs the missing data in the background to minimize service interruption.

Modifying the RAID type for a volume

Changing the RAID group type

You can change the type of RAID protection configured for a volume. When you change a volume's RAID type, Data ONTAP reconfigures all the existing RAID groups to the new type and applies the new type to all subsequently created RAID groups in that volume.

Changing from RAID 4 to RAID DP protection

When you change a volume's RAID protection from RAID 4 to RAID DP, you first need to ensure that hot spare disks of sufficient number and size exist on that volume's filer, because during the conversion, Data ONTAP adds an additional disk to each existing RAID group from the filer's hot spare disk pool and assigns the new disk the dParity disk function for the RAID DP group. In addition, the volume's `raidsize` option, which controls the size of new RAID groups that might be created in the volume, is modified to RAID DP defaults on this appliance.

To change an existing volume's RAID protection from RAID 4 to RAID DP, complete the following steps.

Step	Action
1	Determine the number of RAID groups and the size of their parity disks in the volume in question by issuing the following command. **`vol status` *`volname`* `-r`**
2	Issue the following command to make sure that a hot spare disk exists on the filer for each RAID group listed for the volume in question, and make sure that these hot spare disks match the size and checksum type of the existing parity disks in those RAID groups. **`vol status -s`** If necessary, add hot spare disks of the appropriate number of appropriate number and size and checksum type to the filer. See "Adding new disks to the filer" on page 223.

Step	Action
3	Enter the following command: **`vol options vol-name raidtype raid_dp`** *vol-name* is the volume whose RAID type you are changing. **Example:** The following command changes the RAID type of the volume thisvol to RAID DP: `vol options thisvol raidtype raid_dp`

Associated RAID group size changes: When you change the RAID protection of an existing volume from RAID 4 to RAID DP, the following associated RAID group size changes take place:

- A second parity disk (dParity) is automatically added to each existing RAID group from the hot spare disk pool, thus increasing the size of each existing RAID group by one.

  Note

  If hot spare disks of insufficient number or size to support the RAID type conversion are available on the filer, Data ONTAP will issue a warning before executing the `vol options` *volname* `raidtype raid_dp` command. If you continue the operation, RAID DP protection will be implemented on the volume in question, but some of its RAID groups for which no second parity disk was available will remain degraded. The protection will be no improvement over RAID 4, and no hot spare disks will be available in case of disk failure since all were reassigned as dParity disks.

- The volume's `raidsize` option, which sets the size for any new RAID groups created in this volume, is automatically reset to one of the following RAID DP defaults: on all non-NearStore appliances, 16; on an R100 appliance, 12; on an R150 platform, 14.

  Note

  After the `vol options` *volname raidtype raid_dp* operation is complete, you can manually change the `raidsize` option through the `vol options` *volname* `raidsize` command. See "Changing the size of the most recent and future RAID groups" on page 177.

Changing from RAID DP to RAID 4 protection

To change the an existing volume's RAID protection from RAID DP to RAID 4, complete the following step.

Step	Action
1	Enter the following command: **`vol options vol-name raidtype raid4`** *vol-name* is the volume whose RAID type you are changing. **Example:** The following command changes the RAID type of the volume thatvol to RAID 4: `vol options thatvol raidtype raid4`

Associated RAID group size changes: When you change the RAID protection of an existing volume from RAID DP to RAID 4, Data ONTAP automatically carries out the following associated RAID group size changes:

- In each of the volume's existing RAID groups, the RAID DP second parity disk (dParity) is removed and placed in the hot spare disk pool, thus reducing each RAID group's size by one.
- On NearStore appliances, Data ONTAP changes the volume's `raidsize` option, which determines the size of any new RAID group created in this volume, to the RAID 4 default size on the NearStore R100 (8) or R150 (6) platforms.
- On non-NearStore appliances, Data ONTAP usually changes the setting for the volume's `raidsize` option, which determines the size of any new RAID group created in this volume, to the size of the largest RAID group in the volume.

  Exceptions are the following cases: if the volume's largest RAID group is larger than the maximum RAID 4 group size on non-NearStore platforms (14), then the volume's `raidsize` option will be set to 14; if the volume's largest RAID group is smaller than the default RAID 4 group size on non-NearStore platforms (8), then the volume's `raidsize` option will be set to 8.

Note

After the `vol options` *volname raidtype raid_dp* operation is complete, you can change the raidsize setting through the `vol options` *volname* `raidsize` command. See "Changing the size of the most recent and future RAID groups" on page 177.

Verifying the RAID type

To verify the RAID type of a volume, enter the `vol options` *`vol_name`* or `vol status` *`vol_name`* command.

Modifying the size of RAID groups

Maximum and default RAID group sizes

You can change the size of RAID groups that will be created on a volume.

Maximum and default RAID group sizes vary according to the NetApp platform and type of RAID group protection provided. The default RAID group sizes are the sizes that NetApp generally recommends.

Maximum and default RAID DP group sizes: The following table lists the maximum RAID DP group sizes supported on Network Appliance platforms.

Platform	Maximum group size	Default group size	Minimum group size
R100	12	12	3
R150	12	12	3
R200	14	14	3
FAS250	14	7	3
All other NetApp platforms	28	16	3

Note

If as a result of a software upgrade from an older version of Data ONTAP, volumes that contain RAID 4 groups larger than the maximum 28 disks exist, NetApp recommends that you convert the volumes in question to RAID DP as soon as possible.

Maximum and default RAID 4 group sizes: The following table lists the maximum RAID 4 group sizes supported on NetApp platforms.

Platform	Maximum group size	Default group size	Minimum group size
FAS250	14	7	2
R100	8	8	2

Platform	Maximum group size	Default group size	Minimum group size
R150	6	6	2
R200	7	7	2
All other NetApp platforms	14	8	2

Changing the size of the most recent and future RAID groups

The volume option named `raidsize` is the maximum RAID group size that can be reached by adding disks to a volume.

- You can modify the `raidsize` setting upward to increase the allowable size of the most recently created RAID group. If you cannot use the raidsize setting to decrease the size of the most recently created RAID group.
- The new `raidsize` setting also applies to subsequently created RAID groups in a volume. Either increased or decreased `raidsize` settings will apply to future RAID groups.
- With the exception of the most recently created RAID group, existing RAID groups remain the whatever size they were before the `raidsize` setting was changed.

To change the `raidsize` setting for an existing volume, complete the following step.

Step	Action
1	Enter the following command: **`vol options vol-name raidsize size`** *vol-name* is the volume whose `raidsize` setting you are changing. *size* is the number of disks you want in the most recently created and all future RAID groups in this volume. **Example:** The following command changes the `raidsize` setting of the volume yervol to 16 disks: `vol options yervol raidsize 16`

See “Adding new disks to a RAID group” on page 226 and “Forcibly adding disks to volumes” on page 226 for information about adding disks to RAID groups.

Verifying the raidsize setting

To verify the size of raidsize setting in a volume, enter the `vol options` *`vol_name`* command.

Modifying the size of existing RAID groups

You can use `-g` *raidgroup* option in the `vol add` command to increase the size of an existing RAID group. See “Adding new disks to a RAID group” on page 226 and “Forcibly adding disks to volumes” on page 226 for information about adding disks to existing RAID groups.

Controlling the speed of RAID operations

RAID operations you can control

You can control the speed of the following RAID operations:

- RAID data reconstruction
- Disk scrubbing
- Plex resynchronization
- Synchronous mirror verification

Effects of varying the speed on filer performance

The speed that you select can affect the overall performance of the filer. You can vary the speed of an operation from a speed that minimally slows filer network operations to a speed that can severely slow filer network operations.

Before changing the speed of a RAID operation, you might want to determine the CPU load to see what effect the change might have on network performance.

Determining CPU load

To determine the CPU load on the filer when RAID operations are in progress, complete the following steps.

Step	Action	
1	Enter the following command: `sysstat` **Result:** System activity statistics appear every 15 seconds.	
2	**If...**	**Then...**
	The number under CPU is lower than 70%	Consider increasing the RAID operation speed to hasten reconstruction.
	The number under CPU is 70% or higher	Consider decreasing the RAID operation speed to free CPU resources and possibly improve network performance.

Detailed information

The following sections discuss how to control the speed of RAID operations:

- "Controlling the speed of RAID data reconstruction" on page 181
- "Controlling the speed of disk scrubbing" on page 182
- "Controlling the speed of plex resynchronization" on page 184
- "Controlling the speed of mirror verification" on page 185

Controlling the speed of RAID data reconstruction

About RAID data reconstruction

If a disk fails, the data on it is reconstructed on a hot spare disk if one is available. Because RAID data reconstruction consumes CPU resources, increasing the speed of data reconstruction sometimes slows filer network and disk operations.

Modifying RAID data reconstruction speed

To modify the speed of data reconstruction, complete the following step.

Step	Action
1	Enter the following command: **`options raid.reconstruct.perf_impact`** ***`impact`*** *impact* can be high, medium, or low. High means that the filer uses most of the CPU time available for RAID data reconstruction; this setting can heavily affect filer performance. Low means that the filer uses very little CPU time; this setting lightly affects filer performance. The default speed is medium.

RAID operations affecting RAID data reconstruction speed

When RAID data reconstruction and plex resynchronization are running at the same time, Data ONTAP limits the combined resource utilization to the greatest impact set by either operation. For example, if `raid.resync.perf_impact` is set to medium and `raid.reconstruct.perf_impact` is set to low, the resource utilization of both operations has a medium impact.

Controlling the speed of disk scrubbing

About disk scrubbing

Disk scrubbing means checking the disk blocks of all disks on the filer for media errors and consistency. If Data ONTAP finds media errors or inconsistencies, it fixes them by reconstructing the data from other disks and rewriting the data. Disk scrubbing reduces the chance of a multiple disk failure that is so extensive that RAID reconstruction is impossible.

Data ONTAP enables block checksums to ensure data integrity. If checksums are incorrect, Data ONTAP generates an error message similar to the following:

```
Scrub found checksum error on /vol/vol0/plex0/rg0/4.0 block 436964
```

If RAID 4 is enabled, Data ONTAP scrubs a RAID group only when all the group's disks are operational.

If RAID DP is enabled, Data ONTAP can carry out a scrub even if one disk in the RAID group has failed.

Although disk scrubbing slows the filer somewhat, network clients might not notice the change in filer performance because disk scrubbing starts automatically at 1:00 a.m. on Sunday, when most systems are lightly loaded, and stops after six hours.

Modifying disk scrubbing speed

To modify the speed of disk scrubbing, complete the following step.

Step	Action
1	Enter the following command: **`options raid.scrub.perf_impact`** ***`impact`*** *impact* can be high, medium, or low. High means that the filer uses most of the CPU time available for disk scrubbing; this setting can heavily affect filer performance. Low means that the filer uses very little CPU time; this setting lightly affects filer performance. The default speed is low.

RAID operations affecting disk scrubbing speed

When disk scrubbing and mirror verification are running at the same time, Data ONTAP limits the combined resource utilization to the greatest impact set by either operation. For example, if `raid.verify.perf_impact` is set to medium and `raid.scrub.perf_impact` is set to low, the resource utilization by both operations has a medium impact.

Controlling the speed of plex resynchronization

What plex resynchronization is

Plex resynchronization refers to the process of synchronizing the data of the two plexes of a mirrored volume. When plexes are synchronized, the data on each plex is identical. When plexes are unsynchronized, one plex contains data that is more up to date than that of the other plex. Plex resynchronization updates the out-of-date plex until both plexes are identical.

When plex resynchronization occurs

Data ONTAP resynchronizes the two plexes of a mirrored volume if one of the following occurs:

- One of the plexes was taken offline and then brought online later
- You add a plex to an unmirrored volume

Modifying plex resynchronization speed

To modify the speed of plex resynchronization, complete the following step.

Step	Action
1	Enter the following command: `options raid.resync.perf_impact` ***`impact`*** *impact* can be high, medium, or low. High means that the filer uses most of the CPU time available for plex resynchronization; this setting can heavily affect filer performance. Low means that the filer uses very little CPU time; this setting lightly affects filer performance. The default speed is medium.

RAID operations affecting plex resynchronization speed

When plex resynchronization and RAID data reconstruction are running at the same time, Data ONTAP limits the combined resource utilization to the greatest impact set by either operation. For example, if `raid.resync.perf_impact` is set to medium and `raid.reconstruct.perf_impact` is set to low, the resource utilization by both operations has a medium impact.

Controlling the speed of mirror verification

What mirror verification is

You use mirror verification to ensure that the two plexes of a synchronous mirrored volume are identical. See the synchronous mirror volume management chapter in the *Data Protection Guide* for more information.

Modifying mirror verification speed

To modify the speed of mirror verification, complete the following step.

Step	Action
1	Enter the following command: **`options raid.verify.perf_impact`** ***`impact`*** *impact* can be high, medium, or low. High means that the filer uses most of the CPU time available for mirror verification; this setting can heavily affect filer performance. Low means that the filer uses very little CPU time; this setting lightly affects filer performance. The default speed is low.

RAID operations affecting mirror verification speed

When mirror verification and disk scrubbing are running at the same time, Data ONTAP limits the combined resource utilization to the greatest impact set by either operation. For example, if `raid.verify.perf_impact` is set to medium and `raid.scrub.perf_impact` is set to low, the resource utilization of both operations has a medium impact.

Scheduling or manually running disk scrubbing

About disk scrub scheduling

By default, automatic disk scrubbing is enabled for once a week and begins at 1:00 a.m. on Sunday. However, you can modify this schedule to suit your needs.

- You can reschedule automatic disk scrubbing to take place on other days, at other times, or at multiple times during the week.
- You might want to disable automatic disk scrubbing if disk scrubbing encounters a recurring problem.
- You can specify the duration of a disk scrubbing operation.
- You can start or stop a disk scrubbing operation manually.

Rescheduling disk scrubbing

If you want to reschedule the default weekly disk scrubbing time of 1:00 a.m. on Sunday, you can specify the day, time, and duration of one or more alternative disk scrubbings for the week. To schedule weekly disk scrubbings, complete the following steps.

Step	Action
1	Enter the following command: **`options raid.scrub.schedule`** ***`duration`*`{h|m}@`*`weekday@start_time`**`[,`***`duration`***`{h|m}@`***`weekday@start_time`**`] ...`** *duration* {h|m} is the amount of time, in hours (h) or minutes (m) that you want to allot for this operation. **Note** If no duration is specified for a given scrub, the value specified in the `raid.scrub.duration` option is used. For details, see "Setting the duration of automatic disk scrubbing" on page 188. *weekday* is the day of the week (`sun`, `mon`, `tue`, `wed`, `thu`, `fri`, `sat`) when you want the operation to start. *start_time* is the hour of the day you want the scrub to start. Acceptable values are `0-23`, where `0` is midnight and `23` is 11 p.m. **Example:** The following command schedules two weekly RAID scrubs. The first scrub is for four hours every Tuesday at 2 a.m. The second scrub is for eight hours every Saturday at 10 p.m. **`options raid.scrub.schedule 240m@tue@2,8h@sat@22`**
2	Verify your modification with the following command: **`options raid.scrub.schedule`** The duration, weekday, and start times for all your scheduled disk scrubs appear. **Note** If you want to restore the default automatic scrub schedule of Sunday at 1:00 a.m., reenter the command with an empty value string as follows: **`options raid.scrub.schedule " "`**.

Toggling automatic disk scrubbing

To enable and disable automatic disk scrubbing for the filer, complete the following step.

Step	Action
1	Enter the following command: **`options raid.scrub.enable off | on`** Use `on` to enable automatic disk scrubbing. Use `off` to disable automatic disk scrubbing.

Setting the duration of automatic disk scrubbing

You can set the duration of automatic disk scrubbing. The default is to perform automatic disk scrubbing for six hours (360 minutes). If scrubbing does not finish in six hours, Data ONTAP records where it stops. The next time disk scrubbing automatically starts, scrubbing starts from the point where it stopped.

To set the duration of automatic disk scrubbing, complete the following step.

Step	Action
1	Enter the following command: **`options raid.scrub.duration`** ***`duration`*** *duration* is the length of time, in minutes, that automatic disk scrubbing runs. **Note** If you set *duration* to -1, all automatically started disk scrubs run to completion.

Note

If an automatic disk scrubbing is scheduled through the `options raid.scrub.schedule` command, the duration specified for the `raid.scrub.duration` option applies only if no duration was specified for disk scrubbing in the `options raid.scrub.schedule` command.

Starting and stopping disk scrubbing manually

You can use your UNIX or CIFS host to start a disk scrubbing operation at any time. For example, you can start disk scrubbing by putting `disk scrub start` into a remote shell command in a UNIX cron script.

To start and stop disk scrubbing manually, complete the following step.

Step	Action
1	Enter the following command: **`disk scrub start|stop`** Use `start` to initiate disk scrubbing on all RAID groups regardless of the setting of the `raid.scrub.enable` option. Use `stop` to stop a disk scrubbing operation.

Checking and correcting parity

About checking RAID group parity

RAID group parity is checked by Data ONTAP by default. Data ONTAP also allows you to check RAID group parity manually.

You can check that parity is correct on RAID groups at the RAID group level, plex level, or volume level. Parity checking compares the data disks in a RAID group to the parity disk in a RAID group. If during the parity check Data ONTAP determines that parity is incorrect, Data ONTAP corrects the parity disk contents.

At the RAID group level, you can check only RAID groups that are in an active parity state. If the RAID group is in a degraded, reconstructing, or repairing state, Data ONTAP reports errors.

If you are checking a volume that has some RAID groups in an active parity state and some not in an active parity state, Data ONTAP checks and corrects the RAID groups in an active parity state and reports errors for the RAID groups not in an active parity state.

Checking parity on all volumes

To check parity on all volumes, complete the following step.

Step	Action
1	Enter the following command: `vol scrub start`

Checking parity on a specific volume, plex, or RAID group

To check parity on the RAID groups of a specific volume, plex, or RAID group, complete the following step.

Step	Action
1	Enter the following command: `vol scrub start` *`name`* *name* is the name of the volume, plex, or RAID group.

Examples: In this example, use the following command to check parity on all the RAID groups in the vol2 volume:

```
vol scrub start vol2
```

In this example, use the following command to check parity on all the RAID groups of plex1 of the vol2 volume:

```
vol scrub start vol2/plex1
```

In this example, use the following command to check parity on RAID group 0 of plex1 of the vol2 volume:

```
vol scrub start vol2/plex1/rg0
```

Stopping parity checking

You might need to stop Data ONTAP from checking parity. To stop parity checking, complete the following step.

Step	Action
1	Enter the following command: **`vol scrub stop`** ***`name`*** If *name* is not specified, Data ONTAP stops all parity checking.

Suspending and resuming parity checking

Rather than stopping Data ONTAP from checking and correcting parity, you can suspend checking for a while and resume checking later.

To suspend parity checking, complete the following step.

Step	Action
1	Enter the following command: **`vol scrub suspend`** ***`name`*** If *name* is not specified, Data ONTAP suspends all parity checking.

To resume parity checking, complete the following step.

Step	Action
1	Enter the following command: **`vol scrub resume`** ***`name`*** If *name* is not specified, Data ONTAP resumes all suspended parity checking.

Viewing parity checking status

The parity checking status tells you what percentage of the parity checking has been completed. Parity checking status also displays whether parity checking of a volume, plex, or RAID group is suspended.

To view the status of parity checking, complete the following step.

Step	Action
1	Enter the following command: **`vol scrub status`** ***`name`*** If *name* is not specified, Data ONTAP shows the parity checking status of all RAID groups.

Minimizing media error disruption of RAID reconstructions

About media error disruption prevention

Excessive disk media errors encountered during RAID reconstruction for a single disk failure might cause a filer panic or double disk failure. The following features minimize the risk of filer disruption due to media errors. The features include

- Improved handling of media errors by a WAFL repair mechanism. See “Handling of media errors during RAID reconstruction” on page 194.
- Default continuous media error scrubbing on filer disks. See “Continuous media scrub” on page 195.
- Continuous monitoring of disk media errors and automatic failing and replacement of disks that exceed system-defined media error thresholds. See “Disk media error failure thresholds” on page 198.

Handling of media errors during RAID reconstruction

About media error handling during RAID reconstruction

By default, Data ONTAP, if it encounters media errors during a RAID reconstruction, automatically invokes an advanced mode process (wafliron) that compensates for the media errors and enables Data ONTAP to bypass the errors.

If this process is successful, RAID reconstruction continues, and the volume in which the error was detected remains online.

If you configure Data ONTAP so that it does not invoke this process or if this process fails, Data ONTAP attempts to place the affected volume in restricted mode. If restricted mode fails, the filer panics, and after a reboot, Data ONTAP brings up the affected volume in restricted mode. In this mode, you can manually invoke the wafliron process in advanced mode or schedule downtime for your filer for reconciling the error by running the `WAFL_check` command from the Boot menu.

Purpose of the raid.reconstruction.wafliron.enable option

The `raid.reconstruction.wafliron.enable` option determines whether Data ONTAP automatically starts the wafliron process after detecting medium errors during RAID reconstruction. By default, the option is set to On.

Recommendation: Network Appliance recommends that you leave the `raid.reconstruction.wafliron.enable` option at its default option to increase data availability.

Enabling and disabling the automatic wafliron process

To enable or disable the `raid.reconstruction.wafliron.enable` option, complete the following step.

Step	Action
1	Enter the following command: **`options raid.reconstruction.wafliron.enable on | off`**

Continuous media scrub

About continuous media scrubbing

By default, Data ONTAP runs continuous background media scrubbing for media errors on filer disks. The purpose of the continuous media scrub is to detect and scrub media errors in order to minimize the chance of filer disruption due to media error while a filer is in degraded or reconstruction mode.

Negligible performance impact: Because continuous media scrubbing searches only for media errors, the impact on system performance is negligible.

Note

Media scrubbing is a continuous background process. Therefore, you might observe disk LEDs blinking on an apparently idle system. You might also observe some CPU activity even when no user workload is present. The media scrub attempts to exploit idle disk bandwidth and free CPU cycles to make faster progress. However, any client workload results in aggressive throttling of the media scrub resource.

Not a substitute for a scheduled disk scrub: Because the continuous process described in this section scrubs only media errors, Network Appliance strongly recommends that you continue to run the filer's scheduled complete disk scrub operation, which is described in "Scheduling or manually running disk scrubbing" on page 186. The complete disk scrub carries out parity and checksum checking and repair operations, in addition to media checking and repair operations, on a scheduled rather than a continuous basis.

Adjusting maximum time for a media scrub cycle

You can lessen the CPU resource that would be consumed by the continuous media scrub under a heavy client workload by increasing the maximum time allowed for a media scrub cycle to complete from the default 72 hours.

By default, one cycle of a filer's continuous media scrub can take a maximum of 72 hours to complete. In most situations, one cycle of the continuous media scrub will complete in a much shorter time; however, under heavy client workload conditions, the default 72-hour maximum ensures that whatever the client load on the filer, enough CPU resource will be allotted to the media scrub to ensure that one cycle will complete in no more than 72 hours.

If you want the media scrub operation to consume even less CPU resource under heavy filer client workload, you can increase the maximum number of hours for the media scrub cycle, thus lessening the CPU resource consumed by the media scrub in times of heavy filer usage.

To extend or shorten the maximum time for a media scrub cycle, enter the following command at the filer command line:

```
options raid.media_scrub.deadline max_hrs_per_cycle
```

max_hrs_per_cycle is the maximum number of hours that you want to allow for one cycle of the continuous media scrub. Valid options range from 72 to 336 hours.

Disabling continuous media scrubbing

Network Appliance recommends that you keep continuous media error scrubbing enabled, but you might decide to disable your continuous media scrub in situations in which your filer is carrying out operations with heavy performance impact and in which you have alternative measures (such as volume SyncMirror replication or RAID DP configuration) in place that prevent data loss due to filer disruption or double disk failure.

To disable continuous media scrubbing, enter the following command on the filer command line:

```
options raid.media_scrub.enable off
```

Note

To restart continuous media scrubbing after you have disabled it, enter the following command:

```
options raid.media_scrub.enable on
```

Checking media scrub activity

You can confirm media scrub activity on your filer by entering the following command:

```
vol media_scrub status [/vol[/plex][/raidgroup]] [-v]
```

[*/vol*[/*plex*] [/*raidgroup*]] is the optional pathname to the volume, plex, or RAID group on which you want to confirm media scrubbing activity.

[-v] is the optional parameter for the verbose version of the media scrubbing activity status. The verbose version of status displays the percentage of the current scrub that is complete, the start time of the current media scrub, and the completion time of the last media scrub.

Example 1. Checking of filer-wide media scrubbing: The following command displays media scrub status information for all the volumes on the filer.

```
vol media_scrub status
vol media_scrub /vol0/plex0/rg0 is 28% complete
vol media_scrub /vol2/plex0/rg0 is 3% complete
vol media_scrub /vol3/plex0/rg0 is 3% complete
vol media_scrub /vol3/plex0/rg1 is 3% complete
vol media_scrub /vol3/plex0/rg2 is 3% complete
vol media_scrub /vol3/plex0/rg3 is 3% complete
```

Example 2. Verbose checking of filer-wide media scrubbing: The following command displays verbose media scrub status information for all the volumes on the filer.

```
vol media_scrub status -v
vol media_scrub: status of /vol0/plex0/rg0 :
  Current instance of media_scrub is 32% complete.
  Media scrub started at Tue Aug 12 19:24:00 GMT 2003
  Last full media_scrub completed: Tue Aug 12 19:23:39 GMT 2003

vol media_scrub: status of /vol2/plex0/rg0 :
  Current instance of media_scrub is 8% complete.
  Media scrub started at Tue Aug 12 19:30:00 GMT 2003
  Last full media_scrub completed: Tue Aug 12 19:27:37 GMT 2003
```

Example 3. Checking for media scrubbing on a specific volume: The following command displays media scrub status information for the volume *vol2*.

```
vol media_scrub status /vol2
  vol media_scrub /vol2/plex0/rg0 is 4% complete
  vol media_scrub /vol2/plex0/rg1 is 10% complete
```

Disk media error failure thresholds

About media error thresholds

To prevent filer panics or double disk failure due to excessive media errors encountered during single-disk failure reconstruction, Data ONTAP tracks media errors on each active filer disk and sends a disk failure request to the RAID system if system-defined media error thresholds are crossed on that disk.

Disk media error thresholds that trigger a disk failure request include

- More than 25 media errors (that are not related to disk scrub activity) occurring on a disk within a 10-minute period
- Three or more media errors occurring on the same sector of a disk

If the filer volume is not already running in degraded mode due to single-disk failure reconstruction when the disk failure request is received, RAID will fail the disk in question, swap in a hot spare disk, and begin RAID reconstruction to replace the failed disk.

Failing disks at the thresholds listed in this section ensures that a disk will be replaced before the number of media errors on that disk increases the likelihood of a filer panic or double-disk failure during a single-disk failure reconstruction.

Viewing RAID status

About RAID status

You can use the `volume status` command to check the current RAID status and configuration for your filer volumes. To view RAID status for your filer volumes, complete the following step:

Step	Action
1	To view the RAID status of your volumes, enter the following command: **`vol status [`*`vol_name`*`] -r`** *vol_name* is the name of the volume whose RAID status you want to view. **Note** If you omit a volume name, Data ONTAP displays the RAID status of all the volumes on the filer.

Possible RAID statuses displayed

The `volume status -r` command displays the following possible statuses that pertain to RAID:

- Normal—The volume is unmirrored, and all of its RAID groups are functional.
- Mirrored—The volume is mirrored, and all of its RAID groups are functional.
- Degraded—The volume contains at least one degraded RAID group that is not being reconstructed after single disk failure.
- Double degraded—The volume contains at least one RAID group with double disk failure that is not being reconstructed (this state is possible if RAID DP protection is enabled for the affected volume).
- Mirror degraded—The volume is a mirrored volume, and one of its plexes is offline or resynchronizing.
- Reconstruction xx% complete—At least one RAID group in the volume is being reconstructed after experiencing a single disk failure.
- Double reconstruction xx% complete—At least one RAID group in the volume is being reconstructed after experiencing a double disk failure

(this state is possible if RAID DP protection is enabled for the affected volume).

- Partial—At least one disk was found for the volume, but two or more disks are missing.
- Resyncing—The volume contains two plexes; one plex is resynchronizing with the volume.

Sample volume status output

The following output for `vol status` *`vol_name`* `-r` command indicates that volume vol1 is reconstructing data after a data disk failure was detected in the RAID group rg0.

```
>vol status vol1 -r
Volume vol1 (online, raid_dp, reconstruct, degraded)
    (zoned checksums)
  Plex /vol1/plex0 (online, normal, active)
    RAID group /vol1/plex0/rg0 (reconstruction 23% completed,
      degraded)

RAID Disk Device HA SHELF BAY CHAN Used (MB/blks) Phys (MB/blks)
-------- ------ ----------- ---- ------------ -------------
dparity   v0.2  v0  0    2  FC:A 36/74752     42/87168
parity    v0.4  v0  0    4  FC:A 36/74752     42/87168
(reconstruction 23% completed)
data      v0.3  v0  0    3  FC:A 36/74752     42/87168
data      v0.5  v0  0    5  FC:A 36/74752     42/87168
data      v0.6  v0  0    6  FC:A 36/74752     42/87168
data      FAILED      N/A        36/74752
data      v0.9  v0  0    9  FC:A 36/74752     42/87168
```

Disk Management 5

About this chapter

This chapter discusses how the filer uses disks, how disks are dually connected to the filer for redundancy, how to add disks to the filer so that Data ONTAP can use them, and how to remove disks from the filer and Data ONTAP.

For instructions about adding or removing disks from a disk shelf physically or determining the location of a disk in a disk shelf, see your disk shelf documentation.

Topics in this chapter

This chapter discusses the following topics:

- "Understanding disks on the filer" on page 202
- "Multipath I/O for Fibre Channel disks" on page 205
- "Storage subsystem management" on page 208
- "Adding disks" on page 221
- "Removing disks" on page 227
- "Ascertaining free disk space" on page 231
- "Disk sanitization" on page 233
- "Disk ownership assignment on an FAS270/FAS270c appliance" on page 247

Understanding disks on the filer

About disks

Removable disks and disks that you add are in disk shelves that are connected to a filer. The filer supports SCSI disks in SCSI disk shelves and Fibre Channel disks in Fibre Channel disk shelves.

Note

The NearStore R100 and R150 appliances also support ATA disks in SCSI disk shelves.

Kinds of disks

There are four kinds of disks, as shown in the following table.

Disk	Description
Data	Holds data stored on behalf of clients and data generated about the state of the filer as a result of a malfunction
Hot spare	Does not hold usable data, but is available for addition to a volume
Parity	Stores data reconstruction information
dParity	If RAID DP is enabled, stores the double-parity information

Disk formats

Disks are formatted for the following two types of RAID parity checksums:

- Block checksums—You can use block checksum disks for block checksum volumes and zoned checksum volumes.
- Zoned checksums—You can use zoned checksum disks only for zoned checksum volumes.

You cannot change the format of a disk.

Disk addressing

You identify a disk by its address, which is listed in the Device column of the `sysconfig -d` output. You can also use the `storage show` or `storage show disk -p` command to view the disk addresses. Disk addresses are represented in the following format:

path_id.device_id

path_id is generally the slot number on the filer where the host adapter is attached. However, certain exceptions exist. For example, the *path_id* for the onboard Fibre Channel card is 0a, and the *path_id* for the dual channel adapter is represented by *slot_number* a|b.

device_id is a protocol-specific identifier for the attached disks. For SCSI, the *device_id* is an integer from 0 to 15. For FC-AL, the *device_id* is an integer from 0 to 126.

SCSI disk addresses: For a SCSI disk, *path_id* is the disk's SCSI adapter number and *device_id* is the disk's SCSI ID number. For example, disks attached to the disk adapter 7a are numbered 7a.0 through 7a.14.

Fibre Channel disk addresses: For a Fibre Channel disk, *path_id* is the disk's host adapter number and *device_id* is the disk's ID. To identify the disk ID, see the `fcstat device_map` command output on your filer.

Available space on new disks

When you add a new disk, Data ONTAP reduces the amount of disk space available for user data by rounding down to maintain compatibility across disks from various manufacturers. The available disk space listed by informational commands such as `sysconfig` is, therefore, less for each disk than its rated capacity.

The available disk space on a disk is rounded down as shown in the following table.

Disk class	Available space	Available blocks
4-GB disks	4 GB	8,192,000
9-GB disks	8.6 GB	17,612,800
18-GB disks	17 GB	34,816,000
35-GB disks (block checksum disks)	34 GB	69,632,000
36-GB disks (zoned checksum disks)	34.5 GB	70,656,000
72-GB disks	68 GB	139,264,000
144-GB disks	136 GB	278,528,000

Disk class	Available space	Available blocks
160-GB disks (available on R100 appliances)	136	278,258,000
250-GB disks (available on R150 appliances)	212	434,176,000

Multipath I/O for Fibre Channel disks

Understanding Multipath I/O

The Multipath I/O feature for Fibre Channel disks enables you to connect each disk on a single filer to two different host adapters. As a result, there are two access paths to each disk. One path is designated primary and the other secondary. If the primary path to the disk becomes unavailable, the secondary path is automatically used for accessing the disk.

Advantages of Multipath I/O

By providing redundant paths to the same disk on a single filer, the Multipath I/O feature provides the following advantages:

- Overall reliability and uptime of the storage subsystem of the filer is increased.
- Disk availability is higher.
- Storage subsystem hardware can be maintained with no downtime.

  When a primary host adapter is brought down, all traffic moves from that host adapter to the secondary host adapter. As a result, you can perform maintenance tasks, such as replacing a malfunctioning LRC module or cables connecting that host adapter to the disk shelves, without affecting the storage subsystem service.

Requirements to enable Multipath I/O on the filer

The Multipath I/O feature is enabled automatically, subject to the following restrictions:

- Only the F740, F760, F8XX, and FAS900 series filer platforms support Multipath I/O.

  Note

  None of the NearStore appliance platforms (R100 or R150) support Multipath I/O.

- Only the 2200 (P/N X2040B) or 2212 (X2044A) (2044B and X2050A) host adapter installed in the appropriate slot support Multipath I/O.

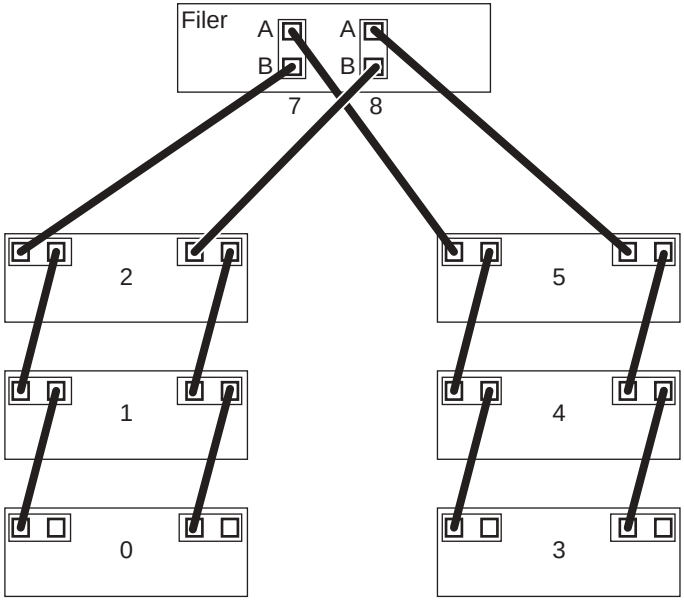
  Although the 2200 and 2212 host adapters can co-exist with older (2100 and 2000) adapters on a filer, Multipath I/O is not supported one the older models.

  To determine the slot number where a host adapter can be installed in your filer, see the *System Configuration Guide* at http://now.netapp.com.

- FC7 and FC8 disk shelves do not support Multipath I/O.
- FC9 and DS14 disk shelves must have 2 LRCs to support Multipath I/O.
- Older 9-GB disks (ST19171FC) and older 18-GB disks (ST118202FC) do not support Multipath I/O.
- Filers configured in clusters (that are not Fabric Metroclusters) do not support Multipath I/O. Multipath I/O setup and clustering setup both require the A and B ports of the disk shelves. Therefore, it is not possible to have both features enabled simultaneously.
- Hardware connections must be set up for Multipath I/O as specified in the corresponding *Fibre Channel StorageShelf* guide.

Example of Multipath I/O setup

The following figure illustrates an example of Multipath I/O setup on a filer. The filer has dual-port host adapters installed in slot 7 and slot 8. The host adapter ports connect to the disk shelf ports as shown.



This setup uses two dual-port host adapters to connect to two disk shelf loops, thus preventing the host adapters from being the single points of failure. Although it is not necessary to have a dual-port host adapter to set up Multipath I/O, Network Appliance recommends this setup because using dual-port adapters conserves PCI slots.

Storage subsystem management

Command for managing the storage subsystem

You can use the `storage` command to manage disks, SCSI and Fibre Channel host adapters, and all other components of the storage subsystem connected to your filer.

With the `storage` command, you can enable or disable a host adapter, view multiple host adapter paths to disks, and view information about hubs, disks, tape devices, and medium changer devices connected to your filer.

The storage command syntax and subcommands

The `storage` command syntax is as follows:

```
storage <subcommand>
```

The subcommands are

```
alias [alias {electrical_name | wwn }]
disable adapter name
enable adapter name
help sub_command
show adapter [-a] [name ]
show hub [-a] [name ]
show disk [-a | -p]
show mc [name]]
show port [name]
show switch [name]
show tape [name]
show tape supported [-v]
stats tape name
stats tape zero name
unalias {alias | -a | -m | -t }
```

For detailed information about the `storage` command and all options available with this command, see the na_storage(1) man page on the filer. The options `alias` and `unalias` are discussed in detail in the *Data Protection Guide*.

Detailed information

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

- "Changing the state of a host adapter" on page 209
- "Viewing storage subsystem information" on page 211

Changing the state of a host adapter

About the state of a host adapter

A host adapter can be enabled or disabled. You can change the state of a host adapter by using the `storage` command.

When to change the state of a host adapter

Disable: You might want to disable a host adapter if

- You are replacing any of the hardware components connected to the host adapter, such as cables and GBICs
- You are replacing a malfunctioning LRC module or bad cables

You can disable a host adapter only if all disks connected to it can be reached through another adapter. Consequently, SCSI adapters and adapters connected to single-attached devices cannot be disabled.

If you try to disable an adapter that is connected to disks with no redundant access paths, you will get the following error message:

"`Some device(s) on host adapter` *`n`* `can only be accessed through this adapter; unable to disable adapter`"

After an adapter connected to dual-connected disks has been disabled, the other adapter is not considered redundant; thus, the other adapter cannot be disabled.

Enable: You might want to enable a disabled host adapter after you have performed maintenance.

Enabling or disabling a host adapter

To enable or disable a host adapter, complete the following steps.

Step	Action
1	Enter the following command to identify the name of the host adapter whose state you want to change: **`storage show adapter`** **Result:** The field that is labeled "The Slot:" lists the host adapter name.

Step	Action	
2	**If you want to...**	**Then...**
	Enable the host adapter	Enter the following command: **`storage enable adapter`** ***`name`*** *name* is the host adapter name.
	Disable the host adapter	Enter the following command: **`storage disable adapter`** ***`name`*** *name* is the host adapter name.

Viewing storage subsystem information

Command for viewing information

You can use the `storage show` command and its options to view information about the storage subsystem components connected to your filer, including the host adapters, disks, tape devices, and medium changer devices.

For detailed information about the options available with the `storage show` command, see the na_storage(1) man page on the filer.

Viewing information about all components

To view information about storage subsystem components, complete the following step.

Step	Action
1	Enter the following command: **`storage show`**

Example: The following example shows information about the host adapters and disks connected to the filer tpubs-cf1:

```
tpubs-cf1> storage show
Slot:            7
Description:     Fibre Channel Host Adapter 7 (QLogic 2100 rev. 3)
Firmware Rev:    1.19.14
PCI Bus Width:   32-bit
PCI Clock Speed: 33 MHz
FC Node Name:    2:000:00e08b:006a15
Cacheline Size:  128
FC Packet Size:  512
Link Data Rate:  2 Gbit
SRAM Parity:     No
External GBIC:   No
State:           Enabled
In Use:          Yes
Redundant:       No
```

```
DISK  SHELF  BAY  SERIAL       VENDOR   MODEL      REV
----- -----  ---  ------       ------   ---------  ----
7.6     0     6   LA774453     SEAGATE  ST19171FC  FB59
7.5     0     5   LA694863     SEAGATE  ST19171FC  FB59
7.4     0     4   LA781085     SEAGATE  ST19171FC  FB59
7.3     0     3   LA773189     SEAGATE  ST19171FC  FB59
7.14    1     6   LA869459     SEAGATE  ST19171FC  FB59
7.13    1     5   LA781479     SEAGATE  ST19171FC  FB59
7.12    1     4   LA772259     SEAGATE  ST19171FC  FB59
7.11    1     3   LA783073     SEAGATE  ST19171FC  FB59
7.10    1     2   LA700702     SEAGATE  ST19171FC  FB59
7.9     1     1   LA786084     SEAGATE  ST19171FC  FB59
7.8     1     0   LA761801     SEAGATE  ST19171FC  FB59
7.2     0     2   LA708093     SEAGATE  ST19171FC  FB59
7.1     0     1   LA773443     SEAGATE  ST19171FC  FB59
7.0     0     0   LA780611     SEAGATE  ST19171FC  FB59
```

Viewing information about host adapters

To view information about host adapters, complete the following step.

Step	Action	
1	**If you want to view...**	**Then...**
	Information about all the host adapters	Enter the following command: **`storage show adapter`**
	Information about a specific host adapter	Enter the following command: **`storage show adapter`** ***`name`*** *name* is the adapter name.

Example 1: The following example shows information about all the host adapters installed in the filer tpubs-cf2:

```
tpubs-cf2> storage show adapter
Slot:            7
Description:     Fibre Channel Host Adapter 7 (QLogic 2100 rev. 3)
Firmware Rev:    1.19.14
PCI Bus Width:   32-bit
PCI Clock Speed: 33 MHz
FC Node Name:    2:000:00e08b:00fb15
Cacheline Size:  128
FC Packet Size:  512
Link Data Rate:  2 Gbit
```

```
SRAM Parity:     No
External GBIC:   No
State:           Enabled
In Use:          Yes
Redundant:       No
Slot:            8
Description:     Fibre Channel Host Adapter 8 (QLogic 2100 rev. 3)
Firmware Rev:    1.19.14
PCI Bus Width:   32-bit
PCI Clock Speed: 33 MHz
FC Node Name:    2:000:00e08b:006b15
Cacheline Size:  128
FC Packet Size:  512
Link Data Rate:  2 Gbit
SRAM Parity:     No
External GBIC:   No
State:           Enabled
In Use:          Yes
Redundant:       No
```

Example 2: The following example shows information about host adapter 8 in the filer tpubs-cf2:

```
tpubs-cf2> storage show adapter 8
Slot:            8
Description:     Fibre Channel Host Adapter 8 (QLogic 2100 rev. 3)
Firmware Rev:    1.19.14
PCI Bus Width:   32-bit
PCI Clock Speed: 33 MHz
FC Node Name:    2:000:00e08b:006b15
Cacheline Size:  128
FC Packet Size:  512
Link Data Rate:  2 Gbit
SRAM Parity:     No
External GBIC:   No
State:           Enabled
In Use:          Yes
Redundant:       No
```

Viewing information about hubs

To view information about hubs, complete the following step.

Step	Action	
1	**If you want to view...**	**Then...**
	Information about all hubs	Enter the following command: **`storage show hub`**
	Information about a specific hub	Enter the following command: **`storage show hub`** ***`name`*** *name* is the hub name.

Example: The following example shows information about hub 8a.shelf1:

```
storage show hub 8a.shelf1
Hub name:    8a.shelf1
Channel:     8a
Loop:        A
Shelf id:    1
Shelf UID:   50:05:0c:c0:02:00:12:3d
Term switch: OFF
Shelf state: ONLINE
ESH state:   OK
                  Loop Invalid  Invalid Clock  Insert Stall Util
Disk  Disk   Port     up       CRC     Word   Delta  Count Count %
Id    Bay    State    Count    Count   Count
-------------------------------------------------------------------
[IN ]        OK   3    0        0      128    1      0     0
[ 16]  0     OK   4    0        0      128    0      0     0
[ 17]  1     OK   4    0        0      128    0      0     0
[ 18]  2     OK   4    0        0      128    0      0     0
[ 19]  3     OK   4    0        0      128    0      0     0
[ 20]  4     OK   4    0        0      128    0      0     0
[ 21]  5     OK   4    0        0      128    0      0     0
[ 22]  6     OK   4    0        0      128    0      0     0
[ 23]  7     OK   4    0        0      128    0      0     0
[ 24]  8     OK   4    0        0      128    0      0     0
[ 25]  9     OK   4    0        0      128    0      0     0
[ 26]  10    OK   4    0        0      128    0      0     0
[ 27]  11    OK   4    0        0      128    0      0     0
[ 28]  12    OK   4    0        0      128    0      0     0
[ 29]  13    OK   4    0        0      128    0      0     0
[OUT]        OK   4    0        0      128    0      0     0
```

```
Hub name:    8b.shelf1
Channel:     8b
Loop:        B
Shelf id:    1
Shelf UID:   50:05:0c:c0:02:00:12:3d
Term switch: OFF
Shelf state: ONLINE
ESH state:   OK
                 Loop  Invalid  Invalid Clock  Insert Stall Util
Disk  Disk   Port     up       CRC     Word   Delta  Count Count %
Id    Bay    State    Count    Count   Count
-------------------------------------------------------------------
[IN ]        OK   3    0        0      128    1      0     0
[ 16]  0     OK   4    0        0      128    0      0     0
[ 17]  1     OK   4    0        0      128    0      0     0
[ 18]  2     OK   4    0        0      128    0      0     0
[ 19]  3     OK   4    0        0      128    0      0     0
[ 20]  4     OK   4    0        0      128    0      0     0
[ 21]  5     OK   4    0        0      128    0      0     0
[ 22]  6     OK   4    0        0      128    0      0     0
[ 23]  7     OK   4    0        0      128    0      0     0
[ 24]  8     OK   4    0        0      128    0      0     0
[ 25]  9     OK   4    0        0      128    0      0     0
[ 26]  10    OK   4    0        0      128    0      0     0
[ 27]  11    OK   4    0        0      128    0      0     0
[ 28]  12    OK   4    0        0      128    0      0     0
[ 29]  13    OK   4    0        0      128    0      0     0
[OUT]        OK   4    0        0      128    0      0     0
```

Note

Hub 8b.shelf1 is also listed by the `storage show hub 8a.shelf1` command in the example because the two hubs are part of the same shelf and the disks in the shelf are dual-ported disks. Effectively, the command is showing the disks from two perspectives.

Viewing information about all disks

To view information about all the disks connected to your filer, complete the following step.

Step	Action
1	Enter one of the following commands: **`storage show disk`** or **`storage show disk -a`** **Note** Both commands display information about all disks in different formats. The `-a` option displays information about all disks in a report form that is easily interpreted by scripts.

Example: The following example shows information about all the disks connected to the filer tpubs-cf1:

```
tpubs-cf1> storage show disk

DISK  SHELF  BAY    SERIAL     VENDOR   MODEL       REV
----  -----  ---    --------   ------   ---------   ----
7.6     0     6     LA774453   SEAGATE  ST19171FC   FB59
7.5     0     5     LA694863   SEAGATE  ST19171FC   FB59
7.4     0     4     LA781085   SEAGATE  ST19171FC   FB59
7.3     0     3     LA773189   SEAGATE  ST19171FC   FB59
7.14    1     6     LA869459   SEAGATE  ST19171FC   FB59
7.13    1     5     LA781479   SEAGATE  ST19171FC   FB59
7.12    1     4     LA772259   SEAGATE  ST19171FC   FB59
7.11    1     3     LA783073   SEAGATE  ST19171FC   FB59
7.10    1     2     LA700702   SEAGATE  ST19171FC   FB59
7.9     1     1     LA786084   SEAGATE  ST19171FC   FB59
7.8     1     0     LA761801   SEAGATE  ST19171FC   FB59
7.2     0     2     LA708093   SEAGATE  ST19171FC   FB59
7.1     0     1     LA773443   SEAGATE  ST19171FC   FB59
7.0     0     0     LA780611   SEAGATE  ST19171FC   FB59
```

Viewing the primary and secondary paths to the disks

To view the primary and secondary paths to all the disks connected to your filer, complete the following step.

Step	Action
1	Enter the following command: `storage show disk -p`

Note

The disk addresses shown for the primary and secondary paths to a disk are aliases of each other.

Example: In the following example, a host adapter is installed in the PCI expansion slot 7 and slot 8 of the filer tpubs-cf1. The slot 8 host adapter is connected to port A of disk shelf 0 and the slot 7 host adapter is connected to port B of disk shelf 1. The output of the `storage show disk -p` command shows the primary and secondary paths to all disks connected to the filer.

```
tpubs-cf1> storage show disk -p
PRIMARY PORT  SECONDARY  PORT  SHELF BAY
------- ----  ---------- ----  ----- ---
7b.0    B     8a.0       A       0    0
8a.1    A     7b.1       B       0    1
8a.2    A     7b.2       B       0    2
7b.3    B     8a.3       A       0    3
7b.4    B     8a.4       A       0    4
8a.5    A     7b.5       B       0    5
8a.6    A     7b.6       B       0    6
7b.8    B     8a.8       A       1    0
8a.9    A     7b.9       B       1    1
7b.10   B     8a.10      A       1    2
7b.11   B     8a.11      A       1    3
7b.12   B     8a.12      A       1    4
7b.13   B     8a.13      A       1    5
8a.14   A     7b.14      B       1    6
```

Viewing information about medium changers

To view information about medium changers attached to your filer, complete the following step.

Step	Action
1	Enter the following command: **`storage show mc [ name ]`** *name* is the name of the medium changer for which you want to view information. If no medium changer name is specified, information for all medium changers is displayed.

Viewing information about switch ports

To view information about ports for switches attached to the filer, complete the following step.

Step	Action
1	Enter the following command: **`storage show port [ name ]`** *name* is the name of the port for which you want to view information. If no port name is specified, information for all ports is displayed.

Viewing information about switches

To view information about switches attached to the filer, complete the following step.

Step	Action
1	Enter the following command: **`storage show switch [ name ]`** *name* is the name of the switch for which you want to view information. If no switch name is specified, information for all switches is displayed.

Viewing information about tape drives

To view information about tape drives attached to your filer, complete the following step.

Step	Action
1	Enter the following command: **`storage show tape tape`** *tape* is the name of the tape drive for which you want to view information. If no tape name is specified, information for all tape drives is displayed.

Viewing supported tape drives

To view information about tape drives that are supported by your filer, complete the following step.

Step	Action
1	Enter the following command: **`storage show tape supported [-v]`** `-v` displays all information about supported tape drives, including their density and compression settings. If no option is given, only the list of supported tape drives is displayed.

Viewing tape drive statistics

To view storage statistics for tape drives attached to the filer, complete the following step.

Step	Action
1	Enter the following command: **`storage stats tape name`** *name* is the name of the tape drive for which you want to view storage statistics.

Resetting tape drive statistics

To reset storage statistics for a tape drive attached to the filer, complete the following step.

Step	Action
1	Enter the following command: `storage stats tape zero` *`name`* *name* is the name of the tape drive.

Adding disks

About adding disks

When you add disks to a filer, you need to insert them in a disk shelf according to the disk shelf manufacturer's instructions. They then become hot spare disks. You also add the disks to a volume with the `vol add` command.

You can use disks of various sizes in a volume according to the following rules:

- If an added disk replaces a failed data disk, its capacity is limited to that of the failed disk.
- If an added disk is not replacing a failed data disk and it is not larger than the parity disk, its full capacity (subject to rounding) is available as a data disk.

When to add disks

You add disks when

- You want to add storage capacity to the filer
- You are running out of hot spare disks

Considerations when adding disks to the filer

Original disk configuration: Network Appliance recommends adding RAID groups and/or disks before the file system is 80% to 90% full.

The number of disks and RAID groups at initial filer configuration affects read and write performance. A greater number of disks means a greater number of independently seeking disk-drive heads reading data. Write performance can also benefit from more disks; however, the difference can be masked by the effect of NVRAM and the manner in which WAFL manages write operations.

Adding more disks to the filer: Network Appliance recommends adding disks in multiples of 8 to 14 for overall improved performance. However, as more disks are configured, the performance increase levels off. Performance is affected more with the fewer disks you add until the striping across all the disks levels out. When the striping levels out, there is an increase in operations per second and a reduced response time.

The ideal time to add more disks to meet future storage requirements is before the applied load places stress on the existing array of disks, even though adding more disks at this time will not significantly improve the filer's current performance.

Adding hot spare disks to the filer

You should periodically check the number of hot spares you have on your filer and add disks if you have fewer than one per shelf. If there are no hot spare disks, you should add disks to disk shelves so they become available as hot spares.

For more information, see "Hot spare disks" on page 162.

Ascertaining the number of hot spare disks

To ascertain how many hot spare disks you have on your filer, complete the following step.

Step	Action

1	Enter the following command:
	`vol status -s`
	Result: If there are hot spare disks, a display like the following appears, with a line for each spare disk, grouped by checksum type:
	`Spare disks for block or zoned checksum volumes`
	`RAID Disk HA.ID HA SHELF BAY CHAN Used (MB/blks) Phys(MB/blks)`
	`--------- ----- ------------ ---- -------------- --------------`
	`spare 7.17 7 2 1 FC:A 0 17046/34910224`
	`spare 7.26 7 3 2 FC:A 0 34209/70061520`
	`...`
	`Spare disks for zoned checksum volumes only`
	`RAID Disk HA.ID HA SHELF BAY CHAN Used (MB/blks) Phys(MB/blks)`
	`--------- ----- ------------ ---- -------------- --------------`
	`spare 7.21 7 2 5 FC:A 0 17366/35566480`
	`spare 7.19 7 2 3 FC:A 0 17366/35566480`

Rules for adding disks to a volume

You add disks to a volume according to the following rules:

- You can add only hot spare disks to a volume.
- You must specify the volume to which you are adding the disks.
- You can add block checksum disks to zoned checksum volumes; you can add only block checksum disks to block checksum volumes.
- The disks must come from the same spare disk pool (if you are using mirrored volumes).

Parity disk determination

If an added disk is larger than an existing parity disk, the added disk replaces the smaller disk as the parity disk, and the smaller disk becomes a data disk. This enforces a Data ONTAP rule that the parity disk must be at least as large as the largest data disk in a RAID group.

Note

In volumes configured with RAID DP (which uses RAID groups that contain both a regular parity disk and a double, or dParity, disk) the larger added disk replaces any smaller regular parity disk, but its capacity will be downsized, if necessary, to avoid exceeding the capacity of any smaller-sized the dParity disk.

Adding new disks to the filer

To add new disks to the filer, complete the following steps.

Note

In the following procedure, the NearStore R100 appliance and the original-capacity R150 appliance use disks on SCSI-attached shelves. The R150 appliance with upgraded 24 TB capacity and the R200 appliance use disks on Fibre-Channel-attached shelves.

Step	Action	
1	**If the disks are...**	**Then...**
	Fibre Channel disks or on Fibre-Channel-attached shelves	Go to Step 2.
	SCSI disks or on SCSI-attached shelves	Enter the following command, and go to Step 2: `disk swap`

Step	Action
2	Install one or more disks according to the hardware guide for your disk shelf. **Note** On FAS270 and FAS270c appliances, a disk ownership assignment might need to be carried out. See “Disk ownership assignment on an FAS270/FAS270c appliance” on page 247. **Result:** The system displays a message confirming that one or more disks were installed and then waits 15 seconds as the disks are recognized. The system recognizes the disks as hot spare disks. **Note** If you add multiple disks, the system might require 25 to 40 seconds to bring the disks up to speed as it checks the device addresses on each adapter.
3	Verify that the disks were added by entering the following command: `vol status -s` **Result:** The number of hot spare disks in the RAID Disk column under Spare Disks increases by the number of disks you installed.

Adding new disks to a volume

To add new disks to a volume, complete the following steps.

Step	Action
1	Verify that hot spare disks are available for you to add by entering the following command: `vol status -s`

Step	Action

2	Add the disks to a volume by entering the following command:
	`vol add vol-name ndisks[@disk-size] | -d disk-name...`
	vol-name is the name of the volume to which you are adding the disks.
	ndisks is the number of disks to use.
	disk-size is the disk size, in gigabytes, to use. You must have at least *ndisks* available disks of the size you specify.
	disk-name is the disk ID of one or more spare disks; use a space to separate disk IDs. The disk ID is under the HA.ID column in the `vol status -s` display.
	Note
	If you want to use block checksum disks in a zoned checksum volume and there are still zoned checksum hot spare disks, use the `-d` option to select the disks.
	Examples: The following command adds four 36-GB disks to the thisvol volume:
	`vol add thisvol 4@36`
	The following command adds the disks 7.17 and 7.26 to the thisvol volume:
	`vol add thisvol -d 7.17 7.26`

Note

When adding disks to a volume, Data ONTAP attempts to distribute I/O loads, such as alternating host adapters, to improve reconstruction performance. You can override this behavior by specifying the disks you want.

Adding new disks to a RAID group

If a volume has more than one RAID group, you can specify the RAID group to which you are adding disks. To add new disks to a specific RAID group of a volume, complete the following step.

Step	Action
1	Enter the following command: `vol add` *`vol-name ndisks`*`[@`*`disk-size`*`]	-d` *`disk-name...`* `-g` *`raidgroup`* *raidgroup* is a RAID group in the volume specified by *vol-name*. **Example:** The following command adds two disks to RAID group 0 of the vol0 volume: `vol add vol0 -g rg0 2`

Note

The number of disks you can add to a specific RAID group is limited by the `raidsize` setting of the volume to which that group belongs. See Chapter 4, "RAID Group Management," on page 159 for more information.

Forcibly adding disks to volumes

The following are situations that would fail if you tried to add disks to a volume:

- The disks specified in the `vol add` command would cause the disks on an unmirrored volume to span adapters.
- The disks specified in the `vol add` command add disks to a RAID group other than the last RAID group, thereby making it impossible for the file system to revert to a version prior to Data ONTAP 6.4.

To force Data ONTAP to add disks in these situations, complete the following step.

Step	Action
1	Enter the following command: `vol add` *`vol-name`* `-f [-g raidgroup] -d` *`disk-name...`* **Note** You must use the `-g` *raidgroup* option to specify a RAID group other than the last RAID group in the volume.

Removing disks

Reasons to remove disks

You remove a disk for the following reasons:

- It is a failed disk. You cannot use this disk again.
- It is a hot spare and you want to use it elsewhere.
- It is a data disk, is producing excessive error messages, and is about to fail. You cannot use this disk again.

Note

You cannot reduce the number of disks in a file system by removing data disks. The only way to reduce the number of data disks in a file system is to copy the data and transfer it to a new file system that has fewer data disks.

Removing a failed disk

To remove a failed disk, complete the following steps.

Note

In the following procedure, the NearStore R100 appliance and the original-capacity R150 appliance use disks on SCSI-attached shelves. The R150 appliance with upgraded 24 TB capacity and the R200 appliance use disks on Fibre-Channel-attached shelves.

Step	Action
1	Find the disk ID of the failed disk by entering the following command: `vol status -f` **Result:** The ID of the failed disk is shown next to the word *failed*. The location of the disk is shown to the right of the disk ID, in the column HA SHELF BAY.

Step	Action	
2	**If the disk is a...**	**Then...**
	Fibre Channel disk or in a Fibre-Channel-attached shelf	Go to Step 3.
	SCSI disk or in a SCSI-attached shelf	Enter the following command and go to Step 3: `disk swap`
3	Remove the disk from the disk shelf according to the disk shelf manufacturer's instructions.	

Removing a hot spare disk

To remove a hot spare disk, complete the following steps.

Note

In the following procedure, the NearStore R100 appliance and the original-capacity R150 appliance use disks on SCSI-attached shelves. The R150 appliance with upgraded 24 TB capacity and the R200 appliance use disks on Fibre-Channel-attached shelves.

Step	Action
1	Find the disk IDs of hot spare disks by entering the following command: `vol status -s` **Result:** The names of the hot spare disks appear next to the word *spare*. The locations of the disks are shown to the right of the disk name.
2	Enter the following command to spin down the disk: `disk remove disk_name` *disk_name* is the name of the disk you want to remove (from the output of Step 1).

Step	Action	
3	**If the disk is...**	**Then...**
	A Fibre Channel disk or in a Fibre Channel attached shelf	Go to Step 4.
	A SCSI disk or in a SCSI-attached shelf	Enter the following command, and go to Step 4: `disk swap`
4	Wait for the disk to stop spinning. See the hardware guide for your disk shelf model for information about how to tell when a disk stops spinning.	
5	Remove the disk from the disk shelf, following the instructions in the hardware guide for your disk shelf model. **Result:** File service resumes 15 seconds after you remove the disk.	

Removing a data disk that is about to fail

To remove a data disk that is about to fail, complete the following steps.

Note

In the following procedure, the NearStore R100 appliance and the original-capacity R150 appliance use disks on SCSI-attached shelves. The R150 appliance with upgraded 24 TB capacity and the R200 appliance use disks on Fibre-Channel-attached shelves.

Step	Action
1	Find the disk name in the log messages that report the errors by looking at the numbers that follow the word *Disk*.
2	At the filer, enter the following command: `vol status -r`
3	Look at the HA.ID column of the output of the `sysconfig -r` command. It shows the disk ID of each disk. The location of the disk appears to the right of the disk ID, in the column HA SHELF BAY.

Step	Action	
4	Enter the following command to fail the disk: **`disk fail disk_name`** *disk_name* is the disk name from the output in Step 1. **Result:** After you complete this step, the system operates in degraded mode, which means that a disk is missing from the RAID group.	
5	**If the disk is...**	**Then...**
	A Fibre Channel disk or in a Fibre Channel attached shelf	Go to Step 6.
	A SCSI disk or in a SCSI-attached shelf	Enter the following command, then go to Step 6: **`disk swap`**
6	Remove the disk from the disk shelf, following the instructions in the hardware guide for your disk shelf model. **Result:** File service resumes 15 seconds after you remove the disk.	

Ascertaining free disk space

Command for ascertaining free disk space

You use the `df` command to verify the amount of free disk space on the filer.

Disk space report discrepancies

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, the sum, 10,984,940 KB (about 11 GB), is significantly less than the total disk space installed (14 GB). The following causes 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.
- Like the UNIX FFS (Fast File System), the filer reserves 10% of the total disk space for efficiency, which `df` does not count as part of the file system space.

Note

Snapshot reserve, if activated, can also cause disk space report discrepancies. For more information, see the *Data Protection: Online Backup and Recovery* guide.

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 might 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.

Getting accurate qtree space information

To get accurate space information about a qtree, complete the following step.

Step	Action
1	Use the filer `df` command instead of the client `df` command.

Disk sanitization

About disk sanitization

Disk sanitization is the process of physically obliterating data by means of overwrite patterns on filer or NearStore appliance disks in a manner that precludes recovery of that data by any known recovery methods. The Data ONTAP disk sanitize feature enables you to carry out disk sanitization by using three successive byte overwrite patterns per cycle and a default six cycles per operation, in compliance with United States Department of Defense and Department of Energy security requirements.

You might choose to sanitize filer or NearStore appliance disks for the following reasons:

- You have some disks that you have deactivated for use in one filer or NearStore appliance and you either want to reuse those disks in another filer or NearStore appliance or simply dispose of them but want to ensure that data currently on those disks is physically unrecoverable in any case.
- You want to selectively sanitize information from specified files or directories in a filer or NearStore appliance volume while preserving the remainder of the data that is stored on that volume.

What this section covers

The section covers the following topics:

- "Disk sanitization limitations" on page 233
- "Licensing disk sanitization" on page 234
- "Sanitizing disks" on page 235
- "If formatting is interrupted" on page 238
- "Selectively sanitizing data" on page 238
- "Reading disk sanitization log files" on page 244

Disk sanitization limitations

The disk sanitization operations described in this section are subject to the following limitations:

- The operations are not supported on gFiler appliances.
- The operations are not supported in takeover mode on clustered systems or on disabled cluster systems.
- Disk sanitization cannot be carried out on disks that were failed due to readability or writability problems.

- Disk sanitization cannot be carried out on disks that have ever belonged to an SEC 17a-4-compliant SnapLock volume.
- The formatting phase of the disk sanitization process is skipped on Advanced Technology Attachment (ATA) drives.

Licensing disk sanitization

The disk sanitization feature for Data ONTAP 6.5 requires installation of a license. Once installed, the license enables the `disk sanitize` command and limited SnapMirror replication abilities that facilitate selective sanitization operations.

Caution

Once installed on a NetApp appliance, the license for disk sanitization is permanent.

To install the disk sanitization license, complete the following step:

Step	Action
1	Enter the following command: **`license add`** ***`sanitize_code`*** *sanitize_code* is the disk sanitization license code that Network Appliance provides.

Sanitizing disks

You can sanitize any disk that has “spare” status. This includes disks that exist on the filer or NearStore appliance as spare disks after the volume that they belong to has been destroyed and disks that were removed from the spare disk pool by the `disk remove` command but have been returned to spare status after an appliance reboot.

To sanitize a disk or a set of disks on a filer or NearStore appliance, complete the following steps:

Step	Action
1	To print a listing of all disks assigned to RAID groups, failed, or existing as spares, enter the following command. **`sysconfig -r`** Do this to verify that the disk or disks that you want to sanitize do not belong to any existing RAID group in any existing volume.

Step	Action
2	Enter the following command to physically sanitize the specified disk or disks of all existing data. **`disk sanitize start [-p pattern1 [-p pattern2] [-p pattern3]] [-c cycle_count] disk_list`** In the preceeding example, *pattern* specifies a cycle of one to three user-defined hex byte overwrite patterns that can be applied in succession to the disks being sanitized. The default hex pattern specification is -p 0x55 -p 0xAA -p 0x36. `-c cycle_count` specifies the number of cycles for applying the specified overwrite patterns. The default value is 6 cycles. The maximum value is 7 cycles. `disk_list` specifies a space-separated list of spare disks to be sanitized. **Example:** The following command applies the default three disk sanitization overwrite patterns for the default six cycles (for a total of 18 overwrites) to the specified disks, 7.6, 7.7, and 7.8. **`disk sanitize start 7.6 7.7 7.8`** **Result:** The specified disks are sanitized, put into the pool of broken disks, and marked as "sanitized." A list of all the sanitized disks is stored in the filer or NearStore appliance's `/etc` directory. **Verification:** To check the status of the disk sanitization process, enter the following command: **`disk sanitize status [disk_list]`** **Note** If you need to abort the sanitization operation, enter **`disk sanitize abort [disk_list]`** **Caution** Do not turn off the filer, disrupt the disk loop, or remove target disks during the sanitization process. If the sanitization process is disrupted, the target disks that are in the formatting stage of disk sanitization will require reformatting before their sanitization can be completed. See "If formatting is interrupted" on page 238.

Step	Action
3	To release sanitized disks from the pool of broken disks for reuse as spare disks, enter the following command: **`disk sanitize release`** ***`disk_list`*** **Caution** The `disk sanitize release` command removes the "sanitized" label from the affected disks and returns them to spare state. Rebooting the filer or removing the disk also removes the sanitized label from any sanitized disks and returns them to spare state. **Verification:** To list all disks on the filer or NearStore appliance and verify the release of the sanitized disks into the pool of spares, enter **`sysconfig -r`**.

Process description: After you invoke the `disk sanitize start` command, Data ONTAP begins the sanitization process on each of the specified drives. The process consists of a disk format operation, followed by the specified overwrite patterns repeated for the specified number of cycles.

Note
The formatting phase of the disk sanitization process is skipped on Advanced Technology Attachment (ATA) drives.

The time to complete the sanitization process for each drive varies with the size of the drive, the number of patterns specified, and the number of cycles specified.

For example, the following command invokes 1 format overwrite pass and 18 pattern overwrite passes of disk 7.3.

```
disk sanitize start -p 0x55 -p 0xAA -p 0x37 -c 6 7.3
```

- If disk 7.3 is 36 GB and each formatting or pattern overwrite pass on it takes 15 minutes, then the total sanitization time is 19 passes times 15 minutes, or 285 minutes or 4.75 hours.
- If disk 7.3 is 73 GB and each formatting or pattern overwrite pass on it takes 30 minutes, then total sanitization time is 19 passes times 30 minutes, or 570 minutes or 9.5 hours.

If disk sanitization is interrupted: If the sanitization process is interrupted by power failure, filer panic, or a user-invoked disk sanitize abort command, the `disk sanitize` command must be re-invoked and the process repeated from the beginning.

If formatting is interrupted: If the formatting phase of disk sanitization is interrupted, the Data ONTAP attempts to reformat any drives that were corrupted by an interruption of the formatting. After a system reboot and once every hour Data ONTAP checks for any sanitization target disk that did not complete the formatting phase of its sanitization. If such a disk is found, Data ONTAP attempts to reformat that disk and writes a message to the console informing the user that a corrupted disk has been found and will be reformatted. After the disk is reformatted, it is returned to the hot spare pool. The user can then rerun the `disk sanitize` command on that disk.

Stopping disk sanitization

You can use the `disk sanitize abort` command to stop an ongoing sanitization process on one or more specified disks. If you use the `disk sanitize abort` command, the specified disk or disks are returned to "spare" state and the "sanitized" label is removed.

Step	Action
1	Enter the following command: **`disk sanitize abort`** ***`disklist`*** Data Ontap displays the message "Sanitization abort initiated." If the specified disks are undergoing the disk formatting phase of sanitization, just prior to the overwrite pattern phase, the abort will not occur until the disk formatting is complete. Once the process is stopped, ONTAP displays the message "Sanitization aborted for *diskname*."

Selectively sanitizing data

Selective data sanitization consists of physically obliterating data in specified files or directories while preserving all other active file data located on the affected volume for continued user access.

Summary of the selective sanitization process: Because data for any one file on a filer or NearStore appliance is physically stored on any number of data disks in the volume containing that data, and because the physical location

of data within a volume can change, sanitization of selected data, such as files or directories, requires that you sanitize every disk in the volume where it is located (after first migrating the volume data that you do *not* want to sanitize to disks on another volume). To selectively sanitize data contained in a volume, you must carry out three general tasks.

1. Delete the selected files or directories (and any volume snapshots that contain those files or directories) from the volume that contains them.
2. Migrate the remaining data (the data that you want to preserve) in the affected volume to a new set of disks in a destination volume on the same filer or NearStore appliance.
3. Destroy the original volume and sanitize all the disks that belong to that volume.

Requirements for selective sanitization: Successful completion of this process requires the following conditions:

- You install a disk sanitization license which enables the disk sanitizing feature and also enables limited SnapMirror functionality, such as the `snapmirror initialize` and `snapmirror migrate` commands for local operations on your filer or NearStore appliance.
- You have enough storage space on your filer or NearStore appliance to create an additional destination volume to which you can migrate the data that you want to preserve from the original volume. This destination volume must have a storage capacity at least as large as that of the original volume.

Volume size and selective sanitization: Because sanitization of any unit of data in a volume still requires you to carry out data migration and disk sanitization processes on that entire volume, Network Appliance recommends that you use small volumes to store data that will require sanitization. Use of small volumes for storage of data requiring sanitization minimizes the time, disk space, and bandwidth that sanitization of that data will require.

Backup and data sanitization: Absolute sanitization of data means physical sanitization of *all* instances of volumes containing sensitive data; it is therefore advisable to maintain your sensitive data in volumes that are *not* regularly backed up to volumes that also back up large amounts of nonsensitive data.

Procedure for selective sanitization: To carry out selective sanitization of data within a volume, complete the following steps:

Step	Action
1	From a Windows or UNIX client, delete from the original volume, the directories or files whose data you want to selectively sanitize from the active file system. Use the appropriate Windows or UNIX command, such as: `rm -rf /nixdir/nixfile.doc`

2	Through the filer CLI, in the original volume, enter one of the following command forms to delete all volume snapshots containing the files or directories that you just deleted.
	◆ To delete all snapshots associated with the volume, enter the following command syntax:
	`snap delete -a vol_name`
	For example: `snap delete -a nixsrcvol`
	◆ To delete a specific snapshot, enter the following command syntax
	`snap delete vol_name snapshot_name`
	For example: `snap delete nixsrcvol nightly0`

3	Enter the following command to determine the size of the volume from which you deleted file, directory, or snapshot data: `vol status vol_name -b` **Example:** `vol status nixsrcvol -b` Calculate the volume size in bytes by multiplying the bytes per block (block size) by the blocks per volume (volume size).

Step	Action
4	Enter the following command to create a volume to which you will migrate undeleted data. This volume must be of equal or greater storage capacity than the volume from which you just deleted file, directory, or snapshot data: **`vol create dest_vol disklist`** **Example:** **`vol create nixdestvol 8@72G`** **Note** The purpose of this new volume is to provide a migration destination that is absolutely free of the data that you want to sanitize. No previous mirror or backup relationship between the new volume and the source volume can exist.
5	Enter the following command to put the volume that you just created in restricted state: **`vol restrict dest_vol`** Restricted state is a requirement of the SnapMirror destination volume. **Example:** **`vol restrict nixdestvol`**

Step	Action
6	Enter the following command to establish a SnapMirror source/destination relationship between the volume whose data you want to sanitize and the volume to which you will migrate the data that you want to preserve: **`snapmirror initialize -S src_vol dest_vol`** *src_vol* is the source volume. *dest_vol* is the destination volume. **Caution** Be sure that you have deleted the files or directories that you want to sanitize (and any affected snapshots) from the source volume *before* you run the `snapmirror initialize` command. **Example:** **`snapmirror initialize -S nixsrcvol nixdestvol`** **Verification:** To confirm that the SnapMirror relationship is established enter the following command: **`snapmirror status dest_vol`**

Step	Action
7	Enter the following command to migrate all the undeleted data from the source volume to the destination volume: **`snapmirror migrate`** ***`src_vol dest_vol`*** **Example:** **`snapmirror migrate nixsrcvol nixdestvol`** **Results:** The `snapmirror migrate` command transfers all content—including the contained qtrees, directories, files, and all NFS filehandles—from the source volume to the destination volume. The process does the following: ◆ Checks the source and destination volumes for readiness. ◆ Stops NFS and CIFS service to the source volume. This prevents changes to the source volume data, which makes it appear to clients as though nothing has changed during the migration. ◆ Runs a regular SnapMirror transfer between the two volumes. ◆ Migrates the NFS filehandles to the destination volume. ◆ Puts the source volume in restricted mode ◆ Makes the destination volume writable.
8	Record the disks currently in the source volume. (After that volume is destroyed, you will sanitize these disks.) To list the disks in the source volume, enter the following command syntax: **`vol status`** ***`src_vol`*** **`-r`** **Example:** **`vol status nixsrcvol -r`** The disks that you are going to sanitize are listed in the Device column of the `vol status` -r output.
9	Enter the following command to take the source volume offline: **`vol offline`** ***`src_vol`*** **Example:** **`vol offline nixsrcvol`**

Step	Action
10	Enter the following command to destroy the source volume: **`vol destroy src_vol`** **Example:** **`vol destroy nixsrcvol`**
11	Enter the following command to rename the destination volume, giving it the name of the source volume that you just destroyed: **`vol rename dest_vol src_vol`** **Example:** **`vol rename nixdestvol nixsrcvol`**
12	Reestablish your CIFS or NFS services: ◆ If the original volume supported CIFS services, restart the CIFS services on the destination volume after migration is complete. ◆ If the original volume supported NFS services, enter the following command: **`Run exportfs -a`** **Result:** Users who were accessing files in the original volume will access those files in the renamed destination volume with no remapping of their connections required.
13	Use the `disk sanitize` command to sanitize the disks that used to belong to the source volume. Follow the procedure described in “Sanitizing disks” on page 235.

Reading disk sanitization log files

The disk sanitization process outputs two types of log files.

- One file, sanitized_disks, lists all the drives that have been sanitized.
- For each disk being sanitized, a file is created where the progress information will be written.

Listing the sanitized disks: The /etc/sanitized_disks file contains the serial numbers of all drives that have been successfully sanitized. For every invocation of the `disk sanitize start` command, the serial numbers of the newly sanitized disks are appended to the file.

A read of the /etc/sanitized_disks file shows output similar to the following:

```
admin1> rdfile /etc/sanitized_disks
```

```
Tue Jun 24 02:54:11 Disk 8a.44 [S/N 3FP0RFAZ00002218446B]
sanitized.
Tue Jun 24 02:54:15 Disk 8a.43 [S/N 3FP20XX400007313LSA8]
sanitized.
Tue Jun 24 02:54:20 Disk 8a.45 [S/N 3FP0RJMR0000221844GP]
sanitized.
Tue Jun 24 03:22:41 Disk 8a.32 [S/N 43208987] sanitized.
```

Reviewing the disk sanitization progress: A progress file is created for each drive sanitized and the results are consolidated to the /etc/sanitization.log file every 15 minutes during the sanitization operation. Entries in the log resemble the following:

```
Tue Jun 24 02:40:10 Disk sanitization initiated on drive 8a.43 [S/N
3FP20XX400007313LSA8]
Tue Jun 24 02:40:10 Disk sanitization initiated on drive 8a.44 [S/N
3FP0RFAZ00002218446B]
Tue Jun 24 02:40:10 Disk sanitization initiated on drive 8a.45 [S/N
3FP0RJMR0000221844GP]
Tue Jun 24 02:53:55 Disk 8a.44 [S/N 3FP0RFAZ00002218446B] format
completed in 00:13:45.
Tue Jun 24 02:53:59 Disk 8a.43 [S/N 3FP20XX400007313LSA8] format
completed in 00:13:49.
Tue Jun 24 02:54:04 Disk 8a.45 [S/N 3FP0RJMR0000221844GP] format
completed in 00:13:54.
Tue Jun 24 02:54:11 Disk 8a.44 [S/N 3FP0RFAZ00002218446B] cycle 1
pattern write of 0x47 completed in 00:00:16.
Tue Jun 24 02:54:11 Disk sanitization on drive 8a.44 [S/N
3FP0RFAZ00002218446B] completed.
Tue Jun 24 02:54:15 Disk 8a.43 [S/N 3FP20XX400007313LSA8] cycle 1
pattern write of 0x47 completed in 00:00:16.
Tue Jun 24 02:54:15 Disk sanitization on drive 8a.43 [S/N
3FP20XX400007313LSA8] completed.
Tue Jun 24 02:54:20 Disk 8a.45 [S/N 3FP0RJMR0000221844GP] cycle 1
pattern write of 0x47 completed in 00:00:16.
Tue Jun 24 02:54:20 Disk sanitization on drive 8a.45 [S/N
3FP0RJMR0000221844GP] completed.
Tue Jun 24 02:58:42 Disk sanitization initiated on drive 8a.43 [S/N
3FP20XX400007313LSA8]
Tue Jun 24 03:00:09 Disk sanitization initiated on drive 8a.32 [S/N
43208987]
Tue Jun 24 03:11:25 Disk 8a.32 [S/N 43208987] cycle 1 pattern write
of 0x47 completed in 00:11:16.
Tue Jun 24 03:12:32 Disk 8a.43 [S/N 3FP20XX400007313LSA8]
sanitization aborted by user.
Tue Jun 24 03:22:41 Disk 8a.32 [S/N 43208987] cycle 2 pattern write
of 0x47 completed in 00:11:16.
```

```
Tue Jun 24 03:22:41 Disk sanitization on drive 8a.32 [S/N 43208987]
completed.
```

Disk ownership assignment on an FAS270/FAS270c appliance

About FAS270/FAS270c disk ownership

The Network Appliance FAS270 and FAS270c appliances consist of a single disk shelf of 14 disks and either one internal filer head (on the FAS270) or two clustered internal filer heads (on the FAS270c). By design, a disk located on this common disk shelf can, if the appliance has two filer heads, be assigned to the ownership of either one of the two filer heads. The ownership of each disk is ascertained by an ownership record written on each disk.

Network Appliance delivers the FAS270 and FAS270c appliances with each disk preassigned to the one FAS270 internal filer head or preassigned to either one of the two FAS270c filer heads.

However, if you add one or more disk shelves to an existing FAS270 or FAS270c appliance, you might have to assign ownership to the disks contained on those shelves.

You can use the `disk show` command and the `disk assign` command to display, assign, and modify disk ownership on the FAS270 and the FAS270c.

Note

The `disk show` command and the `disk assign` command are valid only on the FAS270 and FAS270c appliances.

Viewing the disk ownership

Use the `disk show` command plus options to view the ownership of disks on the FAS270 or the FAS270c.

Step	Action
1	From the console of one of the FAS270 or FAS270c filer heads, enter the following command to display a listing of all disks on the FAS270 or the FAS270c and their owners. `fh1>disk show -v`

Sample output: The following sample output of the `disk show -v` command on an FAS270c shows disks 0b.16 through 0b.29 assigned in odd/even fashion to the internal cluster nodes fh1and fh2. The 14 disks on the add-on disk shelf are still unassigned to either filer head.

```
fh1>disk show -v
 DISK      OWNER            POOL   SERIAL NUMBER
 --------- ---------------  -----  -------------
 0b.43     Not Owned        NONE   41229013
 0b.42     Not Owned        NONE   41229012
 0b.41     Not Owned        NONE   41229011
 0b.40     Not Owned        NONE   41229010
 0b.39     Not Owned        NONE   41229009
 0b.38     Not Owned        NONE   41229008
 0b.37     Not Owned        NONE   41229007
 0b.36     Not Owned        NONE   41229006
 0b.35     Not Owned        NONE   41229005
 0b.34     Not Owned        NONE   41229004
 0b.33     Not Owned        NONE   41229003
 0b.32     Not Owned        NONE   41229002
 0b.31     Not Owned        NONE   41229001
 0b.30     Not Owned        NONE   41229000
 0b.29     fh1 (84165672)   Pool0  41226818
 0b.28     fh2 (84165664)   Pool0  41221622
 0b.27     fh1 (84165672)   Pool0  41226333
 0b.26     fh2 (84165664)   Pool0  41225544
 0b.25     fh1 (84165672)   Pool0  41221700
 0b.24     fh2 (84165664)   Pool0  41224003
 0b.23     fh1 (84165672)   Pool0  41227932
 0b.22     fh2 (84165664)   Pool0  41224591
 0b.21     fh1 (84165672)   Pool0  41226623
 0b.20     fh2 (84165664)   Pool0  41221819
 0b.19     fh1 (84165672)   Pool0  41227336
 0b.18     fh2 (84165664)   Pool0  41225345
 0b.17     fh1 (84165672)   Pool0  41225446
 0b.16     fh2 (84165664)   Pool0  41201783
```

Additional disk show options: The additional disk show options are listed below.

disk show parameter	Information displayed
`disk show -a`	Shows all disks that have an owner
`disk show [-o` *`hostname`* `| -s` *`nvram_id`*`]`	Shows all disks owned by a specified host or its NVRAM ID (the NVRAM number shipped with the filer head)
`disk show -n`	Shows all the disks that do not have an owner

disk show parameter	Information displayed
`disk show -v`	Shows all the visible disks

Assigning disks on the FAS270 or the FAS270c

To assign disks that are currently labeled “not owned,” complete the following steps.

Step	Action
1	In the CLI of the FAS270 or FAS270c filer head to which you want to assign disks, use the `disk show -n` command to view all disks that do not have assigned owners.
2	Use the following command to assign the disks that are labeled “not owned” to one of the filer heads. **`disk assign {`*`disk1`*` [`*`disk2`*`] [...]|all}`** *disk1* [*disk2*] [...] are the names of the disks to which you want to assign ownership. `all` is the option to assign all of the unowned disks to the current filer head. **Example:** The following command assigns six disks on the FAS270c to the filer fh1: `fh1>` **`disk assign 0b.43 0b.41 0b.39 0b.37 0b.35 0b.33`** **Result:** The specified disks are assigned as disks to the specified filer.
3	Use the `disk show -v` command to verify the disk assignments that you have just made.

After you have assigned disks to the one internal FAS270 filer head or to one of the two internal FAS270c filer heads, you can assign those disks to volumes on the filer that owns them or leave them as spare disks on the filer that owns them.

Modifying disk assignments on the FAS270c appliance

You can also use the `disk assign` command to modify the ownership of any disk assignment that you have made. This includes, on the FAS270c, reassigning a disk from one filer head to the other or, on either the FAS270 or FAS270c appliance, changing an assigned disk back to "not owned" status.

Caution

Network Appliance recommends that disk assignments be modified only for spare disks. Disks that have already been assigned to a volume cannot be reassigned without endangering all the data and the structure of that entire volume.

To modify disk ownership assignments on the FAS270 or the FAS270c, complete the following steps.

Step	Action
1	In the CLI of the FAS270 or FAS270c filer head whose disks you want to reassign, use the `vol status -r` command to view the spare disks, whose ownership can safely be changed.
2	Use the following command to modify assignment of the spare disks. **`disk assign disk1 [disk2] [...] -f {-o ownername | -s unowned | -s sysid}`** *disk1* [*disk2*] [...] are the names of the spare disks whose ownership assignment you want to modify. -f is a switch required for changing the assignment of disks that have already been assigned ownership. -o *ownername* specifies the host name of the internal filer head to which you want to reassign the disks in question. -s unowned modifies the ownership assignment of the disks in question back to "not owned." -s *sysid* is the factory-assigned NVRAM number of the internal filer head to which you want to reassign the disks in question. **Example:** The following command unassigns four disks on the FAS270c from the filer fh1: `fh1>`**`disk assign 0b.30 0b.29 0b.28 0b.27 -s unowned`**

Step	Action
3	Use the `disk show -v` command to verify the disk assignment modifications that you have just made.

User Information Maintenance 6

About this chapter

This chapter discusses how you can use Data ONTAP configuration files and Network Information Service (NIS) to maintain user information.

Topics in this chapter

This chapter discusses the following topics:

- "About maintaining user information" on page 254
- "Using configuration files to maintain user information" on page 255
- "Using NIS to maintain user information" on page 266
- "Changing the user name search order" on page 268
- "Using LDAP to maintain user information" on page 270

About maintaining user information

Reason for maintaining user information

User information is maintained by Data ONTAP on the filer to identify users. Identification of users is necessary to determine if users are allowed to access files through file access protocols such as NFS, CIFS, and FTP.

Ways to maintain user information

User information can be maintained in one or more of the following ways:

- In configuration files in the /etc directory on the filer's default volume

  For information about configuration files, see "Using configuration files to maintain user information" on page 255.
- On a Network Information Service (NIS) server

  For information about maintaining user information on an NIS server, see "Using NIS to maintain user information" on page 266.
- On an LDAP-based server

  For information about maintaining user information on an LDAP server, see "Using LDAP to maintain user information" on page 270.

Search order for user information

If you use a combination of the configuration files and an NIS or LDAP server, the order in which the files and the NIS or LDAP server are used is specified by the /etc/nsswitch.conf file. For detailed information about this file, see "Changing the user name search order" on page 268.

Using configuration files to maintain user information

About this section

This section provides information about the purpose and format of the configuration files, how to access these files from a Windows or UNIX client, and how to edit them.

Detailed information

For detailed information about using configuration files to maintain user information, see the following topics:

- "Understanding configuration files" on page 256
- "Accessing the configuration files" on page 259
- "Editing the configuration files" on page 261

Understanding configuration files

Configuration files available

Data ONTAP stores user information in the following four configuration files:

- /etc/passwd
- /etc/group
- /etc/netgroup
- /etc/shadow

Location of configuration files

When Data ONTAP is first installed on the filer, the passwd and group files are automatically created in the /etc directory. If you want to use the shadow and netgroup files, you need to create them in the /etc directory.

Purpose of the passwd and group files

Data ONTAP uses passwd and group files for the following purposes:

- To authenticate (PC)NFS access
- To translate user and group names in the /etc/quotas file to UIDs and GIDs
- To authenticate CIFS users.

  If, during setup, you configure Data ONTAP to use the passwd file to authenticate CIFS users, Data ONTAP uses the encrypted passwords in the passwd file to verify the user login names and passwords.
- To obtain the UNIX credentials of CIFS users

  If, during setup, you configure the filer to be a member of a domain or a Windows NT workgroup, Data ONTAP uses the domain controller or local user accounts to authenticate users. Data ONTAP, however, uses the passwd file for obtaining the UNIX credentials of CIFS users. For more information about UNIX credentials, see the section on how CIFS users obtain UNIX credentials in the CIFS chapter of the *File Access Management Guide*.

Purpose of the netgroup file

Data ONTAP uses the netgroup file to check access permission for remote mounting requests. For more information about the netgroup file, see the na_netgroup(5) man page.

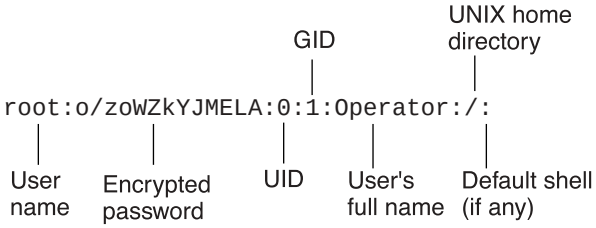
Purpose of the shadow file

Data ONTAP uses the shadow file to securely store the user passwords that would otherwise be in a passwd file.

Generally, the passwd file contains the user names and encrypted passwords of users who are authorized to access a filer. However, sometimes the passwd file needs to be readable by everyone, thus making the password information vulnerable. To overcome this vulnerability, you can use a shadow file. The password field in the passwd file is left blank if you use a shadow file. The shadow file contains encrypted passwords for all the users in the passwd file. For more information about the shadow file, see the na_shadow(5) man page.

Default entry in the passwd file

The following illustration shows the default entry in the passwd file.



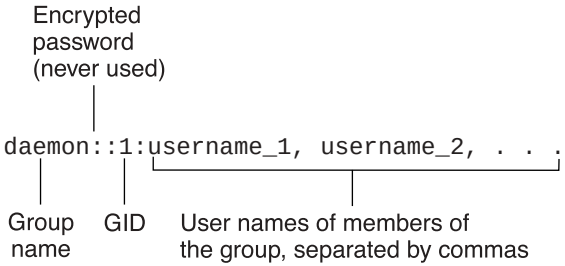
The following table describes the fields in an entry.

Field	Description
User name	Contains the alphanumeric user ID.
Encrypted password	Contains a password for the user. Data ONTAP disregards this field if NIS or a Windows NT domain controller is used to authenticate user names and passwords. This field is blank if you are using a shadow file.
UID	Contains a unique integer that is the UID.
GID	Contains a unique integer that is the user's primary GID.
User's full name	Contains the user's full name. Data ONTAP ignores this field.
UNIX home directory	Contains the path of the user's home directory. Data ONTAP ignores this field.

Field	Description
Default shell	Contains the path of the user's default shell. Data ONTAP ignores this field.

Default entry in the group file

The following illustration shows the default entry in the group file.



Field	Description
Group name	Contains a unique group name.
Encrypted password	Contains a password. Data ONTAP ignores this field.
GID	Contains a unique integer that is the GID.
User names	Contains a comma-separated list of user names that belong to the group.

Accessing the configuration files

No access from filer console

You cannot access the contents of the /etc directory from the console of your filer. You must use an NFS or a CIFS client to access the configuration files.

Accessing the files from a Windows workstation

To access the configuration files from a Windows PC, complete the following steps.

Step	Action
1	Map a drive to the root of the default filer volume using the following Universal Naming Convention (UNC) path: `\\`*`filername`*`\C$` *filername* is the name of your filer. If Windows prompts you for a user name and a password, ◆ Use Administrator as the user name. ◆ If you set up Data ONTAP to use a Windows NT domain controller to authenticate users, use the domain's administrator password. If you set up Data ONTAP to use the passwd file to authenticate users, use the administrator password that you entered when you ran `setup`.
2	Change to the /etc directory.

Accessing the files from a UNIX/NFS client

To access the configuration files from a UNIX/NFS client, complete the following steps.

Step	Action
1	Mount the root of the default filer volume to the client's file system using the following path: ***filername*:/vol/vol0** *filername* is the name of your filer.
2	Change to the /etc directory.

Editing the configuration files

About this section

This section describes how to edit the passwd and group files. For information about the netgroup and shadow files, see the na_netgroup(5) and na_shadow(5) man pages.

Editing restrictions

You cannot edit the contents of configuration files from the console of your filer because Data ONTAP does not have a text editor. You need to edit the configuration files from a UNIX/NFS or Windows/CIFS client.

When editing the configuration files, you must follow the specific format required for each file for each entry.

Procedures for editing the passwd file

The procedure you use to add entries to the /etc/passwd file depends on whether you are using the /etc/passwd file or either a Windows NT domain controller, NIS, or an LDAP server to authenticate users:

- If you are using the /etc/passwd file to authenticate users, follow the instructions in the next section, "Adding entries to the passwd file with passwords."
- If you are using either a Windows NT domain controller, NIS, or an LDAP server to authenticate users, follow the instructions in "Editing entries in the passwd file without passwords" on page 263.

Adding entries to the passwd file with passwords

In this procedure, you add entries to the passwd file using a text editor on a UNIX host, and then transfer the file to the filer. You add entries to the passwd file on a UNIX host because each entry in the file requires an encrypted password and Data ONTAP does not have a command for creating passwords. The format for entries in the passwd file is as follows:

```
username:password:uid:gid:full name:/:default shell
```

To add entries to the passwd file, complete the following steps.

Step	Action	
1	Log in to a UNIX host as the root user.	
2	Change to the /etc directory.	
3	Make a copy of the existing passwd file under a new name, such as passwd.old. **Note** It is important that you make a copy of the passwd file before editing it so that you can restore the original passwd file on this UNIX host after you have created a file that you will use on the filer.	
4	**If you have a passwd file on your filer that...**	**Then...**
	Contains only the default entry	Go to Step 5.
	Contains entries in addition to the default	Copy this file from the filer to the UNIX host in the /etc directory. This action overwrites the passwd file that exists in the /etc directory on the UNIX host.
5	Use a text editor to open the passwd file on the UNIX host.	
6	Add entries to the passwd file using the format shown in "Adding entries to the passwd file with passwords" on page 261. Leave the password and default shell fields empty. You set passwords in Step 8. Data ONTAP does not use the default shell field. **Note** Data ONTAP requires that the passwd file be terminated with a carriage return. Insert a carriage return after the last entry if none exists.	
7	Save the file and exit the text editor.	

Step	Action
8	Set a password for each user by entering the following command after the UNIX host command prompt: **`passwd`** ***`username`*** *username* is the user name of the user you added in Step 6. The `passwd` program prompts you to enter the new password twice.
9	Copy the passwd file from the UNIX host to the /etc directory on the filer.
10	Delete the passwd file from the UNIX host and give the copied file (passwd.old) the new name, passwd.

Deleting entries in the passwd file with passwords

In this procedure, you delete from the passwd file, the entries of users to whom you want to deny access.

To delete entries from the passwd file, complete the following steps.

Step	Action
1	Use a text editor to open the /etc/passwd file on the filer.
2	Delete the entry for each user to whom you want to deny access.
3	Save the file and exit the text editor.

Editing entries in the passwd file without passwords

In this procedure, you edit the passwd file from a Windows PC or a UNIX/NFS client, using a text editor. You do not need to encrypt a password, because a domain controller or NIS authenticates users. The format for entries in the passwd file is as follows:

```
username:password:gid:uid:full name:default shell
```

To edit entries in the passwd file, complete the following steps.

Step	Action
1	Use a text editor to open the /etc/passwd file on the filer.

Step	Action
2	Add an entry for each user who needs to access files. Leave the password and default shell fields empty. Data ONTAP does not use them when users are authenticated by a Windows NT domain controller, NIS, or an LDAP server. **Note** Data ONTAP requires that the passwd file be terminated with a carriage return. Insert a carriage return after the last entry if none exists.
3	Save the file and exit the text editor.

Editing the group file

In this procedure, you edit the group file from either a Windows workstation or a UNIX/NFS client using a text editor. The format for entries in the group file is as follows:

```
groupname::gid:username,...,username
```

To edit the group file, complete the following steps.

Step	Action	
1	Make sure that you are connected to the root filer volume.	
2	Use a text editor to open the /etc/group file.	
3	**If...**	**Then...**
	You want to add a group	Add an entry for the group for which you want to authorize access using the format shown above. **Note** Data ONTAP requires that the group file be terminated with a carriage return. Insert a carriage return after the last entry if none exists.
	You want to delete a group	Delete the entry for the group.

Step	Action
4	Save the file and exit the text editor. **Result:** The passwd and group files on the filer now contain entries for users and groups. Data ONTAP can verify users' rights to access the filer.

Note

There is no upper limit to the number of entries in the /etc/group file. Any one user, however, can belong to only 32 groups.

Using NIS to maintain user information

Prerequisites

Before you can use NIS on your filer, you must

- Configure the filer to use NIS
- Bind the filer to an NIS server

For more information about configuring and binding the filer to use NIS, see the *Data ONTAP Network Administration Guide*.

Additionally, you should set up caching of group information. Using cached group information enables Data ONTAP to quickly determine CIFS users' secondary GIDs when users try to connect to the filer. Without the cached information, Data ONTAP has to access the NIS server for group information each time a CIFS user requests a connection.

Advantage of using NIS

NIS enables you to centrally maintain user information. As a result, you do not have to maintain the configuration files for each filer on your network. If you have several filers on your network, maintaining user information centrally saves you from updating these files on each filer every time you add or delete a user or a group.

Caching NIS group information

To cache NIS group information, complete the following steps.

Step	Action
1	Enter the following command: **`options nis.group_update_schedule`** ***`hours_list`*** *hours_list* can be a comma-separated list of numbers that range from 1 to 24, representing the hours of the day, or `now`, which means that the cache is updated immediately.
2	To make Data ONTAP use the cached group information, enter the following command: **`options nis.group_update.enable on`** The default is not to use cached group information.

Example of caching NIS group information: The following command caches NIS group information immediately and again at 2 a.m. and 2 p.m.:

```
options nis.group_update_schedule 2, 14, now
```

Displaying cached group entries

To list the group entries that are cached locally by Data ONTAP, complete the following step.

Step	Action
1	Enter the following command: **`ypgroup [`*`username`*`]`** If you specify a user name, group entries for that name only will be listed. Use the `ypgroup` command without the user name to list group entries for every user.

Updating cached NIS group information

Anytime you update the group file on the NIS master server, you should also update the cached group information on the filer. If you do not, there is a period of time when the information on the NIS master server and the filer is out of sync.

To update the cached NIS group information on the filer, complete the following step.

Step	Action
1	After updating the NIS master server, enter the following command: **`options nis.group_update_schedule now`** **OR** Add the following command to the script you use to push the NIS master server update to the NIS slave servers: **`options nis.group_update_schedule now`**

Changing the user name search order

How search order is specified

Data ONTAP uses the /etc/nsswitch.conf file to specify the user name search order.

If the nsswitch.conf file does not already exist in the /etc directory, Data ONTAP creates a default nsswitch.conf file when NIS or LDAP authentication is enabled on your filer. The contents of the default nsswitch.conf file are as follows:

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

Note

All entries in the nswitch.conf file except hosts pertain to user information maintenance. For information about the hosts map, see the *Data ONTAP Network Administration Guide*.

The default nsswitch.conf file searches for user information in the following order:

- Configuration files in the /etc directory of the filer's default volume
- NIS

Note

If you want to authenticate users through an LDAP database, add ldap to the passwd entry of the /etc/nsswitch.conf file. For details, see "Editing the /etc/nsswitch.conf file for LDAP" on page 273.

Changing the search order

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

Step	Action
1	Use a text editor to open the /etc/nsswitch.conf file.

Step	Action
2	Edit the file appropriately. Each line in the file must be in the following format: `map: service ...` *map* is one of the following: passwd, netgroup, group, shadow. *service* is one or more of the following: files, nis.
3	Save the file and exit the text editor.

Example

The following example shows an nsswitch.conf file. According to the search order specified in this file, Data ONTAP searches for user information (user name, password, and groups) using only NIS. However, Data ONTAP also looks up information in the /etc/netgroup file for the network groups if the NIS lookup fails.

```
passwd: nis
group: nis
netgroup: nis files
shadow: nis
```

Using LDAP to maintain user information

About using LDAP

An LDAP server enables you to centrally maintain user information. As a result, you do not have to maintain the configuration files for each filer on your network. If you have several filers on your network, maintaining user information centrally saves you from updating these files on each filer every time you add or delete a user or a group.

If you store your UNIX user database on an LDAP server, you can configure your filer to look up user information in the LDAP database rather than using the /etc/password file or NIS.

About configuring LDAP

You can configure the filer to use one or more LDAP servers by setting Data ONTAP options at the command line.

To configure the filer for LDAP authentication, do the following:

- Specify the LDAP base distinguished name using the `ldap.base` option.
- Specify the LDAP servers to use for LDAP queries using the `ldap.servers` option.
- Specify the preferred LDAP servers, if any, using the `ldap.servers.preferred` option.
- Enable LDAP on the filer using the `ldap.enable` option.
- Edit the password line of the /etc/nsswitch.conf file for LDAP.

In addition, you might need to do the following:

- If anonymous authentication does not work in your environment, specify the administrative user name to use for queries for looking up user IDs (UIDs) and group IDs (GIDs) by using the `ldap.name` option.
- If there is a password for the administrative user, specify the administrative password using the `ldap.passwd` option.
- If your LDAP server has been set up to use a port for LDAP queries other than the default port 389, specify the LDAP port using the `ldap.port` option.

Types of LDAP servers supported

Data ONTAP supports the following types of LDAP servers:

- Netscape®
- iPlanet
- OpenLDAP

Extending the RFC 2307 schema

RFC 2307 schema must be extended on the LDAP servers that you want to use for LDAP queries.

For more information refer to RFC 2307 or to documentation by third party directory integration vendors.

Specifying the LDAP base in the distinguished name format

The LDAP base is the name of the LDAP tree in which user information is stored.

To specify which LDAP base distinguished name to use for looking up user names, complete the following step.

Step	Action
1	Enter the following command: **`options ldap.base`** ***`name`*** *name* is the base distinguished name. Use quotes around names with embedded spaces. **Example:** `options ldap.base “o=networkappliance,c=us”`

Specifying LDAP servers

To specify the LDAP servers to be used for LDAP queries, complete the following step.

Step	Action
1	Enter the following command: **`options ldap.servers "name[,name...]"`** *name* is the name of an LDAP server. You can enter multiple server names using a space-separated list enclosed in quotes. **Example:** `options ldap.servers "server1 server2"`

Specifying preferred servers

You might want to specify LDAP servers that are on faster links as the preferred servers.

To specify the LDAP servers that you want the filer to check first, complete the following step.

Step	Action
1	Enter the following command: **`options ldap.servers.preferred "name[,name...]"`** *name* is the name of a preferred LDAP server. You can enter multiple server names using a space-separated list enclosed in quotes. **Example:** `options ldap.servers.preferred "server1,server2"`

Enabling or disabling LDAP

To enable or disable LDAP on your filer, complete the following step.

Step	Action
1	Enter the following command: **`options ldap.enable {on|off}`** Use `on` to enable and `off` to disable LDAP.

Editing the /etc/nsswitch.conf file for LDAP

To edit the /etc/nsswitch.conf file for LDAP, complete the following steps.

Step	Action
1	Open the /etc/nsswitch.conf file on the filer for editing. **Note** For instructions on accessing files in the /etc directory, see “Accessing the configuration files” on page 259.
2	Enter the following at the password line: `ldap` You can optionally add `files` and/or `nis` to the password line, but they must be entered after `ldap` if you want to use LDAP as the primary mechanism to retrieve user information. **Example:** `passwd: ldap files nis`
3	Save the file.

Specifying the admin user name

If anonymous authentication does not work in your environment, you need to specify an administrative user name.

To specify the user name to be used for administrative queries for looking up UIDs and GIDs, complete the following step.

Step	Action
1	Enter the following command: **`options ldap.name`** ***`name`*** *name* is the name to be used for administrative queries. The best practice is for this to be the name of a user with read-only access to the LDAP database. Use quotes around names with embedded spaces. **Example:** `options ldap.name “cn=root,o=networkappliance,c=us”`

Specifying the admin password

To specify a password for the administrative user, complete the following step.

Step	Action
1	Enter the following command: **`options ldap.passwd password`** *password* is the password for the administrative user. The password is displayed as a series of asterisks.

Specifying the LDAP port

You might need to specify the port to use for LDAP queries if the LDAP server has been set up to use a port other than the default for LDAP, port 389.

To specify the LDAP port, complete the following step.

Step	Action
1	Enter the following command: **`options ldap.port N`** *N* is the LDAP port number.

Note

Data ONTAP does not support using LDAP with the Secure Socket Layer (SSL) or Transport Layer Security (TLS) protocol, which is commonly called Secure LDAP or LDAPS. Configuring Data ONTAP to connect to an LDAPS server on port 636 does not provide secure LDAP access for filer queries and, in most cases, will cause all LDAP queries from the filer to fail.

Supporting LDAP authentication of UNIX clients

To support authentication of UNIX clients through an LDAP server make sure that ldap is the first protocol entered on the password line of the /etc/nsswitch.conf file as described in “Editing the /etc/nsswitch.conf file for LDAP” on page 273.

Supporting LDAP authentication of Windows clients

You can authenticate Windows clients through an LDAP server. To support authentication of Windows clients through an LDAP server, complete the following additional operations.

- Run CIFS Setup on the filer to be accessed, and specify NIS/LDAP as the authentication method to be used for CIFS clients on that filer.
- Configure the local security settings of each Windows client to use clear text password (unencrypted password) authentication rather than use Kerberos or other encrypted authentication methods.
- Verify that your Windows clients have their userpassword attribute configured in the LDAP user database.

Supporting LDAP authorization of Windows client access to UNIX files

To support authorization of Windows client access to UNIX files on a filer that uses LDAP authentication, complete the following tasks:

- On the filer to be accessed, verify that the /etc/nsswitch.conf file specifies ldap as one of the passwd entries. See “Editing the /etc/nsswitch.conf file for LDAP” on page 273.
- On the filer to be accessed, verify that every CIFS user who needs to access UNIX files is mapped to an associated UNIX user name in the usermap.cfg file.
- Verify that every associated UNIX username has an entry in the LDAP database.

Supporting LDAP authorization of UNIX client access to an NTFS or mixed file system

To support authorization of UNIX client access to an NTFS or mixed file system on a filer that uses LDAP authentication, complete the following tasks:

- On the filer to be accessed, verify that the /etc/nsswitch.conf file specifies ldap as one of the passwd entries. See “Editing the /etc/nsswitch.conf file for LDAP” on page 273.
- Verify that every UNIX user that needs to access an NTFS or mixed file system has an entry in the LDAP database.
- On the filer to be accessed, verify that every UNIX user that needs to access an NTFS or mixed file system is mapped to an associated CIFS user name in the usermap.cfg file.

Cluster Management 7

About this chapter

This chapter discusses what clusters are and how to manage them.

Topics in this chapter

This chapter discusses the following topics:

- "About clusters" on page 278
- "How a clustered system works" on page 281
- "Cluster operation modes" on page 282
- "Understanding interface configurations for takeover" on page 291
- "Configuring dedicated and standby interfaces" on page 294
- "Configuring mirrored, dual-MAC interfaces" on page 297
- "Enabling cluster failover in a DAFS configuration" on page 300
- "Managing a cluster" on page 307
- "Checking cluster configuration" on page 332

About clusters

What a cluster is

A cluster is a pair of NetApp filer heads (or nodes) that are connected through an interconnect adapter and cables and are configured so that both filer heads share access to a set of Fibre Channel disks, subnets, and tape drives to provide fault tolerance. In most cases the clustered nodes are two separate filers sharing access to the same disk shelves

Note

NetApp also offers combined internal filer head and disk shelf appliances (such as the FAS270c) in which the cluster consists of two internal filer heads that reside in and share the same disk shelf.

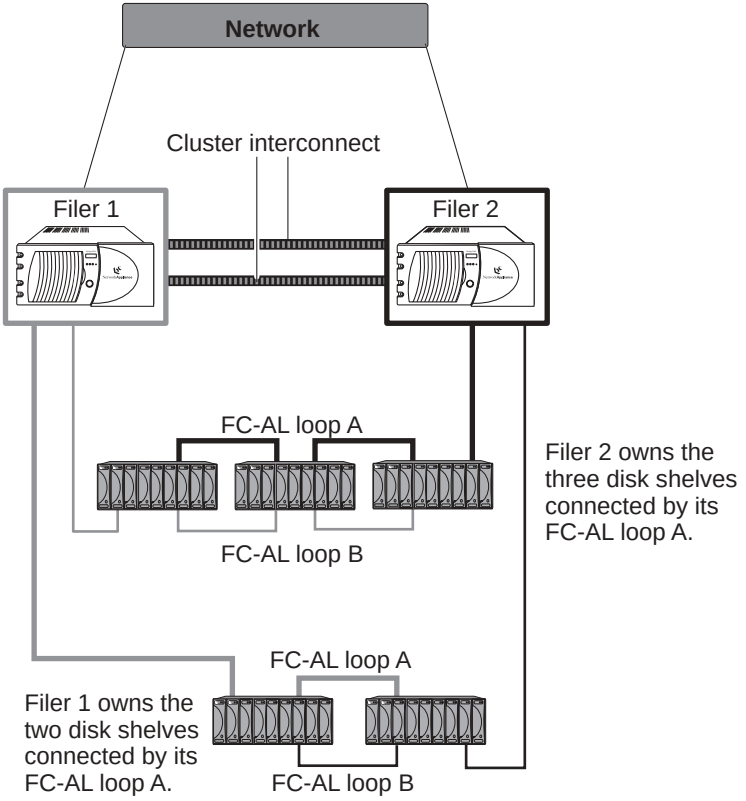
What a cluster does

When one filer in a cluster undergoes a system failure and cannot reboot, the partner filer in the cluster can take over the failed filer functions and serve the data on the failed filer disks. This is known as a takeover.

In addition to enabling automatic takeover, you can also force a takeover at any time. You can cause a clustered filer to immediately halt and be taken over by its partner, allowing scheduled maintenance to be performed on the filer. Alternatively, you can cause a filer to halt while preventing its partner from taking over, allowing scheduled maintenance on both the filer and disks.

Cluster diagram

The following figure diagrams a common configuration of the cluster hardware components.



Cluster characteristics

The filers in a cluster have the following characteristics:

- They are connected to each other through a cluster interconnect consisting of adapters and cable through which they
  - Continually check whether the other filer is functioning
  - Mirror log data for each other's NVRAM
  - Synchronize each other's time
- They use two Fibre Channel Arbitrated Loops (FC-AL loops), in which
  - 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.
- They may reside on the same NT domain or on different domains.

Cluster considerations

The following considerations are necessary to have a successful cluster environment:

- Architecture compatibility

  Both filers should have roughly the same performance and architecture so that no undue performance degradation occurs during takeover.

  The number of Fibre Channel disks shared by the two filers must not exceed the capacity of either filer.
- Disks and disk shelf compatibility

  You can use only FC-AL disks and disk shelves in clustered configurations.

Preparing to configure a cluster

In addition to the cluster considerations described earlier, you must also ensure the following conditions before configuring a cluster:

- System hardware is compatible between filers.
- Disks are connected through Fibre Channel to both filers.
- Interconnect host adapters and cables are installed.
- Filers are attached to the same networks.
- Network card interfaces are configured correctly.
- Cluster software is licensed on both filers.
- Cluster software is enabled on both filers.
- The same system software, such as CIFS, NFS, or SyncMirror, is licensed and enabled on both filers, except in the case of vFiler, which does not have to be licensed on the takeover filer for takeover of the partner filer to succeed.

Note

Filers in a cluster can reside on the same NT domain or on different domains.

How a clustered system works

What happens during a system failure

If a filer is in a cluster that is not configured for immediate takeover on panic and undergoes a system failure but recovers and reboots, the filer behaves exactly like a nonclustered filer.

If a filer is in a cluster that is configured for immediate takeover on panic and undergoes a system failure, or if that filer undergoes a system failure and cannot reboot, the partner filer in the cluster performs an automatic takeover of the failed filer functions and disk drives by creating an emulated filer that

- Assumes the identity of the failed filer
- Accesses the failed filer disks and serves its data to clients

The partner filer maintains its own identity and its own primary functions but also handles the added functionality of the failed filer through the emulated filer.

What happens after takeover

When a takeover occurs, you view the surviving partner as having two identities, its own and its partner's, that exist simultaneously on the same filer. Each identity can access only the appropriate volumes and networks. You can send commands or log in to either filer by using the `rsh` command, allowing remote scripts that invoke filer commands through a Remote Shell connection to continue to operate normally.

Access with rsh: Commands sent to the failed filer through a Remote Shell connection are serviced by the partner filer, as are `rsh` command login requests.

Access with Telnet: If you log in to a failed filer through a Telnet session, you see a message alerting you that your filer failed and to log in to the partner filer instead. If you are logged in to the partner filer, you can access the failed filer or its resources from the partner filer by using the `partner` command.

What happens during a giveback

When the failed filer is functioning again,

- You initiate a `giveback` command that terminates the emulated 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

Cluster operation modes

Types of operation modes

There are three types of operation modes in a clustered environment:

- Normal mode is the condition of the cluster when there is no takeover or giveback in effect.
- Takeover mode is when one filer has taken over its partner.
- Giveback mode is the condition of a cluster as it goes through the process of returning to normal mode.

Detailed information

The following sections discuss the three cluster operation modes:

- "Normal operation mode" on page 283
- "Takeover operation mode" on page 285
- "Giveback operation mode" on page 289

Normal operation mode

Normal mode activities

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 through the cluster interconnect adapters and cables 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.

Action	Description
Monitoring	The cluster monitor software monitors 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 negotiated failover facility determines that one partner should take over the other because of a disk shelf failure. ◆ The partner power-on-self-test (POST) fails. ◆ A panic occurs on the partner, and the `cf.takeover.on_panic` option is set to on. **Note** If the filer that is not transmitting a heartbeat is rebooting, no takeover occurs. Takeover might be disabled if NVRAM mirroring is not functioning. Automatic takeover is disabled if the partner mailbox disks are not accessible.
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	Filer times can deviate from each other. A daemon updates the times on filers periodically to keep them synchronized.

Takeover operation mode

Takeover process

The takeover process is shown in the following table.

Stage	Description
1	One of these four things happens: ◆ An operator issues the `halt` command without the `-f` option. ◆ One filer in a cluster undergoes a hardware or power failure and cannot reboot. That filer is the failed filer, and the partner 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. ◆ One filer in a cluster experiences a disk shelf failure and the negotiated failover function determines that the partner filer should take over for the filer with the failed shelf. **Note** Negotiated failover occurs only once after one or both of the filers in the cluster are booted, establishing communication between the two filers. The two filers exchange disk shelf information and determine whether negotiated failover should occur. The next time a negotiated failover decision can be made is when one or both filers are rebooted, causing the filers to reestablish communication.
2	The takeover filer reserves all disk drives in the cluster, so that they do not receive commands from the failed filer.
3	The takeover filer creates an emulated filer that takes over the resources and the identity of the failed filer and executes the equivalent of a reboot of the emulated filer. For example, it uses network interface information from the failed filer configuration files to duplicate the failed filer network services.

Stage	Description
4	The emulated filer goes through a process similar to the reboot process: ◆ All filer information preserved across filer reboot is preserved in a takeover. ◆ The filer information that is lost in a reboot, such as the 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 on the console of the takeover filer. In the following example, the `cf takeover` command was issued:

```
filer1> Tue Jul 21 23:02:02 GMT [filer1:
cf.misc.operatorTakeover:warning]: Cluster monitor: takeover
initiated by operator
Tue Jul 21 23:02:11 GMT [filer1:
cf.fsm.nfo.acceptTakeoverReq:warning]: Negotiated failover:
accepting takeover request by partner, reason: Operator initiated
cf takeover. Asking partner to shutdown gracefully; will takeover
in at most 180 seconds.
Tue Jul 21 23:02:20 GMT [filer1:
cf.fsm.partnerNotResponding:notice]: Cluster monitor: partner not
responding
```

Then the following messages appear:

```
Tue Jul 21 23:03:06 GMT [filer1:
cf.fsm.nfo.partnerShutdown:warning]: Negotiated failover: partner
has shutdown
Tue Jul 21 23:03:06 GMT [filer1: cf.fsm.stateTransit:warning]:
Cluster monitor: UP --> TAKEOVER
Tue Jul 21 23:03:06 GMT [filer1: cf.fm.takeoverStarted:warning]:
Cluster monitor: takeover started
Tue Jul 21 23:03:10 GMT [filer1: disk_admin:warning]: Resetting all
devices on ISP2100 in slot 7.
```

Reboot messages follow as the emulated filer boots up. Then the following message appears and indicates a successful takeover:

```
Tue Jul 21 23:04:22 GMT [filer1 (takeover):
cf.fm.takeoverComplete:warning]: Cluster monitor: takeover
completed
filer1(takeover)>
```

Note

Negotiated failover messages are not printed when you issue a `cf takeover -f` command, a `cf forcetakeover` command, or a `cf forcetakeover -f` command.

Messages indicating a successful negotiated takeover

During negotiated takeover, the cluster monitor issues messages on the takeover filer similar to the following:

```
filer2> Thu May  4 01:41:15 GMT [ha4:
cf.fsm.nfo.acceptTakeoverReq:warning]: Negotiated failover:
accepting takeover request by partner, reason: partner sees fewer
shelves on its A loop (4) than we do (5).  Asking partner to
shutdown gracefully; will takeover in at most 180 seconds.
```

Additional messages indicate that the takeover filer is waiting for the failed filer to respond. When the failed filer shuts down, the partner filer begins the takeover process:

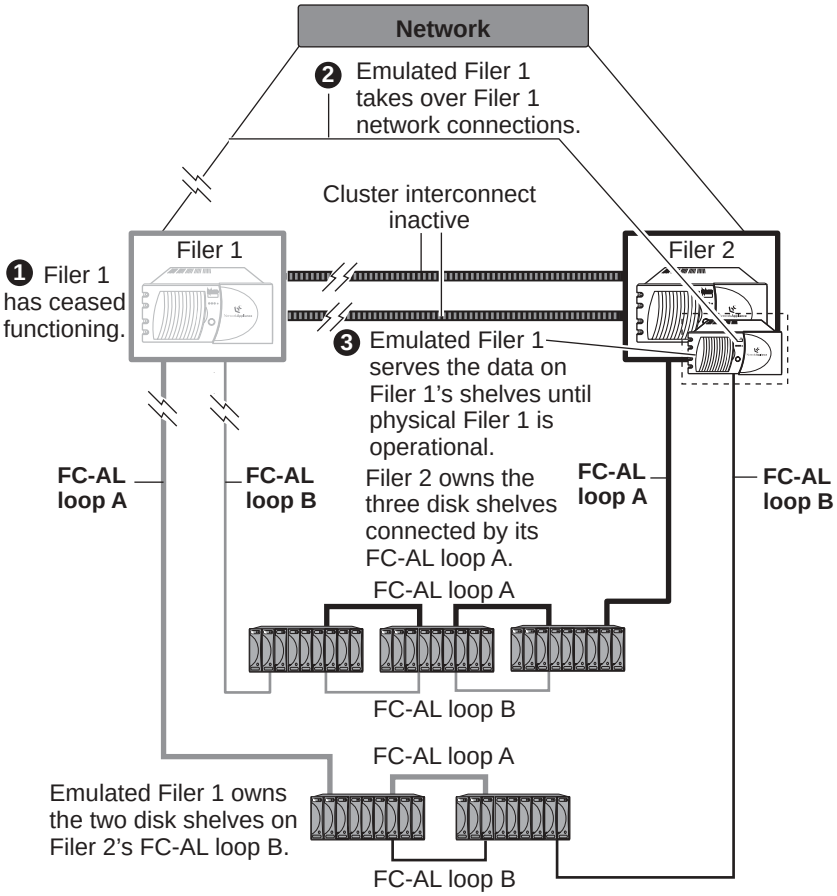
```
Thu May  4 01:41:32 GMT [filer2: cf.fsm.stateTransit:warning]:
Cluster monitor: UP --> TAKEOVER
Thu May  4 01:41:32 GMT [filer2: cf.fm.takeoverStarted:warning]:
Cluster monitor: takeover started
Thu May  4 01:41:36 GMT [filer2: disk_admin:error]: Resetting all
devices on ISP2100 in slot 7.
...
```

The following message indicates a successful negotiated takeover:

```
Thu May  4 01:41:47 GMT [filer2 (takeover):
cf.fm.takeoverComplete:warning]: Cluster monitor: takeover
completed
filer2(takeover)>
```

Takeover result diagram

The following figure shows the state of a cluster after a takeover of Filer 1 by Filer 2.



Giveback operation mode

Giveback process

The giveback process is described in the following table.

Stage	Description
1	The giveback process is initiated in one of two ways: ◆ An operator issues a `cf giveback` command. ❖ If the option `cf.giveback.check.partner` is enabled on the takeover filer, the takeover filer checks whether its partner is in "waiting for giveback" mode before actually executing the `cf giveback` command. ❖ If the `cf giveback -f` command is issued, then the giveback operation will start whether the partner is in "waiting for giveback" mode or not. ◆ If the option `cf.giveback.auto.enable` is enabled on the takeover filer, the takeover filer automatically executes giveback when the state of its partner goes to "waiting for giveback." No manual `cf giveback` command is necessary.
2	The takeover filer updates the failed filer disks.
3	The takeover filer destroys the emulated filer so that no more writes can occur to the failed filer disks.
4	The takeover filer releases the disk reservations that it held on behalf of the failed filer. This action returns control to the failed filer of the disks it owns.
5	The failed filer resumes normal operations, serving its own data. The users of that filer experience the effects of a reboot.

Caution

Avoid initiating the giveback process on the takeover filer before its partner, the failed filer, is ready to resume operations. If the `options cf.giveback.check.partner` command is not enabled or not supported on the takeover server, the giveback process might proceed on the takeover filer, whether or not its partner is in a "waiting for giveback" mode.

Messages indicating a successful giveback

During a typical giveback, the cluster monitor issues messages on the console of the takeover filer:

```
filer1(takeover)> Tue Jul 21 22:37:34 GMT [filer1 (takeover):
cf.misc.operatorGiveback:info]: Cluster monitor: giveback initiated
by operator
Tue Jul 21 22:37:34 GMT [filer1: cf.fm.givebackStarted:warning]:
Cluster monitor: giveback started
Tue Jul 21 22:37:35 GMT [filer1: cf.fm.givebackComplete:warning]:
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 22:37:35 GMT [filer1: cf.fsm.stateTransit:warning]:
Cluster monitor: TAKEOVER --> UP
Tue Jul 21 22:38:35 GMT [filer1: cf.fsm.firmwareStatus:info]:
Cluster monitor: partner booting
Tue Jul 21 22:38:45 GMT [filer1: cf.fsm.firmwareStatus:info]:
Cluster monitor: partner kernel loaded
Tue Jul 21 22:39:35 GMT [filer1: cf.fsm.partnerOk:notice]: Cluster
monitor: partner ok
Tue Jul 21 22:39:35 GMT [filer1: cf.fsm.takeoverDisabled:notice]:
Cluster monitor: takeover of partner disabled (unsynchronized log)
Tue Jul 21 22:39:41 GMT [filer1: cf.fsm.takeoverEnabled:notice]:
Cluster monitor: takeover of partner enabled
Tue Jul 21 22:39:55 GMT [filer1: monitor.globalStatus.ok:info]: The
system's global status is normal
```

The last message indicates that the filers are operating normally and, therefore, either filer can take over for its partner if necessary.

Understanding interface configurations for takeover

About 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 by using the `setup` command to configure interfaces for a new filer or the `ifconfig` command to add interfaces to an existing filer.

Role of network interfaces in a cluster

The role of network interfaces in a cluster is to ensure that each filer in a cluster can assume the identity of its partner on the networks to which the filers are connected:

- Each filer in a cluster must have interface attachments to the same collection of subnets.
- Each filer must have interfaces that can take over each IP address or interface on the partner filer.

In a noncluster configuration, each interface acts in one role: supporting the identity of the local filer.

In a cluster configuration, each interface can act in one of two roles:

- Normal mode, which is the state of the filer when no takeover is being performed in the cluster
- Takeover mode, in which the partner filer has taken over the failed filer

The following table shows the role of each interface type in normal and takeover mode.

Interface type	Normal mode	Takeover mode
Dedicated	Supports the identity of the local filer	Supports the identity of the local filer
Standby	Idle	Supports the identity of the failed filer
Shared	Supports the identity of the local filer	Supports the identity of both the local filer and the failed filer

Takeover configurations

You can have two kinds of interface configurations for takeover:

- A configuration using separate mirrored network interfaces. You use this configuration with cards that support only single Media Access Control (MAC) addresses. For each interface, you configure a dedicated interface on one filer and a standby interface on the other.
- A configuration using NICs with dual MAC capability. The interfaces in a dual-MAC pair, one interface per filer, mirror each other.

FORE/IP and LANE interface configurations

You can configure FORE/IP and LAN Emulation (LANE) interfaces as follows:

- You must configure FORE/IP interfaces as dedicated and standby interfaces.
- For LANE, you can configure ATM adapter interfaces as mirrored interface pairs, provided that after takeover, a given Emulated LAN is configured only once on an adapter. Otherwise, you must use dedicated and standby interfaces.
- You can specify FORE/IP partners by using IP addresses or the FORE/IP interface name.
- You can specify the LANE partner for an ATM adapter by using the ATM adapter interface name.

Interface types and configurations

The following table lists the configurations supported by each type of interface in a cluster.

Interface	Shared	Dedicated	Standby	Partner parameter
Ethernet	X	X	X	IP address or interface name
Gigabit Ethernet		X	X	IP address or interface name
Gigabit Ethernet II	X	X	X	IP address or interface name
FDDI	X	X	X	IP address or interface name
Virtual interface	X	X	X	Virtual interface name

Interface	Shared	Dedicated	Standby	Partner parameter
FORE/IP		X	X	IP address or FORE/IP interface name
ATM adapter for LANE[1]	X[2]	X	X	ATM adapter interface name

[1]The OC-12 adapter does not support clusters on F700-series filers.
[2]ATM adapters in a cluster can be shared for LANE provided that the two adapters do not have Emulated LANs in common.

Configuring dedicated and standby interfaces

Prerequisites for configuring interfaces

To configure a pair of interfaces so that the standby interface can replace a dedicated interface on the partner filer in a cluster, you must have the following items available:

- Four single MAC-address interfaces or ATM adapter interfaces, two on each filer in the cluster
- Two IP addresses or ATM interface names, the address of one interface or ATM interface name on each filer in the cluster

Rules for configuring interfaces

Keep these rules in mind when configuring interfaces:

- You can use interface names for specifying all interfaces.
- You cannot use IP addresses when specifying ATM interfaces for LANE and virtual interfaces.

Configuring interfaces for takeover

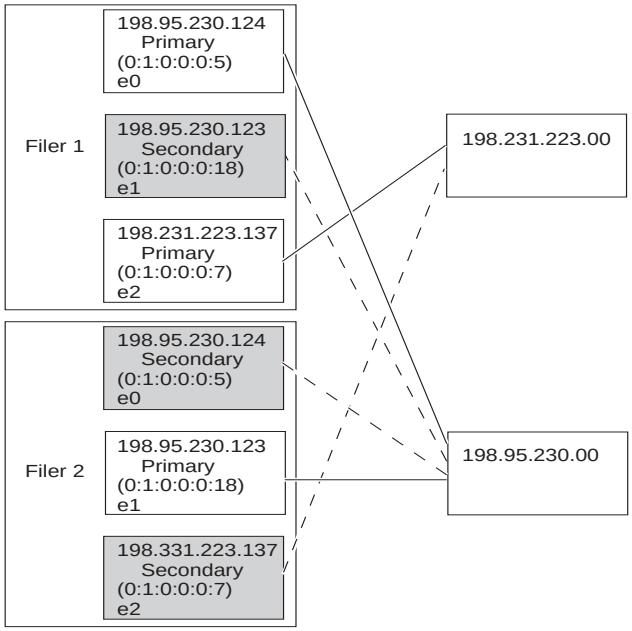
To configure a pair of interfaces for takeover, complete the following steps.

Step	Action
1	Enter the following command on the command line and also enter it in the /etc/rc file of the filer with the dedicated local interface, so that the command is permanent: **`ifconfig`** ***`interface address`*** *interface* is the name of the dedicated local interface. *address* is the IP address or ATM interface name of the dedicated local interface.

Step	Action
2	Enter the following command on the command line and in the /etc/rc file of the filer you used in Step 1 of this procedure: **`ifconfig interface partner address`** *interface* is the interface name of the standby interface. *address* is the IP address or ATM interface name of the dedicated interface of the other filer in the cluster. **Note** You cannot use IP addresses when specifying ATM interfaces for LANE and virtual interfaces.
3	Enter the following command on the command line and in the /etc/rc file of the other filer in the cluster: **`ifconfig interface address`** *interface* is the interface name of the dedicated local interface. *address* is the IP address of the dedicated local interface.
4	Enter the following command on the command line and in the /etc/rc file of the other filer in the cluster: **`ifconfig interface partner address`** *interface* is the name of the standby interface. *address* is the IP address or ATM interface name of the dedicated interface in Step 1. **Note** You cannot use IP addresses when specifying ATM interfaces for LANE and virtual interfaces. For detailed information on configuring ATM interfaces, see the ATM configuration chapter in the *Network Management Guide*.

Takeover configuration with dedicated and hot standby interfaces

In the following configuration illustration, you use two network cards for each interface. Each card can use only one MAC address. One card acts as a dedicated interface and the other acts as a hot standby interface.



If Filer 1 fails, the following occurs on takeover:

- Interface e0 on Filer 1 stops functioning, but e0 on Filer 2 substitutes for the failed interface and handles the Filer 1 network connection with the 230 network.
- Interface e2 on Filer 1 stops functioning, but e2 on Filer 2 substitutes for the failed interface and handles the Filer 1 network connection with the 223 network.

If Filer 2 fails, e1 on Filer 2 stops functioning, but e1 on Filer 1 substitutes for the failed interface and handles the Filer 2 network connection with the 230 network.

Configuring mirrored, dual-MAC interfaces

Prerequisites for configuring dual-MAC interfaces

To configure a pair of failover, 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 must have available the following items:

- Two dual-MAC interfaces, one installed on each filer in the cluster
- Two IP addresses, the address of one interface on each filer in the cluster

Emulated LAN considerations

For LANE, when an adapter is shared, an Emulated LAN 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 Emulated LAN, each filer needs a minimum of two ATM cards to support full takeover.

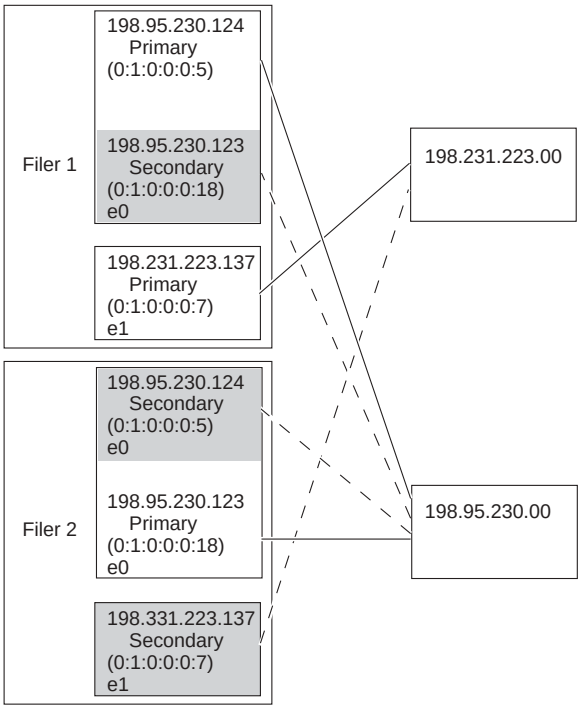
Configuring dual-MAC interfaces

To configure a pair of dual-MAC interfaces for takeover, complete the following step.

Step	Action
1	Enter the following command on the command line of both filers in a cluster and in the /etc/rc file of both filers so that the command is permanent: **`ifconfig`** ***`interface address`*** **`partner`** ***`address`*** *interface* is the name of the local interface. The first *address* is the IP address of the local interface. The second *address* is the IP address or interface name that the local interface services when the local filer has taken over the other filer in the cluster. **Note** When you configure virtual interfaces for takeover, you must specify the interface name and not the IP address.

Takeover configuration with shared interfaces

In the following configuration illustration, you use two network cards that support two MAC addresses to mirror two interfaces, as shown with the interfaces that connect to the 230 net.



With the configuration shown in the preceding figure, if Filer 1 fails, the following occurs on takeover:

- Interface e0 on Filer 1 stops functioning, but the secondary address on e0 on Filer 2 handles the Filer 1 network connection with the 230 network.
- Interface e1 on Filer 1 stops functioning, but e1 on Filer 2 substitutes for the failed interface and handles the Filer 1 network connection with the 223 network.

If Filer 2 fails, e0 on Filer 2 stops functioning; e0 on Filer 1 substitutes for the failed interface, however, and handles the Filer 2 network connection with the 230 network.

Emulated LAN example

The following command specifies that adapter a4 will take over all the Emulated LANs on the partner a5 interface.

```
ifconfig a4 partner a5
```

Note

If a4 and partner a5 are configured with common Emulated LANs, a4 does not take over those Emulated LANs.

A cluster pair might be configured with Emulated LANs as follows.

Filer and adapter	Interfaces and Emulated LANS	Filer and adapter	Interfaces and Emulated LANS
filer1-a4	el0 default el1 nineKMTU	filer1-a5	el2 eighteenKMTU
filer2-a4	el0 default el1 nineKMTU	filer2-a5	el2 eighteenKMTU

To achieve full takeover, you would use the first command on one filer and the second command on the second filer:

```
ifconfig a4 partner a5
ifconfig a5 partner a4
```

Virtual interface example

The following is an example of a virtual interface configuration command:

```
ifconfig vif0 partner vif0
```

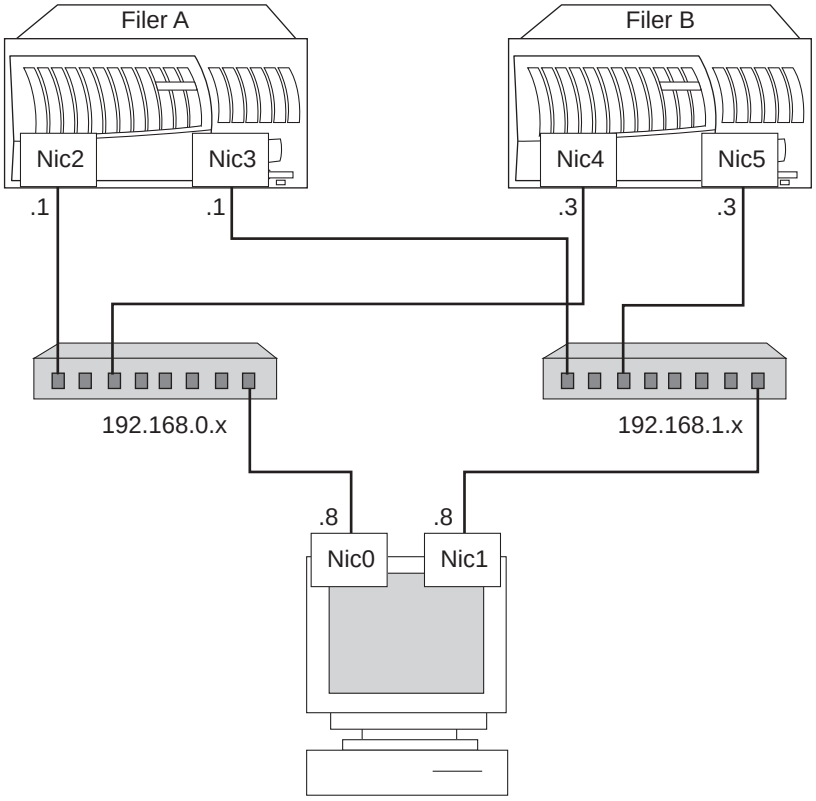
Enabling cluster failover in a DAFS configuration

About the failover path

When a filer in a cluster fails, the DAFS Database Accelerator client must find the failover paths over the VI network to the takeover filer. Each VI NIC on the client must connect to a VI NIC on the takeover filer. The client uses the IP addresses of VI NICs on the takeover filer to identify the failover path. The IP addresses of the VI NICs are listed in the client's /etc/hosts file.

Configuration example

The following figure shows a sample configuration in which a client with two VI NICs connects to a clustered pair of filers through switches.



In this configuration, each VI NIC on a filer has a corresponding failover NIC on the takeover filer. When a filer fails, the following happens:

- If Filer A fails, Filer B is the takeover filer. NIC4 on Filer B is the failover path for NIC2. NIC5 on Filer B is the failover path for NIC3.
- If Filer B fails, Filer A is the takeover filer. NIC2 on Filer A is the failover path for NIC4. NIC3 is the failover path for NIC5.

Each VI NIC on the client uses the failover path to reach the takeover filer. For example, if Filer A fails, then the following happens:

- NIC0 on the client connects to Filer B via NIC4.
- NIC1 on the client connects to Filer B via NIC5.

Enabling failover paths

To enable failover paths between the DAFS Database Accelerator client and the takeover filer, you must configure the following:

- On the client, define the IP addresses of failover VI NICs in the client's /etc/hosts file.
- On the filer, use the `dafs attach` command to define a failover VI NIC as a partner interface.

Detailed information

For detailed information, see the following sections:

- "Designating failover addresses in etc/hosts" on page 302
- "Enabling failover paths on the filer" on page 305

Designating failover addresses in etc/hosts

Failover paths in the etc/hosts file

The /etc/hosts file lists the following information for failover paths:

- IP addresses of the VI NICs on each filer
- The failover host names and addresses for each filer

Failover names for VI NICs

In the /etc/hosts file, the `-failover` designation is included in the host name of a VI NIC to indicate that it is a failover path.

Example: `filerA-failover-viip0`

Sample etc/hosts file

The following example shows the client's /etc/hosts entries for the configuration shown in "Configuration example" on page 300.

```
# VIP_ADDRESS

192.168.0.1 filerA-viip0

192.168.1.1 filerA-viip1

192.168.0.3 filerB-viip0

192.168.1.3 filerB-viip1

# Failover addresses

192.168.0.3 filerA-failover-viip0

192.168.1.3 filerA-failover-viip1

192.168.0.1 filerB-failover-viip0

192.168.1.1 filerB-failover-viip1
```

Explanation of sample /etc/hosts file

The `# VIP_ADDRESS` entries list the local IP addresses and symbolic names of the VI NICs installed on each filer in the clustered pair. The `# Failover addresses` entries list the failover address for each local `VIP_ADDRESS`.

In this sample /etc/hosts file, each NIC IP address corresponds to a local host name and a failover host name. For example, the NIC IP address 192.168.0.1 on filer A has the following host names:

- `filerA-viip0`—The local NIC host name
- `filerB-failover-viip0`—A failover NIC host name that indicates that Filer A is the takeover filer when Filer B fails

How the client identifies the failover path

If Filer A fails, the client performs a lookup for a host name with the failover designator. The lookup returns the following addresses:

- `192.168.0.3 filerA-failover-viip0`
- `192.168.1.3 filerA-failover-viip1`

These addresses correspond to the following NICs on filerB:

- `filerB-viip0`
- `filerB-viip1`

Editing the etc/hosts file

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

Step	Action

1	Use a text editor such as `vi` to open the client's /etc/hosts file.
	Example:
	`# cd /etc`
	`# vi hosts`

Step	Action
2	For each VI NIC, add a line in the following format: ***ip_address hostname-viipN*** *ip_address* is the IP address of the VI NIC. *hostname* is the filer's local host name. `viip` is VI over IP. *N* is the number of the NIC (0, 1, 2, 3). **Example:** `192.168.0.1 filerA-viip0`
3	Add lines in the following format to designate the failover addresses: ***ip_address hostname*`-failover-viip`*N*** *ip_address* is the IP address of the VI NIC. *hostname*`-failover-viip`*N* is the symbolic name of the failover VI NIC plus the failover designator. ◆ *hostname* is the filer's host name. ◆ `viip` is VI over IP. ◆ *N* is the number of the NIC (0, 1, 2, 3). **Example:** `1 92.168.0.3 filerA-failover-viip0`

Enabling failover paths on the filer

Using the dafs attach command

After you configure failover addresses in the /etc/hosts file on the DAFS client, you must use the `dafs attach` command on the filer to define a failover VI NIC as a partner interface. You can also create partner interfaces while using the `dafs setup` command.

For more information about the `dafs setup` command and the `dafs attach` command, see the DAFS chapter of the *File Access Management Guide*.

To enable failover paths on the filer, complete the following steps.

Step	Action

1	On the filer command line, enter the following command:
	`dafs attach` *vinic* **`-d`** *physical_device* **`-p`**
	vinic is the logical partner interface name of the VI NIC to be attached to the DAFS service on the filer. Use the e*N* naming convention, where *N* is the number of the NIC.
	`-d` *physical_device* is the physical device name of the VI NIC. If the physical device is not specified, then by default the physical device name is the same as the logical *vinic* name.
	`-p` indicates that the *vinic* is attached as a partner interface. The partner devices are attached but remain in the inactive partner state until a CFO takeover occurs.
	Example:
	`filer> dafs attach pe4 -d e4 -p`
	`filer>`

Step	Action

2	After attaching the partner devices, you can verify the status of all known VI NICs by entering the following command on the filer command line:
	`dafs status -n`
	Result: The display shows the state of the DAFS service and a list of all VI NICs that use the DAFS service and their current states.
	Example:
	`filer> dafs status -n`
	`DAFS service is started`
	`NIC: e4  Device: e4  State: ATTACHED`
	`NIC: pe4  Device: e4  State: INACTIVE PARTNER`
	`NIC: e5  Device: e5  State: ATTACHED`
	In the examples, the partner devices pvf5 and pe4 are attached, but remain in the inactive partner state until a CFO takeover occurs.

Managing a cluster

About managing a cluster

You need to know how to manage a cluster in normal mode, takeover mode, and giveback mode. Managing a cluster entails enabling and disabling takeover, initiating takeover, managing an emulated filer, and initiating a giveback.

Detailed information

The following topics describe the tasks involved in managing a cluster:

- "Managing a cluster in normal mode" on page 308
- "Enabling negotiated failover between filers" on page 316
- "Managing a cluster in takeover mode" on page 318
- "Managing an emulated filer" on page 320
- "Backing up and restoring from the emulated filer" on page 325
- "Performing a giveback" on page 326
- "Enabling automatic giveback" on page 329
- "Configuring automatic giveback" on page 330

Managing a cluster in normal mode

Normal mode management tasks

You manage a cluster in normal mode by performing the following actions:

- Monitoring the status of the cluster as a whole
- Viewing information about the cluster
- Enabling and disabling takeover on a cluster to perform software upgrades or other maintenance
- Halting a filer in a cluster without causing a takeover
- Performing license operations on the cluster feature
- Performing a takeover on the partner filer
- Enabling and disabling the negotiated failover feature on a cluster

Monitoring cluster status

To find out whether the cluster feature is enabled and whether the other filer in the cluster is up, complete the following step.

Step	Action

1	Enter the following command:
	`cf status [-t]`
	`-t` displays the status of the node as time master or slave.
	Example:
	`filer1> cf status`
	`Cluster enabled, filer2 is up.`
	This verifies the link between the clustered filers and 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.
- Data ONTAP cannot access the cluster mailbox disks.

Description of cluster status messages

The following table shows the messages that the `cf status` command can display.

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 a cluster interconnect adapter.
`Interconnect is down.`	The cluster interconnect adapter cannot access the partner. This might be due to cabling problems or the partner might be down.
`Interconnect is up.`	The cluster interconnect adapter is active and can transmit data to the partner.
`X link is down; Y link is up.`	The ServerNet X link (upper cable) is currently down. This condition might 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 ServerNet Y link (lower cable) is currently down. This condition might occur briefly when a cluster node is booting. If this condition persists, check the cable seating.
`partner_name_1 has detected a mailbox disk error, take over of partner_name_2 disabled.`	One filer in the cluster is unable to access multiple mailbox disks. Check access to both the local and partner root volumes and mirrors, if they exist. Also check for disk or FC-AL problems or offline storage adapters.

Message	Meaning
partner_name_2 may be down and has disabled take over by *partner_name_1*.	One filer in the cluster may be down. Check its state and reboot if necessary.
Version mismatch	The partner node has an incompatible version of Data ONTAP.
partner_name_1 is attempting takeover of *partner_name_2*. giveback is in module *n* of *N* modules.	A takeover is being attempted (includes information on how far the takeover has progressed).
partner_name_1 has taken over *partner_name_2*	A takeover has occurred.
partner_name_1 has taken over *partner_name_2,* giveback in progress. giveback is in module *n* of *N* modules.	A giveback is being attempted (includes information on how far the giveback has progressed).
partner_name_1 has taken over *partner_name_2, partner_name_2* is ready for giveback.	The takeover filer has received information that the failed filer is ready for giveback.
partner_name_1 has taken over *partner_name_2, partner_name_2* is ready for giveback. Automatic giveback is disabled due to exceeding retry count.	The takeover filer has received information that the failed filer is ready for giveback, but giveback cannot take place because the number of retries has exceeded the limit.

Note

Even if one of the ServerNet interconnect links (X or Y) is down, cluster communication continues over the link that is up.

Displaying the partner's name

To find out the name of the other filer in the cluster, complete the following step.

Step	Action
1	Enter the following command: `cf partner`

Cluster configuration information

In a cluster, the `sysconfig` and `vol status` command outputs are special in that they display information about both filers.

For each filer, the `sysconfig` command output displays 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.

Displaying disk information on a cluster

To display disk information for both filers in a cluster, complete the following step.

Step	Action

1	Enter one of the following commands:
	`sysconfig -r`
	or
	`vol status -r`

Enabling and disabling takeover

You might want to disable takeover if you are doing maintenance that typically causes a takeover and then reenable takeover when you have finished maintenance.

To enable or disable takeover on a cluster, complete the following step.

Step	Action
1	Enter the following command: **`cf enable|disable`** Use `cf enable` to enable takeover or `cf disable` to disable takeover. You can enable or disable takeover by using the command from either filer in the cluster.

Enabling and disabling immediate takeover of a panicked partner

You can enable a clustered filer to take over immediately when its partner panics. If the panicked filer needs to do a core dump, and a spare disk is available, the core dump is directed to the spare disk and the partner will take over immediately (within 60 seconds) without having to wait until the core dump is completed.

Note

Until the core dump to the spare disk is complete, the spare disk will not be allocated for use as a file system disk.

If the panicked filer needs to do a core dump and no spare disk is available, the filer does not perform a core dump; instead, it prints out a message that a spare disk is needed. When the filer reboots, it logs the fact that no core dump took place.

To enable or disable immediate takeover of a panicked partner, complete the following step.

Note

For immediate takeover to occur, you must enable cluster takeover by using the `cf enable` command.

Step	Action
1	Enter the following command: `options cf.takeover.on_panic [ on | off ]` Use `on` to enable immediate takeover of a failed partner or `off` to disable immediate takeover. If this option is disabled, normal takeover procedures will occur. **Note** By default this option is turned off, except when the iSCSI and FCP licenses have been previously enabled. If this command is entered on one filer in a cluster, the value applies to both filers in the cluster. The setting of this option is persistent across reboots.

Halting a filer without takeover

You can halt the filer and prevent its partner from taking over. For example, you might need to perform maintenance on both the filer and its disks and want to avoid an attempt by the partner filer to write to those disks.

To halt a filer in a cluster without causing a takeover, complete the following step.

Step	Action
1	Enter the following command: `halt -f`

Licensing and unlicensing the cluster feature

You can license or unlicense the cluster feature only in normal mode with takeover disabled.

To unlicense the cluster feature, complete the following steps.

Step	Action	
1	**If the filer has...**	**Then...**
	Taken over its partner	Enter the following command to perform a giveback: **`cf giveback`**
	Not taken over its partner	Go to Step 2.
2	Enter the following command to disable clustering: **`cf disable`**	
3	Enter the following command to unlicense clustering: **`license delete cluster`**	
4	Enter the following command to reboot the filer and make the change take effect: **`reboot`**	

Takeover commands and descriptions

You can initiate a takeover on a filer in a cluster to perform maintenance on that filer while still serving the data on its disks to users. The following table lists the commands you can use when initiating a takeover and a description of those commands.

Command	Description
`cf takeover`	Initiates a takeover of the partner of the local filer. Takeover is aborted if a core dump is in progress on the partner (only if the `cf.takeover.on_panic` option is set to false).
`cf takeover -f`	Initiates a takeover of the partner of the local filer regardless of whether the other filer is dumping its core (only if the `cf.takeover.on_panic` option is set to false).

Command	Description
`cf forcetakeover`	Tells the cluster monitor to ignore some configuration problems that would otherwise prevent a takeover, such as unsynchronized NVRAM due to a faulty cluster interconnect connection. It then initiates a takeover of the partner of the local filer. **Caution** Use `cf forcetakeover` only in certain controlled circumstances, such as when you cannot get `cf takeover` to succeed. You can lose uncommitted data when you use this command.
`cf forcetakeover -d`	Initiates a takeover of the local partner even in the absence of a quorum of partner mailbox disks. The `cf forcetakeover -d` command is valid only if the cluster_remote license has been enabled. **Caution** Use the `-d` option only after you have verified that the partner is down. **Note** The `-d` option can be used in conjunction with RAID mirroring to recover from disasters in which one partner in a cluster is not available. For more information, see the *Data Protection Guide*.

Enabling negotiated failover between filers

What negotiated failover is

Negotiated failover is a feature that enables each filer in a cluster to negotiate takeover from one filer to its partner in the event of a disk shelf failure.

How negotiated failover works

When communication between clustered filers is first established through the cluster interconnect adapters, the filers exchange a list of disk shelves that are visible on the A and B loops of each filer. If a discrepancy is found between the lists on each filer, the clustered filers perform the following negotiated failover process.

Stage	Description
1	If a filer sees that the B loop shelf count on its partner is greater than its local A loop shelf count, the local filer concludes that it is impaired because it sees fewer of its disk shelves than its partner sees.
2	The local filer asks its partner to initiate a takeover. If the partner filer is not itself impaired, the partner accepts the takeover request and asks the impaired filer to shut down gracefully.
3	The partner filer takes over after the impaired filer shuts down or after a timeout of approximately three minutes.

Enabling and disabling negotiated failover

To enable and disable negotiated failover on a cluster, complete the following step.

Step	Action
1	Enter the following command: **`cf nfo enable|disable *disk_shelf*`** Use `enable` to turn negotiated failover on, and use `disable` to turn negotiated failover off. Negotiated failover is disabled by default.

Specifying the time period before takeover

You can set an option to specify the number of seconds one partner in a cluster can be unresponsive before the other partner takes over.

Step	Action
1	Enter the following command: **`options cf.takeover.detection.seconds`** ***`number_of_seconds`*** The valid values for *number_of_seconds* are 5 through 180; the default is 20.

Both partners do not need to have the same value for this option. Thus, you can have one partner that takes over more quickly than the other.

Caution

If the specified time is less than 20 seconds, false takeover can occur. Use caution when assigning a takeover time of 20 seconds or less.

Determining negotiated failover status

To determine negotiated failover status on a cluster, complete the following step.

Step	Action
1	Enter the following command: **`cf nfo status`** **Results:** The `cf nfo status` command can display the following information: ◆ Whether negotiated failover is enabled or disabled ◆ Whether a takeover is imminent due to a negotiated failover ◆ Whether a takeover has occurred due to a negotiated failover, and the reason for the takeover

Managing a cluster in takeover mode

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

In takeover mode, statistics for some 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 emulated filer. Therefore, the display does not differentiate between the operations on the takeover filer and the operations on the emulated 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.

Statistics are treated differently in takeover mode than in normal mode for the following commands:

```
httpstat
ifstat
netstat
nfsstat
sysstat
```

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 do not distinguish between filers.

Managing an emulated filer

What an emulated filer is

An emulated filer is a software copy of the failed filer that is hosted by the takeover filer. You access the emulated filer in partner mode by using the `partner` command, as described in the following sections.

Management exceptions for emulated filers

You manage an emulated filer as you do any other filer, including managing disks, with the following exceptions, which are described in greater detail later in this section:

- An emulated filer can access only failed filer disks.
- Some commands are unavailable.
- Some displays differ from normal displays.

Accessing the emulated filer from the takeover filer

You access the emulated 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 emulated filer, and takeover mode, in which you manage the takeover filer.
- The `partner` command with a filer command as an argument executes the filer command on the emulated filer in partner mode and then returns to takeover mode.

Using the partner command without arguments: To toggle between partner mode and takeover mode, complete the following step.

Step	Action

1	From the partner prompt, enter the following command:
	`partner`
	Results: The prompt changes to the partner mode prompt, which has the following form:
	`emulated_filer/takeover_filer>`
	Example: 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)>`

Using the partner command with arguments: To use the `partner` command with a filer command as an argument, complete the following step.

Step	Action

1	From the partner prompt, enter the following command:
	`partner` ***`command`***
	command is the filer command you want to initiate on the emulated filer.
	Example: `filer1(takeover)> partner cf status`
	`filer2 has been taken over by filer1.`
	`filer1(takeover)>`

Accessing the emulated filer remotely

You can also access the emulated filer remotely without using the `partner` command.

Using Telnet: With Telnet, you can directly access only the takeover filer.

To access the emulated filer by using a Telnet session, complete the following steps.

Step	Action
1	Log in to the takeover filer.
2	Enter the following command: **`partner`**

Using a Remote Shell connection: To execute a command on the emulated filer using `rsh`, complete the following step.

Step	Action
1	Enter the following command: **`rsh failed_filer command`** *failed_filer* is the name of the failed filer. *command* is the filer command you want to run. **Example:** In the following example, filer2 is the failed filer. `rsh filer2 df`

Emulated filer command exceptions

Almost all the commands that are available to a normal filer are available on the emulated filer. Some commands, however, are either unavailable or behave differently in emulated filer mode. The commands in the following subsections are the exceptions.

Unavailable commands: The following commands are not available on an emulated 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 following table lists the commands that behave differently on an emulated filer than they do on a takeover filer and describes the behavior.

Command	Difference
`ifconfig [`*`interface`*`]`	Displays the following: ◆ Emulated filer interface mappings based on the failed filer's /etc/rc file rather than the takeover filer interface mappings ◆ Emulated filer interface names rather than the interface names of the takeover filer that has performed the takeover ◆ Only interfaces that have been configured, rather than all interfaces, configured or not, as displayed on the takeover filer
`mt`	Uses the tape devices on the takeover filer 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 its partner filer in the cluster.

Command	Difference
`sysconfig`	When it displays hardware information, the `sysconfig` command displays information only about the hardware that is attached to the takeover filer. It does not display information about the 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.
`vol status`	When it displays hardware information, the `vol status` command displays information only about the hardware that is attached to the takeover filer. It does not display information about the 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 emulated filer

Emulated filer tape drives inaccessible

You can perform dumps and restores on the emulated filer. However, neither the takeover filer nor the emulated filer can access the tape drives or other peripheral devices on the failed filer.

Prerequisites to performing dumps on an emulated 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 emulated filer are not scheduled to occur during the same time period; the filer and the emulated filer cannot access the tape drives simultaneously.

Performing dumps on an emulated filer

You can perform dumps on an emulated filer by using the `dump` command on the takeover filer in partner mode. For information about performing dumps on a filer, see the *Data Protection Guide*.

Performing restores on an emulated filer

To perform restores on the emulated filer, complete the following steps.

Step	Action
1	Put the restore tape into a drive accessible by the takeover filer.

2	Use the `restore` command in partner mode or as an argument to the `partner` command.
	Example using the restore command in partner mode:
	`filer1 (takeover)> partner`
	`filer1/filer2> restore [`*`options`* `[`*`arguments`*`]]`
	Example using the restore command as an argument to the partner command:
	`filer1 (takeover)> partner restore [`*`options`* `[`*`arguments`*`]]`

Performing a giveback

About performing a giveback

To restore a cluster to normal operation, complete one of the following giveback operations from the takeover filer.

Initiating normal giveback

If you want to initiate giveback to the failed filer, enter the following command on the command line of the takeover filer:

```
cf giveback
```

Note

After a giveback, the takeover filer's ability to take over its partner automatically is not reenabled until the partner reboots successfully. If the partner fails to reboot, you can enter the `cf takeover` command to initiate a takeover of the partner manually.

Processes that can prevent giveback: The `cf giveback` command does not succeed if any of the following events are taking place on the emulated filer at the same time:

- Dump and restore processes
- Outstanding CIFS sessions
- RAID disk additions
- RAID reconstruction, RAID disk scrubbing
- Volume creation
- Disk ownership assignment (carried out on an FAS270c)
- Disks being added to a volume (vol add)
- Snapshot creation, deletion, or renaming
- Quota initialization
- Advanced mode repair operations such as WAFLiron
- Filer panics
- Backup dump and restore operations
- SnapMirror transfers (if the partner is a SnapMirror destination)
- SnapVault restorations
- Disk sanitization operations

Note

If some of the listed events are in progress, you can terminate them by running the giveback process through the `cf giveback -f` command. See "Terminating takeover filer processes to initiate giveback" on page 327.

Terminating takeover filer processes to initiate giveback

Under some circumstances, you might want to terminate a process early in order to initiate a giveback operation.

To terminate a process early and initiate giveback to the failed filer, enter the following command:

```
cf giveback -f
```

Caution

If `cf giveback -f` is invoked and fails, you might need to reboot both cluster partners.

Forcing giveback

The takeover filer might detect an error condition on the failed filer that typically prevents a complete giveback, such as data not being flushed from NVRAM to the failed filer's disks.

To force the takeover filer to give back the resources of the failed filer even if the takeover filer detects an error that typically prevents a complete giveback, enter the following command:

```
cf forcegiveback
```

When the failed filer reboots as a result of a forced giveback, it displays the following message if data from the emulated filer did not fully synchronize with the failed 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 if you use this command.

What happens when a giveback is interrupted

If the takeover filer experiences a failure or a power outage during the giveback process, the giveback process stops and the takeover filer returns to takeover mode when the failure is repaired or the power is restored. Some operations, such as quota initialization, always cause a giveback to fail.

Enabling automatic giveback

What automatic giveback is

Automatic giveback begins when one node of a cluster is in takeover mode and the failed node is repaired and reboots. The repaired node boots into Data ONTAP, and the node in takeover mode detects this and initiates a giveback.

Note

This option works only with platforms booted from CompactFlash.

How to enable automatic giveback

To enable automatic giveback, complete the following step.

Step	Action
1	Enter the following command: **`options cf.giveback.auto enable off | on`** The `on` value enables automatic giveback. The `off` value disables automatic giveback. This option is Off by default.

Configuring automatic giveback

Shortening giveback time

You can shorten negotiated giveback time by enabling an option (`cf.giveback.check.partner`) that checks whether the failed filer is already recovered and waiting for giveback—that is, waiting for disk reservation release—when you type **`cf giveback`** at console of the takeover filer. If the failed filer is waiting for giveback, giveback proceeds as normal; if not, giveback is canceled.

Note

This option works only with platforms booted from CompactFlash.

How to enable shorter giveback time

To enable shorter giveback time, complete the following step.

Step	Action
1	Enter the following command: **`options cf.giveback.check.partner on | off`** The `on` value enables this option. The `off` value disables this option. This option is Off by default. This option is On by default on platforms booted from CompactFlash.

Setting takeover delay time for CIFS clients

This option specifies the number of minutes to delay an automatic giveback before terminating CIFS clients that have open files. During the delay, the system periodically sends notices to the affected workstations. If you specify 0, CIFS clients are terminated immediately.

When this option is Off, the automatic giveback is deferred until all CIFS files have been closed. This option is used only if automatic giveback is On.

How to set takeover delay time for CIFS clients

To set takeover delay time for CIFS clients, complete the following step.

Step	Action
1	Enter the following command: **`options cf.giveback.auto.cifs.terminate.minutes`** ***`minutes`*** Valid values for *minutes* are 0 to 999, inclusive. The default is 5 minutes.

Setting takeover to terminate long-running processes

This option, when on, specifies that automatic giveback should immediately terminate long-running operations (dump/restore, vol verify, and so on) when initiating an automatic giveback. When this option is off, the automatic giveback is deferred until the long-running operations have completed. This option is used only if automatic giveback is On.

How to set takeover to terminate long-running processes

To set takeover to terminate long-running processes, complete the following step.

Step	Action
1	Enter the following command: **`options cf.giveback.auto.terminate.bigjobs on | off`** The `on` argument enables this option. The `off` argument disables this option. This option is On by default.

Checking cluster configuration

About the Cluster Configuration Checker

Network Appliance provides a Cluster Configuration Checker utility that you can download and run to ensure against the most common configuration causes of failover problems between clustered filers.

Where you can run this utility

You can run the utility, `cf-config-check.cgi`, as a command from a UNIX shell, or you can install the CGI script on a UNIX Web server and invoke it from a Web browser.

rsh requirements

In order to run the Cluster Configuration Checker utility, you require rsh access to both clustered filers.

What the Cluster Configuration Checker finds

When run, the utility detects the following error and warning conditions in the clustered filers:

- Service license errors. If the services licensed on the two clustered filers are not identical, some services might become unavailable on failover.
- Options settings errors. If the options settings on the two clustered filers are not identical, some options will change on failover.
- Network interface errors. If the network interface settings are incorrect, client connections might be lost on failover.
- Data ONTAP version differences between the clustered filers.
- Clustering licensed but not enabled on one of the clustered filers.
- FCP is licensed on the cluster but the `fcp set cfmode` settings on the cluster partners are inconsistent.

Downloading and running the checker

To download and run the Cluster Configuration Checker, log in to the NOW Web site (http://now.netapp.com) and go to Software Downloads > Tools & Utilities. Click “Cluster Configuration Checker (cf-config-check.cgi).”

Follow the directions for downloading and running the utility.

Disk Space Management Using Quotas 8

About this chapter

This chapter describes how to restrict and track the disk space and number of files used by a user, group, or qtree.

Topics in this chapter

This chapter discusses the following topics:

- "Understanding quotas" on page 334
- "When quotas take effect" on page 337
- "Understanding default quotas" on page 338
- "Understanding derived quotas" on page 339
- "How Data ONTAP identifies users" on page 342
- "Understanding the /etc/quotas file" on page 345
- "Turning quotas on or off" on page 361
- "Turning quota message logging on or off" on page 363
- "Creating or resizing quotas" on page 365
- "Deleting quotas" on page 369
- "Effects of qtree changes on quotas" on page 370
- "Understanding quota reports" on page 372
- "Effects of hard quotas on clients" on page 382

Understanding quotas

Reasons for specifying quotas

You specify a quota for the following reasons:

- To limit the amount of disk space or the number of files that can be used by a quota target
- To track the amount of disk space or the number of files used by a quota target, without imposing a limit
- To warn users when their disk space or file usage is high

Quota targets

A quota target can be

- A user, as represented by a UNIX ID or a Windows ID.
- A group, as represented by a UNIX group name or GID. Data ONTAP does not apply group quotas based on Windows IDs.
- A qtree, as represented by the path name to the qtree.

The quota target determines the quota type. For example, a quota whose target is a user is called a user quota, and a quota whose target is a group is called a group quota.

Exception: A quota for a qtree is called a tree quota, not a qtree quota.

Tree quotas

If you apply a tree quota to a qtree, the qtree is similar to a disk partition, except that you can change its size at any time. When applying a tree quota, Data ONTAP limits the disk space and number of files regardless of the owner of the disk space or files in the qtree. No users, including the root and members of the BUILTIN\Administrators group, can write to the qtree if the write exceeds the tree quota.

Quota specifications

Quota specifications are stored in the /etc/quotas file, which you can edit at any time.

User and group quotas are applied on a per-volume or per-qtree basis. You cannot specify a user or group quota for the entire filer or for multiple volumes.

Example: You can specify that a user named jsmith can use up to 10 GB of disk space in the /vol/vol0 volume, or that a group named engineering can use up to 50 GB of disk space in the /vol/vol0/projects qtree.

Explicit quotas

If the quota specification references the name or ID of the quota target, the quota is an explicit quota. For example, if you specify a user name, jsmith, as the quota target, the quota is an explicit user quota. If you specify the path name of a qtree, /vol/vol0/engineering, as the quota target, the quota is an explicit tree quota.

Default quotas and derived quotas

The disk space used by a quota target can be restricted or tracked even if you do not specify an explicit quota for it in the /etc/quotas file. The target might be under the effect of a quota that is derived from a default quota. If a quota is applied to a target but the name or ID of the target does not appear in an /etc/quotas entry, the quota is called a derived quota. For more information about default quotas, see “Understanding default quotas” on page 338. For more information about derived quotas, see “Understanding derived quotas” on page 339.

Hard quotas, soft quotas, and threshold quotas

A hard quota is a limit that cannot be exceeded. If an operation, such as a write, causes a quota target to exceed a hard quota, the operation will fail. When this happens, a warning message is logged to the filer console and an SNMP trap is issued.

A soft quota is a limit that can be exceeded. When a soft quota is exceeded, a warning message is logged to the filer console and an SNMP trap is issued. When the soft quota returns to normal, another syslog message and SNMP trap are generated. In Data ONTAP, you can specify both hard and soft quota limits for the amount of disk space used (in blocks) and the number of files created.

A threshold quota is similar to a soft quota restricted to specifying quotas in blocks. When a threshold quota is exceeded, a warning message is logged to the filer console and an SNMP trap is issued.

A single type of SNMP trap is generated for all types of quota events. You can find details on SNMP traps in the filer’s /etc/mib/netapp.mib file.

Syslog messages about quotas contain qtree ID numbers rather than qtree names. You can correlate qtree names to the qtree ID numbers in syslog messages by using the `qtree status -i` command. For more information, see “Displaying qtree status” on page 147.

Notification about threshold crossings

SNMP traps can be used to arrange e-mail notification when quotas are exceeded. You can access and adapt a sample quota notification script on the NOW site at now.netapp.com in the Tools and Utilities section.

When quotas take effect

Prerequisite for quotas to take effect

You must turn on quotas on a per-volume basis before Data ONTAP applies quotas to quota targets.

Note

You do not need to turn on quotas by using the /etc/rc file.

About quota initialization

After you turn on quotas, Data ONTAP performs quota initialization. This involves scanning the entire file system in a volume and reading from the /etc/quotas file to compute the disk usage for each quota target.

Quota initialization is necessary under the following circumstances:

- You add an entry to the /etc/quotas file, but the quota target for that entry is not currently tracked by the filer.
- You change user mapping in the /etc/usermap.cfg file and you use the QUOTA_PERFORM_USER_MAPPING entry in the /etc/quotas file. For more information about QUOTA_PERFORM_USER_MAPPING, see "Special entries for mapping users" on page 356.
- You change the security style of a qtree from UNIX to either mixed or NTFS.
- You change the security style of a qtree from mixed or NTFS to UNIX.

Quota initialization can take a few minutes. The amount of time required depends on the size of the file system.

When a change to a quota size takes effect

You can change the size of a quota that is being enforced. Resizing an existing quota, whether it is an explicit quota specified in the /etc/quotas file or a derived quota, does not require quota initialization. To learn about derived quotas, see "Understanding derived quotas" on page 339.

Understanding default quotas

About default quotas

You can create a default quota for users, groups, or qtrees. A default quota applies to quota targets that are not explicitly referenced in the /etc/quotas file.

How to override a default quota

If you do not want Data ONTAP to apply a default quota to a particular target, you can create an entry in the /etc/quotas file for that target so that the explicit quota overrides the default quota.

Where default quotas are applied

You apply a default user or group quota on a per-volume or per-qtree basis.

You apply a default tree quota on a per-volume basis. For example, you can specify that a default tree quota be applied to the vol/vol0 volume, which means that all qtrees created in the /vol/vol0 volume are subject to this quota but that qtrees in other volumes are unaffected.

Understanding derived quotas

About derived quotas

Data ONTAP derives the quota information from the default quota entry in the /etc/quotas file and applies it if a write request affects the disk space or number of files used by the quota target. A quota applied due to a default quota, not due to an explicit entry in the /etc/quotas file, is referred to as a derived quota.

Derived user quotas from a default user quota

When a default user quota is in effect, Data ONTAP applies derived quotas to all users in the volume or qtree to which the default quota applies, except those users who have explicit entries in the /etc/quotas file. Data ONTAP also tracks disk usage for root and BUILTIN\Administrators in that volume or qtree.

Example: If a default user quota entry specifies that users in the /vol/vol0 volume are limited to 10 GB of disk space and a user named jsmith creates a file in that volume, Data ONTAP applies a derived quota to jsmith to limit that user's disk usage in the /vol/vol0 volume.

Derived group quotas from a default group quota

When a default group quota is in effect, Data ONTAP applies derived quotas for all UNIX groups in the volume or qtree to which the quota applies, except those groups that have explicit entries in the /etc/quotas file. Data ONTAP also tracks disk usage for the group with GID 0 in that volume or qtree.

Example: If a default group quota entry specifies that groups in the /vol/vol0 volume are limited to 10 GB of disk space and a file is created that is owned by a group named writers, Data ONTAP applies a derived quota to the writers group to limit its disk usage in the /vol/vol0 volume.

Derived tree quotas from a default tree quota

When a default tree quota is in effect, derived quotas apply to all qtrees in the volume to which the quota applies, except those qtrees that have explicit entries in the /etc/quotas file.

Example: If a default tree quota entry specifies that qtrees in the /vol/vol0 volume are limited to 10 GB of disk space and a qtree named projects is created in the /vol/vol0 volume, Data ONTAP applies a derived quota to the /vol/vol0/projects qtree to limit its disk usage.

Default user or group quotas derived from default tree quotas

When a qtree is created in a volume that has a default qtree quota defined in the /etc/quotas file, and that default quota is applied as a derived quota to the qtree just created, Data ONTAP also applies derived default user and group quotas to that qtree.

- If a default user quota or group quota is already defined for the volume containing the newly created qtree, Data ONTAP automatically applies that quota as the derived default user quota or group quota for that qtree.
- If no default user quota or group quota is defined for the volume containing the newly created qtree, then the effective derived user or group quota for that qtree is unlimited. In theory, a single user with no explicit user quota defined can use up the newly defined qtree's entire qtree quota allotment.
- You can replace these initial derived default user quotas or group quotas that Data ONTAP applies to the newly created qtree, by adding explicit or default user or group quotas for the qtree just created to the /etc/quotas file.

Example of a default user quota for a volume applied to a qtree:

Suppose the default user quota in the /vol/vol0 volume specifies that all users are limited to 10 GB of disk space and the default tree quota in the vol/vol0 volume specifies that all qtrees are limited to 100 GB of disk space. If you create a qtree named projects in the /vol/vol0 volume, a default tree quota limits the projects qtree to 100 GB. Also, Data ONTAP applies a derived default user quota, which limits the amount of space used by each user who does not have an explicit user quota defined in the /vol/vol0/projects qtree to 10 GB.

You can change the limits on the default user quota for the /vol/vol0/projects qtree or add an explicit quota for a user in the /vol/vol0/projects qtree by using the `quota resize` command.

Example of no default user quota for a volume applied to a qtree:

Similarly, if no default user quota is defined for the /vol/vol0 volume, the default tree quota for the vol/vol0 volume specifies that all qtrees are limited to 100 GB of disk space, and if you create a qtree named projects, Data ONTAP does not apply a derived default user quota that limits the amount of disk space that users can use on the /vol/vol0/projects tree quota. In theory, a single user with no explicit user quota defined, can use as much as all 100 GB of a qtree's quota if no other user writes to disk space on the new qtree first.

In addition, UID 0, BUILTIN\Administrators, and GID 0 have derived quotas. These derived quotas do not limit the disk space and the number of files. They only track the disk space and the number of files owned by these IDs.

Even with no default user quota defined, no user with files on a qtree can use more disk space in that qtree than is allotted to that qtree as a whole.

Advantages of specifying default quotas

Specifying default quotas offers the following advantages:

- You can automatically apply a limit to a large set of quota targets without typing an entry in the /etc/quotas file for each target. For example, if you want no user to use more than 10 GB of disk space, you can specify a default user quota of 10 GB of disk space instead of creating an entry in the /etc/quotas file for each user.
- You can be flexible in changing quota specifications. Because Data ONTAP already tracks disk and file usage for quota targets of derived quotas, you can easily change the specifications of these derived quotas.

  For example, you can create a default user quota for the /vol/vol0 volume that limits all users to 10 GB of disk space. After quota initialization, this default quota and its derived quotas go into effect.

  If you later decide that a user named jsmith should have a greater quota, you can add an /etc/quotas entry that specifies that jsmith is limited to 20 GB of disk space, overriding the default 10-GB limit. After making the change to the /etc/quotas file, all you need to do is resize the quota, which takes less time than quota reinitialization, to make the jsmith entry effective. Without a default user quota, the newly created jsmith entry requires a quota reinitialization to be effective.

How Data ONTAP identifies users

Two types of user IDs

When applying a user quota, Data ONTAP distinguishes one user from another based on the ID, which can be a UNIX ID or a Windows ID.

Format of a UNIX ID

If you want to apply user quotas to UNIX users, specify the UNIX ID of each user in one of the following formats:

- The user name, as defined in the /etc/passwd file or the NIS password map, such as jsmith.
- The UID, such as 20.
- A file or directory whose UID matches the user. In this case, you should choose a path name that will last as long as the user account remains on the system.

Note

Specifying a file or directory name only enables Data ONTAP to obtain the UID. Data ONTAP does not apply quotas to the file or directory, or to the volume in which the file or directory resides.

Restrictions on UNIX user names: A UNIX user name must not include a backslash (\) or an @ sign, because Data ONTAP treats names containing these characters as Windows names.

Special UID: You cannot impose restrictions on users whose UID is 0. You can specify a quota only to track the disk space and number of files used by this UID.

Format of a Windows ID

If you want to apply user quotas to Windows users, specify the Windows ID of each user in one of the following formats:

- A Windows name specified in pre-Windows 2000 format. For details, see the section on specifying a Windows name in the CIFS chapter of the *File Access Management Guide*.

  If the domain name or user name contains spaces or special characters, the entire Windows name must be in quotation marks, such as “tech support\john#smith”.
- A security ID (SID), as displayed by Windows in text form, such as S-1-5-32-544.

- A file or directory that has an ACL owned by the SID of the user. In this case, you should choose a path name that will last as long as the user account remains on the system.

  Note

  For Data ONTAP to obtain the SID from the ACL, the ACL must be valid. If a file or directory exists in a UNIX-style qtree or if the filer uses UNIX mode for user authentication, Data ONTAP applies the user quota to the user whose UID matches that of the file or directory, not to the SID.

How Windows group IDs are treated

Data ONTAP does not support group quotas based on Windows group IDs. If you specify a Windows group ID as the quota target, the quota is treated like a user quota.

The following list describes what happens if the quota target is a special Windows group ID:

- If the quota target is the Everyone group, a file whose ACL shows that the owner is Everyone is counted under the SID for Everyone.
- If the quota target is BUILTIN\Administrators, the entry is considered a user quota for tracking only. You cannot impose restrictions on BUILTIN\Administrators. If a member of BUILTIN\Administrators creates a file, the file is owned by BUILTIN\Administrators and is counted under the SID for BUILTIN\Administrators.

How quotas are applied to users with multiple IDs

A user can be represented by multiple IDs. You can specify a single user quota entry for such a user by specifying a list of IDs as the quota target. A file owned by any of these IDs is subject to the restriction of the user quota.

Example: A user has the UNIX UID 20 and the Windows IDs corp\john_smith and engineering\jsmith. For this user, you can specify a quota where the quota target is a list of the UID and Windows IDs. When this user writes to the filer, the specified quota applies, regardless of whether the write originates from UID 20, corp\john_smith, or engineering\jsmith.

Quota targets listed in different quota entries are considered separate targets, even though the IDs belong to the same user.

Example: You can specify one quota that limits UID 20 to 1 GB of disk space and another quota that limits corp\john_smith to 2 GB of disk space, even though both IDs represent the same user. Data ONTAP applies quotas to UID 20 and corp\john_smith separately.

If the user has another Windows ID, engineering\jsmith, and there is no applicable quota entry, files owned by engineering\jsmith are not subject to restrictions, even though quota entries are in effect for UID 20 and corp\john_smith.

Root users and quotas

A root user is subject to tree quotas, but not user quotas or group quotas.

When root carries out a file or directory ownership change or other operation (such as the UNIX `chown` command) on behalf of a nonroot user, Data ONTAP checks the quotas based on the new owner but does not report errors or stop the operation even if the nonroot user's hard quota restrictions are exceeded. The root user can therefore carry out operations for a nonroot user (such as recovering data), even if those operations temporarily result in that nonroot user's quotas being exceeded.

Once the ownership transfer is carried out, however, a client system will report a disk space error for the nonroot user who is attempting to allocate more disk space while the quota is still exceeded.

Understanding the /etc/quotas file

About this section

This section provides information about the /etc/quotas file so that you can specify user, group, or tree quotas.

Detailed information

This section discusses the following topics:

- "Overview of the /etc/quotas file" on page 346
- "Fields of the /etc/quotas file" on page 349
- "Sample quota entries" on page 354
- "Special entries for mapping users" on page 356
- "How disk space owned by default users is counted" on page 360

Overview of the /etc/quotas file

Contents of the /etc/quotas file

The /etc/quotas file consists of one or more entries, each entry specifying a default or explicit space or file quota limit for a qtree, group, or user.

The fields of a quota entry in the /etc/quotas file are

```
quota_target type[@/vol/dir/qtree_path] disk [files] [threshold]
[soft_disk] [soft_files]
```

The fields of an /etc/quotas file entry specify the following:

- *quota_target* specifies an explicit qtree, group, or user to which this quota is being applied or an asterisk (*) value that applies this quota as a default to all members of the type specified in this entry that do not have an explicit quota.
- *type* [*@/vol/dir/qtree_path*] specifies the type of entity (qtree, group, or user) to which this quota is being applied. If the type is user or group, this field can optionally restrict this user or group quota to a specific volume, directory, or qtree.
- *disk* is the disk space limit that this quota imposes on the qtree, group, user, or type in question.
- *files* (optional) is the limit on the number of files that this quota imposes on the qtree, group, or user in question.
- *threshold* (optional) is the disk space usage point at which warnings of approaching quota limits are issued.
- *soft_disk* (optional) is a soft quota space limit that, if exceeded, issues warnings rather than rejecting space requests.
- *soft_files* (optional) is a soft quota file limit that, if exceeded, issues warnings rather than rejecting file creation requests.

Note

For a detailed description of the above fields, see “Fields of the /etc/quotas file” on page 349.

Sample /etc/quotas file entries

The following sample quota entry assigns to user jsmith explicit limitations of 500 MB of disk space and 10,240 files in the /vol/rls volume and directory.

```
#Quota target  Type           Disk     Files    thold
#------------  -------------  ----     ------   -------
jsmith         user@/vol/rls  500m      10k      -
```

The following sample quota entry assigns to groups in the /vol/cad directory a default quota of 750 megabytes of disk space and 85,000 files per group. This quota applies to any group in the /vol/cad directory that does not have an explicit quota defined.

```
#Quota target  Type           Disk     Files    thold
#------------  -------------  ----     ------   -------
*              group@/vol/cad 750M     85K      -
```

Note

A line beginning with a pound sign (#) is considered a comment.

Each entry in the /etc/quotas file can extend to multiple lines, but the Files, Threshold, Soft Disk, and Soft Files fields must be on the same line as the Disk field. If they are not on the same line as the Disk field, they are ignored.

Order of entries

Entries in the /etc/quotas file can be in any order. After Data ONTAP receives a write request, it grants access only if the request meets the requirements specified by all /etc/quotas entries. That is, if a quota target is affected by several /etc/quotas entries, the most restrictive entry applies.

Rules for a user or group quota

The following rules apply to a user or group quota:

- If you do not specify a path name to a volume or qtree to which the quota is applied, the quota takes effect in the root volume.
- You cannot impose restrictions on particular quota targets. For the following targets, you can specify quotas entries for tracking purposes only:
  - User with UID 0
  - Group with GID 0
  - BUILTIN\Administrators
- A file created by a member of the BUILTIN\Administrators group is owned by the BUILTIN\Administrators group, not by the member. When determining the amount of disk space or the number of files used by that

user, Data ONTAP does not count the files that are owned by the BUILTIN\Administrators group.

Character coding of the /etc/quotas file

For information about character coding of the /etc/quotas file, see “Contents of the /etc directory” on page 60.

Fields of the /etc/quotas file

Quota Target field

The quota target specifies the user, group, or qtree to which you apply the quota. If the quota is a user or group quota, the same quota target can be in multiple /etc/quotas entries. If the quota is a tree quota, the quota target can be specified only once.

For a user quota: Data ONTAP applies a user quota to the user whose ID is specified in any format described in "How Data ONTAP identifies users" on page 342.

For a group quota: Data ONTAP applies a group quota to a GID, which you specify in the Quota Target field in any of these formats:

- The group name, such as publications
- The GID, such as 30
- A file or subdirectory whose GID matches the group, such as /vol/vol1/archive

Note

Specifying a file or directory name only enables Data ONTAP to obtain the GID. Data ONTAP does not apply quotas to that file or directory, or to the volume in which the file or directory resides.

For a tree quota: The quota target is the complete path name to an existing qtree (for example, /vol/vol1/home).

For default quotas: Use an asterisk (*) in the Quota Target field to specify a default quota. The quota is applied to the following users, groups, or qtrees:

- New users or groups that are created after the default entry takes effect. For example, if the maximum disk space for a default user quota 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 /etc/quotas file. For example, if the maximum disk space for a default user quota is 500 MB, users for whom you have not specified a user quota in the /etc/quotas file can use up to 500 MB of disk space.

Type field

The Type field specifies the quota type, which can be

- User or group quotas, which specify the amount of disk space and the number of files that particular users and groups can own.
- Tree quotas, which specify the amount of disk space and the number of files that particular qtrees can contain.

For a user or group quota: The following table describes the possible values you can specify in the Type field, depending on the volume or the qtree to which the user or group quota is applied.

Quota type	Value in the Type field	Sample entry in the Type field
User quota in a volume	`user@/vol/`*`volume`*	`user@/vol/vol1`
User quota in a qtree	`user@/vol/`*`volume/qtree`*	`user@/vol/vol0/home`
Group quota in a volume	`group@/vol/`*`volume`*	`group@/vol/vol1`
Group quota in a qtree	`group@/vol/`*`volume/qtree`*	`group@/vol/vol0/home`

For a tree quota: The following table describes the values you can specify in the Type field, depending on whether the entry is an explicit tree quota or a default tree quota.

Entry	Value in the Type field
Explicit tree quota	`tree`
Default tree quota	`tree@/vol/`*`volume`* **Example:** `tree@/vol/vol0`

Disk field

The Disk field specifies the maximum amount of disk space that the quota target can use. The value in this field represents a hard limit that cannot be exceeded. The following list describes the rules for specifying a value 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.

Note

The Disk field is not case-sensitive. Therefore, you can use K, k, M, m, G, or g.

- The maximum values you can enter in the Disk field are
  - 4,294,967,295k
  - 4,194,303m
  - 4,095g

Note

If you omit the k, m, or g, Data ONTAP assumes a default value of k. For example, you can enter the maximum Disk field value as 4,294,967,295k or 4,294,967,295.

- The value cannot be specified in decimal notation.
- If you want to track the disk usage but do not want to impose a hard limit on disk usage, type a hyphen (-).
- Do not leave the Disk field blank. The value that follows the Type field is always assigned to the Disk field; thus, for example, Data ONTAP regards the following entries as equivalent:

#Quota Target	type	disk	files
/export	tree		75K
/export	tree	75K	

Files field

The Files field specifies the maximum number of files that the quota target can use. The value in this field represents a hard limit that cannot be exceeded. The following list describes the rules for specifying a value in this field:

- K is equivalent to 1,024, M means 2^{20}, and G means 2^{30}. You can omit the K, M, or G. For example, if you type 100, it means that the maximum number of files is 100.

Note

The Files field is not case-sensitive. Therefore, you can use K, k, M, m, G, or g.

- The maximum values you can enter in the Files field are
  - 4,294,967,295

- 4,194,303k
- 4,095m
- 3g

- The value cannot be specified in decimal notation.
- If you want to track the number of files but do not want to impose a hard limit on the number of files that the quota target can use, type a hyphen (-). If the quota target is root, or if you specify 0 as the UID or GID, you must type a hyphen.
- A blank in this field means there is no restriction on the number of files that the quota target can use. A blank is equivalent to a hyphen (-) in this field. If you leave this field blank, you cannot specify values for the Threshold, Soft Disk, or Soft Files fields.
- The Files field must be on the same line as the Disk field. Otherwise, the Files field is ignored.

Threshold field

The Threshold field specifies the disk space threshold. If a write causes the quota target to exceed the threshold, the write still succeeds but a warning message is logged to the filer console and an SNMP trap is generated. Use the Threshold field to specify disk space threshold limits for CIFS.

The following list describes the rules for specifying a value in this field:

- The format of the Threshold field is the same as the format of the Disk field.
- The maximum values you can enter in the Threshold field are
  - 4,294,967,295k
  - 4,194,303m
  - 4,095g

Note

If you omit the k, m, or g, Data ONTAP assumes the default value of k. For example, you can enter the maximum Threshold field value as 4,294,967,295k or 4,294,967,295.

- The value cannot be specified in decimal notation.
- The Threshold field must be on the same line as the Disk field. Otherwise, the Threshold field is ignored.
- If you do not want to specify a threshold limit on the amount of disk space the quota target can use, enter a hyphen (-) in this field or leave blank.

Soft Disk field

The Soft Disk field specifies the amount of disk space that the quota target can use before a warning is issued. If the quota target exceeds the soft limit, a warning message is logged to the filer console and an SNMP trap is generated. When the soft disk limit returns to normal, another syslog message and an SNMP trap are generated.

The following list describes the rules for specifying a value in this field:

- The format of the Soft Disk field is the same as the format of the Disk field.
- The maximum value you can enter in the Soft Disk field is 4,294,967,295 K.
- The value cannot be specified in decimal notation.
- If you do not want to specify a soft limit on the amount of disk space that the quota target can use, type a hyphen (-) in this field (or leave this field blank if no value for the Soft Files field follows).
- The Soft Disk field must be on the same line as the Disk field. Otherwise, the Soft Disk field is ignored.

Soft Files field

The Soft Files field specifies the number of files that the quota target can use before a warning is issued. If the quota target exceeds the soft limit, a warning message is logged to the filer console and an SNMP trap is generated. When the soft files limit returns to normal, another syslog message and an SNMP trap are generated.

The following list describes the rules for specifying a value in this field.

- The format of the Soft Files field is the same as the format of the Files field.
- The maximum value you can enter in the Soft Files field is 4,294,967,295.
- The value cannot be specified in decimal notation.
- If you do not want to specify a soft limit on the number of files that the quota target can use, type a hyphen (-) in this field or leave the field blank.
- The Soft Files field must be on the same line as the Disk field. Otherwise, the Soft Files field is ignored.

Sample quota entries

Explicit quotas

The following list contains examples of explicit quotas:

- `jsmith  user@/vol/rls  500m  10k`

  The user named jsmith can use 500 MB of disk space and 10,240 files in the /vol/rls volume.

- `jsmith,corp\jsmith,engineering\"john smith",`
  `S-1-5-32-544  user@/vol/rls  500m  10k`

  The user, represented by four IDs, can use 500 MB of disk space and 10,240 files in the /vol/rls volume.

- `writers  group@/vol/cad/proj1  150m`

  The writers group can use 150 MB of disk space and an unlimited number of files in the /vol/cad/proj1 qtree.

- `/vol/cad/proj1  tree 750m 75k`

  The proj1 qtree in the /vol/cad volume can use 750 MB of disk space and 76,800 files.

- `root  user@/vol/rls  -  -`

  Data ONTAP tracks the amount of disk space and the number of files in the /vol/rls volume owned by root.

- `builtin\administrators  user@/vol/rlse  -  -`

  Data ONTAP tracks the amount of disk space and the number of files in the /vol/rlse volume owned by BUILTIN\Administrators. Files created by members of BUILTIN\Administrators are tracked as files owned by BUILTIN\Administrators.

The following sample /etc/quotas file contains default quotas:

```
#Quota Target         type                  disk      files
*                     user@/vol/cad          50M      15K
*                     group@/vol/cad        750M      85K
*                     tree@/vol/cad         100M      75K
jdoe                  user@/vol/cad/proj1   100M      75K
msmith                user@/vol/cad          75M      75K
msmith                user@/vol/cad/proj1    75M      75K
```

The following list explains the effects of these /etc/quotas entries:

- Any user not otherwise mentioned in this file can use 50 MB of disk space and 15,360 files in the /vol/cad volume.
- Any group not otherwise mentioned in this file can use 750 MB of disk space and 87,040 files in the /vol/cad volume.
- Any qtree in the /vol/cad volume not otherwise mentioned in this file can use 100 MB of disk space and 76,800 files.
- If a qtree is created in the /vol/cad volume (for example, a qtree named /vol/cad/proj2), Data ONTAP enforces a derived default user quota and a derived default group quota that have the same effect as these quota entries:

  ```
  *   user@/vol/cad/proj2     50M   15K
  *   group@/vol/cad/proj2   750M   85K
  ```

- If a qtree is created in the /vol/cad volume (for example, a qtree named /vol/cad/proj2), Data ONTAP tracks the disk space and number of files owned by UID 0 and GID 0 in the /vol/cad/proj2 qtree.
- Although a user name jdoe is given an explicit quota of 100 MB of disk space and 76,800 files in the /vol/cad/proj1 qtree, jdoe is restricted to the default limit 50 MB of disk space and 15,360 files. The user named jdoe requires an explicit user quota for the /vol/cad volume to override the default limit. See the entries for the user named msmith for an example.
- A user named msmith can use 75 MB of disk space and 76,800 files in the /vol/cad volume because an explicit quota for this user exists in the /etc/quotas file, overriding the default limit of 50 MB of disk space and 15,360 files.
- A user named msmith can use 75 MB of disk space and 76,800 files in the /vol/cad/proj1 qtree because explicit quotas for this user override the default user quota for the /vol/cad volume and the derived default user quota for the /vol/cad/proj1 qtree.

Special entries for mapping users

Special entries in the /etc/quotas file

The /etc/quotas file supports two special entries whose formats are different from the entries described in "Fields of the /etc/quotas file" on page 349. These special entries enable you to quickly add Windows IDs to the /etc/quotas file. If you use these entries, you can avoid typing individual Windows IDs.

These special entries are

- QUOTA_TARGET_DOMAIN
- QUOTA_PERFORM_USER_MAPPING

Special entry for changing UNIX names to Windows names

The QUOTA_TARGET_DOMAIN entry enables you to change UNIX names to Windows names in the Quota Target field. Use this entry if both of the following conditions apply:

- The /etc/quotas file contains user quotas with UNIX names.
- The quota targets you want to change have identical UNIX and Windows names. For example, a user whose UNIX name is jsmith also has a Windows name that is jsmith.

Format: The following is the format of the QUOTA_TARGET_DOMAIN entry:

`QUOTA_TARGET_DOMAIN` ***`domain_name`***

Effect: For each user quota, Data ONTAP adds the specified domain name as a prefix to the user name. Data ONTAP stops adding the prefix when it reaches the end of the /etc/quotas file or another QUOTA_TARGET_DOMAIN entry without a domain name.

Example: The following example illustrates the use of the QUOTA_TARGET_DOMAIN entry:

```
QUOTA_TARGET_DOMAIN corp
roberts     user@/vol/rls          900M      30K
smith       user@/vol/rls          900M      30K
QUOTA_TARGET_DOMAIN engineering
daly        user@/vol/rls          900M      30K
thomas      user@/vol/rls          900M      30K
QUOTA_TARGET_DOMAIN
stevens     user@/vol/rls          900M      30K
```

Explanation of example: The string corp\ is added as a prefix to the user names of the first two entries. The string engineering\ is added as a prefix to the user names of the third and fourth entries. The last entry is unaffected by the QUOTA_TARGET_DOMAIN entry. The following entries produce the same effects:

```
corp\roberts       user@/vol/rls         900M      30K
corp\smith         user@/vol/rls         900M      30K
engineering\daly   user@/vol/rls         900M      30K
engineering\thomas user@/vol/rls         900M      30K
stevens            user@/vol/rls         900M      30K
```

Special entry for mapping names

The QUOTA_PERFORM_USER_MAPPING entry enables you to map UNIX names to Windows names or vice versa. Use this entry if both of the following conditions apply:

- There is a one-to-one correspondence between UNIX names and Windows names.
- You want to apply the same quota to the user whether the user uses the UNIX name or the Windows name.

Note

The QUOTA_PERFORM_USER_MAPPING entry does not work if the QUOTA_TARGET_DOMAIN entry is present.

How names are mapped: Data ONTAP consults the /etc/usermap.cfg file to map the user names.

Format: The QUOTA_PERFORM_USER_MAPPING entry has the following format:

```
QUOTA_PERFORM_USER_MAPPING [ on | off ]
```

Data ONTAP maps the user names in the Quota Target fields of all entries following the `QUOTA_PERFORM_USER_MAPPING on` entry. It stops mapping when it reaches the end of the /etc/quotas file or when it reaches a `QUOTA_PERFORM_USER_MAPPING off` entry.

Example: The following example illustrates the use of the QUOTA_PERFORM_USER_MAPPING entry:

```
QUOTA_PERFORM_USER_MAPPING on
roberts        user@/vol/rls         900M      30K
corp\stevens   user@/vol/rls         900M      30K
QUOTA_PERFORM_USER_MAPPING off
```

Explanation of example: If the /etc/usermap.cfg file maps roberts to corp\jroberts, the first quota entry applies to the user whose UNIX name is roberts and whose Windows name is corp\jroberts. A file owned by a user with either user name is subject to the restriction of this quota entry.

If the usermap.cfg file maps corp\stevens to cws, the second quota entry applies to the user whose Windows name is corp\stevens and whose UNIX name is cws. A file owned by a user with either user name is subject to the restriction of this quota entry.

The following entries produce the same effects:

```
roberts,corp\jroberts        user@/vol/rls         900M      30K
corp\stevens,cws             user@/vol/rls         900M      30K
```

Importance of one-to-one mapping: If the name mapping is not one-to-one, the QUOTA_PERFORM_USER_MAPPING entry produces confusing results, as illustrated in the following examples.

Example of multiple Windows names for one UNIX name: Suppose the /etc/usermap.cfg file contains the following entries:

```
domain1\user1 => unixuser1
domain2\user2 => unixuser1
```

Data ONTAP displays a warning message if the /etc/quotas file contains the following entries:

```
QUOTA_PERFORM_USER_MAPPING on
domain1\user1 user 1M
domain2\user2 user 1M
```

The /etc/quotas file effectively contains two entries for unixuser1. Therefore, the second entry is treated as a duplicate entry and is ignored.

Example of wildcard entries in usermap.cfg: Confusion can result if the following conditions exist:

- The /etc/usermap.cfg file contains the following entry:
  ```
  *\* *
  ```
- The /etc/quotas file contains the following entries:
  ```
  QUOTA_PERFORM_USER_MAPPING on
  unixuser2 user 1M
  ```

Problems arise because Data ONTAP tries to locate unixuser2 in one of the trusted domains. Because Data ONTAP searches domains in an unspecified order, unless the order is specified by the `cifs.search_domains` option, the result becomes unpredictable.

What to do after you change usermap.cfg: If you make changes to the /etc/usermap.cfg file, you must turn quotas off and then turn quotas back on for the changes to take effect. For more information about turning quotas on and off, see "Turning quotas on or off" on page 361.

How disk space owned by default users is counted

Disk space used by the default UNIX user

For a Windows name that does not map to a specific UNIX name, Data ONTAP uses the default UNIX name defined by the `wafl.default_unix_user` option when calculating disk space. Files owned by the Windows user without a specific UNIX name are counted against the default UNIX user name if either of the following conditions applies:

- The files are in qtrees with UNIX security style.
- The files do not have ACLs in qtrees with mixed security style.

Disk space used by the default Windows user

For a UNIX name that does not map to a specific Windows name, Data ONTAP uses the default Windows name defined by the `wafl.default_nt_user` option when calculating disk space. Files owned by the UNIX user without a specific Windows name are counted against the default Windows user name if the files have ACLs in qtrees with NTFS security style or mixed security style.

Turning quotas on or off

About turning quotas on or off

You can turn quotas on or off for one volume at a time. You cannot enter one command for turning quotas on or off for all volumes.

Prerequisites for turning quotas on

To turn quotas on, you must meet the following prerequisites:

- If the /etc/quotas file contains user quotas that use Windows IDs as targets, CIFS must be running.
- All UNIX and Windows names listed in the Quota Target field of the /etc/quotas file must be valid. For example, if the file contains a misspelled user name, Data ONTAP cannot complete quota initialization and turn on quotas.

Turning quotas on

To turn on quotas, complete the following step.

Step	Action
1	Enter the following command: **`quota on`** ***`volume`*** **Example:** The following example turns on quotas on a volume named cad: `quota on cad`

Turning quotas off

To turn quotas off, complete the following step.

Step	Action
1	Enter the following command: **`quota off`** ***`volume`*** **Example:** The following example turns off quotas on a volume named cad: `quota off cad`

Canceling a quota on command

To cancel a `quota on` command, follow this step.

Step	Action
1	Enter the following command: **`quota off`** ***`volume`*** **Result:** The `quota off` command cancels the `quota on` command and stops quota initialization.

Turning quota message logging on or off

About turning quota message logging on or off

You can turn quota message logging on or off for a volume or for all volumes. You can optionally specify a time interval during which quota messages will not be logged.

Turning quota message logging on

To turn quota message logging on, complete the following step.

Step	Action
1	Enter the following command: **`quota logmsg on [ interval ] [ -v vol | all ]`** *interval* is the time period during which quota message logging is disabled. The interval is a number followed by d, h or m for days, hours, and minutes, respectively. Quota messages are logged at the end of each interval. If no interval is specified, Data ONTAP logs quota messages every 60 minutes. For continuous logging, specify 0 for the interval. Use `-v` *vol* to specify a volume name. Use `all` to apply the interval to all volumes in the filer. **Note** If you specify a short interval, less than five minutes, quota messages might not be logged exactly at the specified interval rate because Data ONTAP buffers quota messages before logging them.

Turning quota message logging off

To turn quota message logging off, complete the following step.

Step	Action
1	Enter the following command: **`quota logmsg off`**

Displaying settings for quota message logging

To display the current settings for quota message logging, complete the following step.

Step	Action
1	Enter the following command: `quota logmsg`

Creating or resizing quotas

Ways to change quotas

You can change quotas in the following ways:

- Create a quota for a quota target that is currently not tracked by your filer. For example, you can add a user quota for jsmith, whose disk usage is currently unrestricted and untracked by any quota.
- Resize a quota that is currently tracked by the filer. You can accomplish the following tasks through quota resizing:
  - Change the size of a quota that has an explicit entry in the /etc/quotas file.
  - Override a derived quota, which does not have an explicit entry in the /etc/quotas file. For example, you can impose a less-restrictive quota on a particular user who is otherwise restricted by a default user quota.

You can determine from the quota report whether your filer is tracking disk usage for a particular user, group, or qtree. A quota in the quota report indicates that the filer is tracking the disk space and the number of files owned by the quota target. For more information about quota reports, see “Understanding quota reports” on page 372.

Creating a quota for a quota target that the filer is not tracking

If a user, group, or qtree is not under the effect of any /etc/quotas entry, complete the following steps to create a quota for the user, group, or qtree.

Step	Action
1	Create an entry in the /etc/quotas file for the user, group, or qtree.

<table>
<tr><th>Step</th><th colspan="2">Action</th></tr>
<tr><td rowspan="4">2</td><td>If quotas are...</td><td>Then...</td></tr>
<tr><td>Currently turned off</td><td>Enter the following command to turn quotas on:
<tt>quota on volume</tt></td></tr>
<tr><td>Currently turned on</td><td>Enter the following commands to reinitialize quotas:
<tt>quota off volume;quota on volume</tt></td></tr>
<tr><td colspan="2">Result: Data ONTAP initializes quotas for the specified volume. Quota initialization can take a few minutes. After the initialization, all entries in the /etc/quotas file and all derived quotas take effect.</td></tr>
<tr><td>3</td><td colspan="2">Enter the following command to check the status of quota initialization:
<tt>quota</tt>
Result: Data ONTAP displays the percentage of quota initialization completed. After Data ONTAP finishes quota initialization, it displays the following message:
<tt>vol_name: quotas are on</tt></td></tr>
</table>

Example: Suppose the /etc/quotas file currently contains the following entry:

```
jsmith     user@/vol/cad     100M    15K
```

The following entry applies quotas to quota targets other than jsmith:

```
*      user@/vol/cad      500M      15K
```

The following commands activate the newly created entry:

```
quota off cad;quota on cad
```

After quota initialization, all users, except for jsmith, are limited to 500 MB of disk space and 15,360 files in the /vol/cad volume.

Resizing an existing quota

To resize a quota, complete the following steps.

Step	Action	
1	**If...**	**Then...**
	The quota being enforced is an explicit quota in the /etc/quotas file	Change the specification of the explicit quota.
	The quota being enforced is a derived quota	Create an entry in the /etc/quotas file to override the derived quota.
2	Enter the following command: `quota resize volume` **Result:** The change you made in Step 1 takes effect.	

Example of resizing an explicit quota: Suppose the /etc/quotas file contains the following entry:

```
jsmith      user@/vol/cad      50M      15K
```

If you want jsmith to have a more restrictive quota, change the existing entry to the following entry:

```
jsmith      user@/vol/cad      20M      15K
```

Then enter the following command:

```
quota resize cad
```

The new quota for jsmith goes into effect immediately.

Example of resizing a derived quota: Suppose the /etc/quotas file contains the following entry:

```
*      user@/vol/cad      50M      15K
```

If you want jsmith to have a less-restrictive quota, add the following explicit quota for jsmith:

```
jsmith      user@/vol/cad      100M      15K
```

Then enter the following command:

```
quota resize cad
```

The new quota for jsmith goes into effect immediately.

Recommendation after an extensive revision of the /etc/quotas file

After you edit the /etc/quotas file extensively, it is better to reinitialize quotas than to resize quotas. Although quota reinitalization takes a longer time than quota resizing, quota reinitialization ensures that all quotas listed in the file become effective. Quota resizing works only for quota targets that are currently tracked by Data ONTAP. If your revisions to the /etc/quotas file include entries for quota targets that the filer is not tracking, the `quota resize` command ignores them. For each entry that it skips, the `quota resize` command prints the following message:

`quota resize: warning: /etc/quotas:`*n*`:skipping new entry`

n is a line number.

Deleting quotas

About quota deletion

You can remove quota restrictions for a quota target in two ways:

- Delete the /etc/quotas entry pertaining to the quota target.
- Change the /etc/quotas entry so that there is no restriction on the amount of disk space or the number of files owned by the quota target. After the change, Data ONTAP continues to keep track of the disk space and the number of files owned by the quota target but stops imposing the restrictions on the quota target. The procedure for removing quota restrictions in this way is the same as that for resizing an existing quota.

Deleting a quota

To delete a quota, complete the following steps.

Step	Action
1	Delete from the /etc/quotas file the entry pertaining to the quota target.
2	Enter the following command: `quota resize volume`

Effects of qtree changes on quotas

Effect of deleting a qtree on tree quotas

When you delete a qtree, all quotas that are applicable to that qtree, whether they are explicit or derived, are automatically deleted.

If you create a new qtree with the same name as the one you deleted, the quotas previously applied to the deleted qtree are not applied automatically to the new qtree. If a default tree quota exists, Data ONTAP creates new derived quotas for the new qtree. However, explicit quotas in the /etc/quotas file do not apply until you reinitialize quotas.

Effect of renaming a qtree on tree quotas

When you rename a qtree, Data ONTAP keeps the same ID for the tree. As a result, all quotas applicable to the qtree, whether they are explicit or derived, continue to be applicable.

Effects of changing qtree security style on user quota usages

Because ACLs apply in qtrees using NTFS or mixed security style but not in qtrees using UNIX security style, changing the security style of a qtree through the `qtree security` command might affect how a UNIX or Windows user's quota usages for that qtree are calculated.

Example: If NTFS security is in effect on qtree A and an ACL gives Windows user Windows/joe ownership of a 5MB file, then user Windows/joe is charged 5 MB of quota usage on qtree A.

If the security style of qtree A is changed to UNIX and Windows user Windows/joe is default mapped to UNIX user joe, the ACL that charged 5 MB of diskspace against the quota of Windows/joe is ignored when calculating the quota usage of UNIX user joe.

Caution

To make sure quota usages for both UNIX and Windows users are properly calculated after you use the `qtree security` command to change style change, turn quotas for the volume containing that qtree off and then back on again using the `quota off` *volume* and `quota on` *volume* commands.

If you change the security style from UNIX to either mixed or NTFS, previously hidden ACLs become visible, any ACLs that were ignored become effective again and the NFS user information is ignored. If no ACL existed before, then the NFS information is used in the quota calculation.

Note

Only UNIX group quotas apply to qtrees. Changing the security style of a qtree, therefore, does not affect the quota usages that groups are subject to.

Understanding quota reports

About this section

This section provides information about quota reports.

Detailed information

The following sections provide detailed information about quota reports:

- "Types of quota reports" on page 373
- "Overview of the quota report format" on page 374
- "Detailed information about the ID and Quota Specifier fields" on page 376
- "How to display disk space thresholds" on page 379
- "How to display soft limits" on page 380
- "Displaying a quota report" on page 381

Types of quota reports

Types of quota reports

You can display these types of quota reports:

- A quota report for all volumes with quotas turned on. It contains the following types of information:
  - Default quota information, which is the same information as that in the /etc/quotas file
  - Current disk space and the number of files owned by a user, group, or qtree that has an explicit quota in the /etc/quotas file
  - Current disk space and the number of files owned by a user, group, or qtree that is the quota target of a derived quota, if the user, group, or qtree currently uses some disk space
- A quota report for a specified path name. It contains information about all the quotas that apply to the specified path name.

  For example, in the quota report for the /vol/cad/specs path name, you can see the quotas to which the disk space used by the /vol/cad/specs path name is charged. If a user quota exists for the owner of the /vol/cad/specs path name and a group quota exists for the /vol/cad volume, both quotas appear in the quota report.

Overview of the quota report format

Contents of the quota report

The following table lists the fields displayed in the quota report and the information they contain.

Heading	Information
Type	Quota type: user, group, or tree.
ID	User ID, UNIX group name, qtree name, or an asterisk (for a default quota).
Volume	Volume to which the quota is applied.
Tree	Qtree to which the quota is applied.
K-Bytes	Current amount of disk space used by the quota target and the maximum amount of disk space allowed for the quota target. If the quota is a default quota, the amount of disk space used is always 0. If a disk space threshold is specified for the quota target, you can optionally display the disk space threshold in this field. If a soft disk limit is specified for the quota target, you can also display the soft disk limit in this field.
Files	Current number of files used by the quota target and the maximum number of files allowed for the quota target. If the quota is a default quota, the number of files used is always 0. If a soft files limit is specified for the quota target, you can also display the soft files limit in this field.
Quota Specifier	For an explicit quota, it shows how the quota target is specified in the /etc/quotas file. For a derived quota, the field is blank.

Note

The exact format of the ID, K-Bytes, Files, and Quota Specifier fields depends on the command used to generate the report.

Detailed information about the ID and Quota Specifier fields

Factors affecting the contents of the fields

The information contained in the ID and Quota Specifier fields can vary according to these factors:

- Type of user—UNIX or Windows—to which a quota applies.
- The specific command used to generate the quota report.

Contents of the ID field

In general, the ID field of the quota report displays a user name instead of a UID or SID; however, the following exceptions apply:

- For a quota with a UNIX user as the target, the ID field shows the UID instead of a name if no user name for the UID is found in the password database or if you specifically request the UID by including the `-q` option in the `quota reports` command.
- For a quota with a Windows user as the target, the ID field shows the SID instead of a name if either of the following conditions applies:
  - The SID is specified as a quota target and the SID no longer corresponds to a user name.
  - The filer cannot find an entry for the SID in the SID-to-name map cache and cannot connect to the domain controller to ascertain the user name for the SID when it generates the quota report.

Formats available for the fields

The information in the ID and Quota Specifier fields is available in three formats:

- A default format generated by the `quota report` command
- An enhanced format generated by the `quota report -u` command
- An enhanced format generated by the `quota report -x` command

Default format

The `quota report` command without options generates the default format for the ID and Quota Specifier fields.

The ID field: If a quota target contains only one ID, the ID field displays that ID. Otherwise, the ID field displays one of the IDs from the list.

The ID field displays information in the following formats:

- For a Windows name, the first seven characters of the user name with a preceding backslash are displayed. The domain name is omitted.
- For a SID, the last eight characters are displayed.

The Quota Specifier field: The Quota Specifier field displays an ID that matches that in the ID field. The ID is displayed the same way the quota target is specified in the /etc/quotas file.

Examples: The following table shows what is displayed in the ID and Quota Specifier fields based on the quota target in the /etc/quotas file.

Quota target in the /etc/quotas file	ID field of the quota report	Quota Specifier field of the quota report
CORP\john_smith	\john_sm	CORP\john_smith
CORP\john_smith,NT\js	\john_sm or \js	CORP\john_smith or NT\js
S-1-5-32-544	5-32-544	S-1-5-32-544

Enhanced format with quota report -u

The `quota report -u` command is useful if you have quota targets that are multiple IDs. It provides more information in the ID and Quota Specifier fields than the default format.

The ID field: The ID field displays all the IDs listed in the quota target of a user quota in the following format:

- On the first line, the format is the same as the default format.
- Each additional name in the quota target is displayed on a separate line in its entirety.

The Quota Specifier field: The Quota Specifier field displays the same list of IDs as specified in the quota target.

Example: The following table shows what is displayed in the ID and Quota Specifier fields based on the quota target in the /etc/quotas file. In this example, the SID maps to the user name NT\js.

Quota target in /etc/quotas	ID field of the quota report	Quota Specifier field of the quota report
CORP\john_smith,S-1-5-21-123456-7890-1234-1166	\john_sm NT\js	CORP\john_smith,S-1-5-21-123456-7890-1234-1166

Enhanced format with quota report -q

The `quota report -q` command displays the quota target's UNIX UID or GID in numeric form. No lookup of the name associated with the target ID is done.

For Windows IDs, the textual form of the SID is displayed.

How to display disk space thresholds

Commands for including disk space thresholds

If the /etc/quotas file includes disk space thresholds, you can use the `-t` option of the `quota report` command to display the thresholds in the K-Bytes field of the quota report.

The quota report generated by the `quota report -x` command also includes disk space thresholds.

If you use both the `-x` and the `-t` options, the resulting quota report is the same as the one generated by the `quota report -x` command.

Note

Disk space fields shown in a quota report can appear incorrect if you do not specify disk space limits and thresholds as a multiple of 4 KB. This happens because disk space fields are always rounded up to the nearest multiple of 4 KB to match disk space limits that are translated into 4-KB disk blocks.

How to display soft limits

Command for including soft disk and soft files limits

If the /etc/quotas file includes soft disk or soft files limits, you can use the `-s` option of the `quota report` command to display the soft disk limits in the K-Bytes field and the soft files limits in the Files field of the quota report.

Displaying a quota report

Displaying a quota report for all quotas

To display a quota report for all quotas, complete the following step.

Step	Action
1	Enter the following command: **`quota report [-s] [ -u | -x ] [ -t ] [-q]`**

Displaying a quota report for a specified path name

To display a quota report for a specified path name, complete the following step.

Step	Action
1	Enter the following command: **`quota report [-s] [ -u | -x ] [ -t ] [-q]`** ***`path_name`*** *path_name* is a complete path name to a file or directory, such as /vol/vol0/etc.

Effects of hard quotas on clients

How Data ONTAP calculates disk space

In general, when determining whether a hard quota will be exceeded, Data ONTAP counts only the blocks that are actually used in a file. In a file in which blocks are not written contiguously, the number of blocks actually used can be smaller than the file size.

There is an exception to this rule, however, after Data ONTAP performs a zero-byte write beyond the end of the file as requested by a CIFS client. In such a case, Data ONTAP reserves file system blocks and updates quotas based on the file size, not the number of blocks actually used. Subsequently, any CIFS and NFS write that extends this file uses these semantics: If the file is extended beyond the current end of file and the additional quota reservation based on the file size exceeds the hard quota block limit, the request fails. The write request also fails if the additional block reservation exceeds the number of available file system blocks.

Total capacity shown on Windows clients

When a Windows client displays the properties of a drive that is mapped to a share on the filer, the capacity of the drive represents the disk space that the user can use after all applicable quotas are taken into consideration.

Example: A user, CORP\jsmith, is restricted to 10 GB in the /vol/vol0 volume, and a hard tree quota restricts the /vol/vol0/test qtree to 5 GB. Suppose a share named QA corresponds to the /vol/vol0/test qtree and the CORP\jsmith user maps the H: drive to the QA share. If the CORP\jsmith user displays the properties of the H: drive, the capacity of the drive is 5 GB, which is the maximum disk space that the corp\jsmith user can use in the /vol/vol0/test qtree.

Data ONTAP messages when quotas are exceeded

When it receives a write request, Data ONTAP first determines whether the file to be written is in a qtree. If the write would exceed the hard tree quota, Data ONTAP logs the following error message:

`tid` *`tree_id`*`: tree quota exceeded on volume` *`vol_name`*

tree_id is the number in the ID column of the quota report for an explicit tree quota.

If the qtree is not full but the write would cause either the user or group hard quota to be exceeded, Data ONTAP logs one of the following errors:

```
uid user_id: disk quota exceeded on volume vol_name
gid group_id: disk quota exceeded on volume vol_name
```

user_id is the UID and *group_id* is the GID.

Errors returned to clients

To the client, Data ONTAP 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 users about hard quotas being exceeded.

On NFS clients: If a write from an NFS client to a filer causes a hard 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 hard 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 hard quota to be exceeded, the error message that the user sees depends on the application. For example, on a SunOS 4.x client, if a user tries to save a file by using `vi` when that user's disk quota is reached, an error message similar to the following is displayed:

```
Disc quota exceeded[Warning - /toaster/home/file1 is incompleted]
```

On CIFS clients: If a write from a CIFS client to a filer causes a hard quota to be exceeded, the user experience depends on the operating system and the application. The following examples show possible error messages:

- An application might display a message as follows:

  `Cannot write file filename`

- When a user tries to copy a file to the filer through Windows Explorer in Windows 95, the following error message is displayed:

  `Cannot create or replace` *`filename`*`: Cannot read from the source file or disk.`

System Information and Performance 9

About this chapter

This chapter describes commands that give you information about the filer and how it is performing.

Note

Some options for different commands can gather the same information. For example, the `vol status -r` command and `sysconfig -r` command gather the same RAID information and present it in the same format.

Topics in this chapter

This chapter discusses the following topics:

- "Filer configuration information" on page 386
- "Storage information" on page 389
- "Volume information" on page 391
- "File statistics for volumes" on page 392
- "Options for the filestats command" on page 396
- "Checksum upgrade status" on page 399
- "Filer statistics" on page 401
- "Environmental status information" on page 404
- "Fibre Channel information" on page 407
- "NFS and RPC statistics" on page 408
- "CIFS statistics" on page 414
- "DAFS statistics" on page 423
- "Improving filer performance" on page 424

Filer configuration information

Version information

The `version` command displays the version of Data ONTAP currently running on a filer.

Hardware configuration

The `sysconfig` command displays information about the filer's hardware configuration. The exact types of information displayed depend on the command options.

For more information about the `sysconfig` command, see the na_sysconfig(1) man page.

Expansion cards

The `sysconfig -c` command checks that expansion cards are in the appropriate slots and reports any configuration errors. If there are no configuration errors, the `sysconfig -c` command reports the following:

```
sysconfig: There are no configuration errors
```

Disk information

The `sysconfig -d` command displays product information about each disk in the filer.

RAID and checksum information

The `sysconfig -r` command displays the status of plexes and volumes; the RAID configuration; and checksum information about the parity disks, data disks, and hot spare disks, if any. This information is useful for the following purposes:

- Locating a disk referenced in a console message
- Determining how much space on each disk is available to the filer
- Determining the status of disk operations, such as RAID scrubbing, reconstruction, parity verification, adding a hot spare, and disk failure
- Determining the number of spare disks
- Determining a checksum type for a volume

You can obtain the information displayed by `sysconfig -r` from SNMP, using the Network Appliance custom MIB. For information about SNMP, see the *Network Management Guide*.

Tape drive information

The `sysconfig -t` command displays device and configuration information for each tape drive on the system. You can use this command to determine the capacity of the tape drive and the device name before you use the `dump` and `restore` commands.

Volume information

The `sysconfig -V` command displays RAID group and disk information about each volume.

Tape library information

The `sysconfig -m` command displays tape library information. Before you can use this option, ensure that the filer was booted with the autoload setting of the tape library off.

Overall filer information

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.

The `sysconfig -a` command displays the same information as the `-v` option, but the information is more detailed.

Note

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.

Comprehensive filer information

The `sysconfig -A` command displays filer information gathered by the following commands, one after the other:

- `sysconfig`
- `sysconfig -c`

- `sysconfig -d`
- `sysconfig -V`
- `sysconfig -r`
- `sysconfig -m`

Therefore, when you use the `sysconfig -A` command, Data ONTAP lists information about configuration errors, disk drives, medium changers, RAID details, tape devices, and volume details.

Storage information

Command for displaying storage information

You use the `storage show` command to display information about storage components for a filer, including information about all disks and adapters.

For more information on the `storage` command, see the na_storage(1) man page and Chapter 5, “Disk Management,” on page 201.

Adapter information

The `storage show adapter` command displays information about all storage adapters currently in your filer. The storage `show adapter` *name* command displays information only for the adapter specified.

Hub information

The `storage show hub` command displays information about all hubs attached to your filer. The `storage show hub` *name* command displays information only for the hub specified.

Disk information

The `storage show disk` command displays information about all disks currently attached to your filer. The `storage show adapter -p` command displays the primary and secondary paths to a disk device.

Medium changer information

The `storage show mc` command displays information about all medium changers (tape libraries) attached to your filer. The `storage show mc` *name* command displays information for only the media changer specified.

Switch port information

The `storage show port` command displays information about all ports on all switches attached to your filer. The `storage show port` *name* command displays information only for the port specified.

Switch information

The `storage show switch` command displays information about all switches attached to your filer. The `storage show switch` *name* command displays information only for the switch specified.

Tape information

The `storage show tape` command displays information about all tape devices attached to your filer. The `storage show tape` *name* command displays information only for the tape drive specified.

Tape statistics

The `storage stats tape` *name* command displays statistics about the tape drive specified. You must specify the tape drive name.

Reset tape drive statistics

The `storage stats tape zero` *name* command resets all statistics for the tape drive specified. You must specify the tape drive name.

Volume information

Command for displaying volume information

You use the `vol status` command to display information about volume configurations. The types of information displayed depend on the command options. When you specify a volume, such as `vol status vol0`, the information for that volume is displayed; when you do not specify a volume, the status of all volumes in the filer is displayed.

For more information about the `vol` command, see the na_vol(1) man page.

Volume state information

With no options, the `vol status` command displays a concise synopsis of volume states. This includes the volume name, whether it is online or offline, and other states, such as partial and degraded.

Block information

The `vol status -b` command displays information about volume and file system block sizes. The information displayed is useful in determining whether volumes are the appropriate size for SnapMirror.

Disk information

The `vol status -d` command displays information about disks. The disk information is the same as the information from the `sysconfig -d` command.

RAID and checksum information

The `vol status -r` command displays RAID, plex, and checksum information for a volume. The display is the same as the `sysconfig -r` display.

RAID information for each group

The `vol status -v` command displays information about each RAID group within a volume and the settings of the volume options.

Volume language

The `vol status -l` command displays the language used by each volume.

File statistics for volumes

How Data ONTAP provides file statistics

The `filestats` command quickly displays 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 gives you 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 [-g] [-u] [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 “Options for the filestats command” on page 396.

Getting a file statistics summary

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.

Example with no options specified

The following example shows sample output from the `filestats` command, without any options, for the hourly.1 snapshot on vol0. Output of this option includes a breakdown of files by size, age, user ID, and group ID, and the cumulative number of inodes for each value:

```
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
```

Note

The # character preceding user IDs or group IDs in the UID and GID sections of the `filestats` command output indicates that the IDs cannot be found in the /etc/passwd and /etc/hosts files on the filer.

Using 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%) 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 of determining 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, see the *Data Protection Guide*.

Options for the filestats command

Options to use with the filestats command

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

- `-g`
- `-u`
- `ages`
- `timetype`
- `sizes`
- `style`

About the group list option

The `-g` option of the `filestats` command enables you to generate separate file usage summaries for each group ID. For each group ID, a separate table containing information about file sizes and ages is listed.

About the user list option

The `-u` option of the `filestats` command enables you to generate separate file usage summaries for each user ID. For each group ID, a separate table containing information about file sizes and ages is listed.

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

To display files with ages under 900 seconds (15 minutes), under 4 hours, and under 7 days, enter the following command:

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

The age section of 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
```

For an example of the entire `filestats` command output, see “Example with no options specified” on page 393.

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
`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 suffixes at the end of a number you specify:

- `K` (kilobytes).
- `M` (megabytes).
- `G` (gigabytes).
- `*` (a special value for listing all unique file sizes, one line per unique size). Using the * suffix can result in output of several thousands of lines.

Example of the sizes option

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 file size section of 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
```

For an example of the entire `filestats` command output, see “Example with no options specified” on page 393.

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

Checksum upgrade status

The checksum upgrade process

The checksum upgrade process is a long process that occurs when you upgrade Data ONTAP on your filer. The following is a brief description of how Data ONTAP upgrades the checksum.

Stage	Description
1	WAFL upgrade—Data ONTAP moves allocated blocks that are in the checksum space out of the checksum space. After Data ONTAP moves all the blocks in the active file system out of the checksum space, Data ONTAP waits for all snapshots created before the upgrade to be deleted. Automatically scheduled snapshots are deleted automatically after a period of time. You must delete user-created snapshots manually.
2	RAID upgrade, first phase—Data ONTAP initializes all checksum blocks to a known state.
3	RAID upgrade, second phase—Data ONTAP seeds the checksum blocks with the correct checksums of the corresponding data blocks. This phase occurs during the first disk scrub after Data ONTAP initializes all checksum blocks (first phase).

Examples of checksum detailed displays

The `vol status -c` command provides details about a block checksum upgrade.

The following example shows checksum status during the WAFL upgrade phase:

```
vol status -c
Volume  Checksum status
  vol0  Checksums initializing:
            WAFL upgrade in progress
```

The following example shows snapshots that must be deleted before the WAFL upgrade can finish:

```
vol status -c
Volume  Checksum status
  vol0  Checksums initializing:
```

```
            WAFL upgrade will complete when the following
            snapshots are deleted:
                date             name
                ------------     ------------
                Jan 21 18:00     snap.0
                Sep 23 10:04     snap.2
```

The following example shows checksum status during RAID upgrade, phase 1:

```
vol status -c
Volume  Checksum status
  vol0  Checksums initializing:
            RAID upgrade phase 1 (of 2) in
            progress (25% complete).
```

The following example shows checksum status during RAID upgrade, phase 2:

```
vol status -c
Volume  Checksum status
  vol0  Checksums initializing:
            RAID upgrade phase 2 (of 2) will be performed
            during next disk scrub.
```

The following example shows a completed checksum upgrade:

```
vol status -c
Volume  Checksum status
  vol0  Checksums active:
```

Filer statistics

Commands for displaying filer statistics

You use the `sysstat` command to display current filer statistics and the `uptime` command to display filer statistics that have accumulated since the last system boot.

About the sysstat command

The `sysstat` command displays statistics about the following:

- CPU utilization
- File server protocol operations
- Network traffic
- Read and write operations on disks and tapes
- Buffer cache age

The `sysstat` output is particularly useful for studying the performance of your filer.

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:

- If CPU utilization is high, you might need to add another filer or upgrade to a more powerful one.
- If the network traffic shown in the Net Kb/s column is at or near the capacity of the network interfaces (1,000 Kb/s for 10Base-T Ethernet adapters, 10,000 Kb/s for FDDI or 100Base-T Ethernet adapters, and 15,000 Kb/s for ATM OC-3 adapters), you might need to add another network adapter card to the system.

Note

For cluster systems, the `sysstat` command reflects the filer that is executing the commands, not the partner.

Example: The following is an example of the default output format for the `sysstat` command with no options.

```
sysstat

CPU   NFS  CIFS  HTTP   Net kB/s  Disk kB/s    Tape kB/s   Cache
                        in  out   read write   read write  age
```

Display interval

By default, the filer displays statistics every 15 seconds; you can specify the interval, in seconds, as an argument to the `sysstat` command.

Extended utilization statistics

The `sysstat -u` command displays extended utilization statistics instead of the default display. The statistics include the total number of operations per second for NFS, CIFS, and HTTP; the cache age; the cache hit percentage; the percentage of time the filer spends in a consistency point; the causes of the consistency points; and the disk utilization percentage of the busiest disk.

Note

Disk utilization is an access-based value and not the percentage of the disk used. High percentages for disk utilization might indicate a disk bottleneck. A value of 60% can indicate a disk bottleneck, and a value of 80% is a definite sign of one.

Sysstat statistics in extended format

The `sysstat -x` command displays all system activity statistics in an extended output format. This produces output that is longer than 80 columns and is intended to be displayed on a wide viewer or to be captured to a file for later analysis. For the information displayed in each column, see the man page for the `sysstat` command.

Filer statistics summary

The `sysstat -s` command displays a summary of the gathered statistics after the `sysstat` command terminates. The summary displays maximum and minimum values, and an average of values.

Note

Except for the average percentage of cache hits, the summary of averages for percentage values is an average of percentages, not a true mean of the data.

Example: Entering the following command displays extended utilization statistics followed by a summary:

```
sysstat -u -s
```

Command for terminating the sysstat command

You can terminate the `sysstat` command by either typing Ctrl-C or setting the interval count by using the `sysstat -c` command.

Example: The following command displays 20 samples of extended utilization statistics:

```
sysstat -u -c 20
```

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, CIFS, HTTP, and DAFS operations the system has performed since it was last booted.

Example: A typical example of the `uptime` display is as follows:

```
uptime

8:54am up 2 days 22:23, 3122520 NFS ops, 130 CIFS ops, 62 HTTP ops,
0 DAFS ops
```

Note

For cluster systems, the `uptime` command reflects the filer that is executing the commands, not the partner.

Environmental status information

Command for displaying environmental status

You use the `environment` command to display information about the filer environment. The exact types of information displayed depend on the command options.

For more information about the `environment` command, see the na_environment(1) man page.

Information provided by the environment command

The `environment` command displays the following types of information about the filer environment:

- Shelf status and temperature
- Filer component information
- Filer temperature
- Devices attached to the filer

You can query information about the following items:

- Disk shelves
- The filer backplane
- The filer CPU
- The filer power supply
- The filer temperature

When the environment command is run

Data ONTAP runs the `environment` command under the following conditions:

- Once every hour. In this case, no output is displayed or logged unless abnormal conditions exist.
- Whenever an environment threshold in the filer is crossed.
- When you enter the command from the command line.

  You run this command manually to monitor the filer subsystems, especially when you suspect a problem and when reporting abnormal conditions to Network Appliance.

All environment information

The `environment status` command with no options displays all filer environment information.

Shelf environment status

The `environment status shelf` command displays the shelf environmental status for all shelves. The `environment status shelf` *adapter* command displays shelf information for shelves attached to a specific adapter.

Chassis environment status

The `environment chassis` command displays the environmental status of all chassis components.

Backplane environment status

The `environment chassis backplane` command displays the environmental status of the filer backplane.

CPU environment status

The `environment chassis cpu` command displays the environmental status of the filer CPUs.

Power environment status

The `environment chassis power` command displays the environmental status of the filer power supplies.

Temperature environment status

The `environment chassis temp` command displays the environmental status of the filer chassis temperature.

Command for monitoring UPS devices

The `ups` command allows the filer's environmental monitoring software to monitor uninterruptible power supply (UPS) devices attached to the filer. When power loss is first detected from a UPS device, a syslog warning message and SNMP trap are issued. After a UPS device switches to battery power, subsequent syslog messages and SNMP traps are generated when there is less than five minutes of battery power left (warning state) and again when there is less than one minute of battery power left (critical state). When a filer enters the critical state, it automatically performs an orderly shutdown before power is completely lost.

You can add UPS devices to be monitored, enable or disable monitoring of UPS devices, or display the status of UPS devices.

UPS status

The `ups status` command displays the status of all UPS devices.

Adding a UPS device to be monitored

To add a UPS device to be monitored by the filer's environmental monitoring software, complete the following step.

Step	Action
1	Enter the following command: **`ups add [ -c community ] IP address`** Use -c *community* if you want to specify the community for the UPS device. *IP address* is the IP address of the UPS device.

Enabling or disabling monitoring of UPS devices

To enable or disable monitoring of one or more UPS devices, complete the following step.

Step	Action
1	**`ups [ disable | enable ] [ all | IP address ]`** *IP address* is the IP address of a specific UPS device you want to disable or enable. **Example:** The following command disables monitoring of all UPS devices. `ups disable all` **Note** The `ups enable all` command does not enable previously disabled UPS devices.

Fibre Channel information

Command for displaying Fibre Channel information

You use the `fcstat` command to display the following types of information:

- Link statistics for all disks on a loop
- Internal statistics kept by the Fibre Channel driver
- Information showing the correlation of disks with disk shelves and the relative physical positions of drives on a loop

For more information about the `fcstat` command, see the na_fcstat(1) man page.

Link statistics for all disks on a loop

The `fcstat link_stats` command displays link statistics for disks on a loop. This display includes the link failure count, the loss of sync count, the loss of signal count, the invalid CRC count, the frame in count, and the frame out count.

Internal Fibre Channel driver statistics

The `fcstat fcal_stats` command displays internal statistics kept by the Fibre Channel driver. The Fibre Channel driver maintains statistics about various error conditions, exception conditions, and handler code paths executed.

Relative physical drive position

The `fcstat device_map` command displays the relative physical positions of drives on a loop and the mapping of devices to disk shelves.

NFS and RPC statistics

Command for displaying NFS and RPC statistics

You use the `nfsstat` command to display 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 of the nfsstat command

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

```
nfsstat [ interval ] | [ -h [ ip_address | host_name ] ] | [ -l ] |
[ -z ] | [ -c ] [ -t ]
nfsstat -h [ ip_address | host_name ]
nfsstat -l
nfsstat -z
nfsstat -c
nfsstat -t
nfsstat -d
```

Options of the nfsstat command

The following table describes the options of the `nfsstat` command.

Option	Description
None	When you do not specify an option, the command displays statistical information gathered since the last time the filer was rebooted or the statistics counters were reset.
interval	When you specify an interval, the command displays statistics continually. The interval specifies the number of seconds the command waits between updates.

Option	Description
`-h`	Displays statistics on a per-client basis. You can 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. If no argument is provided, the `-h` option displays statistics for all known clients.
`-l`	Reports the number of NFS requests that have been received from each NFS client.
`-z`	Resets the NFS event counters.
`-c`	Displays reply cache statistics.
`-t`	Displays statistics since boot time rather than since the last time they were zeroed.
`-d`	Displays reply cache statistics and statistics about incoming messages and allocated mbufs.

When to enable per-client statistics

Enable the per-client statistics collection mode (`nfsstat -h`) before you start the filer, preferably by using the `nfs.per_client_stats.enable` option in the /etc/rc file. Otherwise, `nfsstat -l` reports incorrect percentages and displays statistics that include clients that have generated RPC calls but no NFS calls to the client. You can use `nfsstat -z` to clear this inconsistency, but be aware that doing so will clear all NFS statistics on the filer.

Enabling the nfs.per_client_stats.enable option

To enable the `nfs.per_client_stats.enable` option, complete the following step.

Step	Action
1	Enter the following command: **`options nfs.per_client_stats.enable on`**

Example with no options

The following lines show the output of the `nfsstat` command when you specify no options and the statistics on the filer have been reset after use of the `nfsstat -z` command:

```
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%

Read request stats (version 2)
0-511    512-1023   1K-2047   2K-4095   4K-8191   8K-16383
0        0          0         0         0         0
16K-32767  32K-65535  64K-131071
0          0          0

Write request stats (version 2)
0-511    512-1023   1K-2047   2K-4095   4K-8191   8K-16383
0        0          0         0         0         0
16K-32767  32K-65535  64K-131071
0          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%

Read request stats (version 3)
0-511    512-1023   1K-2047   2K-4095   4K-8191   8K-16383
0        0          0         0         0         0
16K-32767  32K-65535  64K-131071
0          0          0

Write request stats (version 3)
0-511    512-1023   1K-2047   2K-4095   4K-8191   8K-16383
0        0          0         0         0         0
16K-32767  32K-65535  64K-131071
0          0          0

Server nfs V4: (24 calls)
access         close   commit    create   delegpurge  delegret
0 0%           0 0%    0 0%      0 0%     0 0%        0 0%
getattr        getfh   link      lock     lockt       ocku
3 0%           3 0%    0 0%      0 0%     0 0%        0 0%
lookup         lookupp nverify   open     openattr    open_confirm
6 0%           0 0%    0 0%      0 0%     0 0%        0 0%
open_downgrade putfh   putpubfh  putrootfh read       readdir
0 0%           0 0%    0 0%      3 0%      0 0%       0 0%
readlink       remove  rename    renew    restorefh   savefh
0 0%           0 0%    0 0%      0 0%     0 0%        0 0%
secinfo        setattr setclntid setclntid_cfm verify write
0 0%           0 0%    9 0%      0 0%          0 0%   0 0%
```

Example with 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 with 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%

Server nfs V4: (24 calls)
access          close    commit   create  delegpurge  delegret
0 0%            0 0%     0 0%     0 0%    0 0%        0 0%
getattr         getfh    link     lock    lockt       locku
3 0%            3 0%     0 0%     0 0%    0 0%        0 0%
lookup          lookupp  nverify  open    openattr    open_confirm
6 0%            0 0%     0 0%     0 0%    0 0%        0 0%
open_downgrade  putfh    putpubfh putrootfh read      readdir
0 0%            0 0%     0 0%     3 0%      0 0%      0 0%
readlink        remove   rename   renew   restorefh   savefh
0 0%            0 0%     0 0%     0 0%    0 0%        0 0%
secinfo         setattr  setclntid setclntid_cfm verify write
0 0%            0 0%     9 0%     0 0%           0 0%   0 0%
```

Example with the -c option

The following command displays similar information as `nfsstat` with no options, as well as the reply cache statistics shown in this example:

```
nfsstat -c
```

```
NFS reply cache statistics:
TCP:
In progress   Delay hits   Misses   Idempotent   Non-idempotent
0             0            0        0            0
UDP:
In progress   Delay hits   Misses   Idempotent   Non-idempotent
0             0            0        0            0
```

The -t option

The following command displays output since the last system boot:

```
nfsstat -t
```

Example with the -d option

The following command displays information similar to that displayed by the `nfsstat` command with no options; it also displays statistics about reply cache, incoming messages, and allocated mbufs, as shown in this example:

```
nfsstat -d
```

```
nfs cache size=49152, hash size=131072

num msg=0, too many mbufs=0, rpcErr=0, svrErr=0
no msg queued=0, no msg paused=0, no msg dropped=0
no msg unqueued=0, no msg discarded=0
sbfull queued=0, sbfull unqueued=0, sbfull discarded=0
no mbuf queued=0, no mbuf dropped=0
no mbuf unqueued=0, no mbuf discarded=0
(cumulative) active=0/491 req mufs=0
tcp input flowcontrol receive=0, xmit=0
```

CIFS statistics

Command for displaying 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.

Syntax of the cifs stat command

```
cifs stat [interval] [ -u user ] [ -h host ] [ -v[v] ]
cifs stat -c
cifs stat -z
```

Options of the cifs stat command

The following table describes the options of the `cifs stat` command.

Option	Description
None	If you do not specify any options, the command displays CIFS statistics that have accumulated since the last reboot.
interval	If you specify a time interval, the command displays CIFS statistics continuously. The interval specifies the number of seconds the command waits between updates.
`-u` *user*	Displays statistics for one or more user accounts. The reported values are the sum of all clients matching the specified user accounts. To use the `-u` option, you must have the `cifs.per_client_stats.enable` option turned on.
`-h` *host*	Displays statistics for one or more hosts. The reported values are the sum of all clients matching the specified hosts. To use the `-h` option, you must have the `cifs.per_client_stats.enable` option turned on.

Option	Description
`-v[v]`	If you are using the `-u` or `-h` option to display per-client statistics, you might want to use `-v` to display the number of clients for which statistics are being reported. `-vv` displays the actual clients as well as the number of clients. To use the `-v[v]` option, the `cifs.per_client_stats.enable` option must be turned on.
`-c`	Displays both blocking and nonblocking CIFS operations. The default is to show only blocking operations. You cannot use this option in conjunction with any of the options for displaying per-client statistics.
`-z`	Resets the operation counters to zero.

Example of cifs stat output with a specified interval

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 first line contains the statistics accumulated since the last boot. 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

Field	Description
`Other`	Other operations, such as deletes and logoffs

Troubleshooting and performance statistics

The following table describes additional `cifs stat` output fields that are useful for troubleshooting and performance management.

Field	Description
`Dup Disconnect`	This variable counts the number of times that the Server Message Block (SMB) layer issues a session disconnect. Usually this indicates that the session termination was initiated by the filer.
`FidHashAlloc`	If a client tries to open more than the threshold number of files, Data ONTAP allocates a new hash table for the new file IDs. This variable counts the number of times this occurs.
`FoldAttempts`	The number of times Data ONTAP attempts to fold a file with the version in a snapshot. When a file is updated, typically only a subset of the file changes. So, only the changed parts of the file are actually written to disk if folding is enabled.
`FoldDuplicates`	The number of duplicate folds on the fold queue.
`FoldOverflows`	The number of times that the fold queue overflowed.
`FoldRenameFailures`	The number of times that a fold queue entry was unsuccessfully renamed.
`FoldRenames`	The number of times that an entry on the queue of files waiting for folding was renamed.

Field	Description
`FoldWAFLTooBusy`	The number of times that WAFL concurrent folds were exceeded.
`Mapped Null User`	The number of times a “null” or “blank” user was successfully mapped. This mapping is done only for IP addresses specified in the usermap.cfg file.
`mBufWait`	Each SMB request needs one or more mbufs. If an mbuf is not available, the request is not processed until one becomes available—that is, Data ONTAP waits for an mbuf. This variable counts the number of times Data ONTAP waits for mbufs to become available.
`NBT Disconnect`	The number of times that the NetBIOS layer issues a session disconnect. Usually, this indicates that the session termination was initiated by the client. NBT Disconnect can also occur when the filer’s TCP/IP layer loses the connection to the client due to a networking problem between the filer and the client.
`NoSupport`	The number of times an unsupported SMB operation is attempted on the filer.
`OpLkBkDelayedBreak`	When an oplock break is sent too soon after an open request completes, some clients lock up. This variable counts the number of times that Data ONTAP must delay before sending an oplock break. The delay is controlled by the `cifs.oplocks.opendelta` option.
`OpLkBkIgnoredAck`	The number of times the filer ignores an oplock break acknowledgment sent from the client. For example, the filer might ignore an acknowledgment that is late.

Field	Description
OpLkBkNoBreakAck	The number of times there was a timeout due to no oplock break acknowledgment being sent by the client to the filer.
OpLkBkNoBreakAck95	The number of oplock break failures to acknowledge caused by Windows 95 clients.
OpLkBkNoBreakAckNT	The number of oplock break failures to acknowledge caused by Windows NT clients.
pBlkWait	A pBlk is the primary control block that holds the context for each SMB request that passes through Data ONTAP. If a pBlk is not available, the request is not processed until one becomes available—that is, Data ONTAP waits for a pBlk. This variable counts the number of times Data ONTAP waits for a pBlk to become available.
PDCupcalls	When a user attempts to map to the filer, Data ONTAP passes all information about the user to a domain controller. This variable counts the number of times Data ONTAP calls into domain controllers for user authentication.
SharingErrorRetries	The number of times an open request that conflicts with another open request is retried before Data ONTAP returns a sharing violation.

Field	Description
`Short Msg Prevent`	When the TCP stack in Windows NT 4, under high load with many outstanding operations, encounters a client bug that causes the TCP window to become zero and stay at zero, the session is lost. To prevent this situation, the filer implements the Nagle algorithm, which requires only full TCP packets to be transmitted even if it causes delays in transmission. This variable reports the number of CIFS sessions for which the Nagle algorithm has been used.
`SMB Disconnect`	The number of times the SMB layer issues a session disconnect. Usually, this indicates that the session termination was initiated by the filer.

When to enable per-client statistics

You can enable the filer to collect per-client statistics at any time without restarting CIFS by turning on the `cifs.per_client_stats.enable` option. This option must be turned on to use the `-u` and `-h` options of the `cifs stat` command and to use the `cifs top` command. Turning this option on can noticeably affect filer performance. If this option is turned off, any existing per-client statistics are deleted.

Turning on the cifs.per_client_stats.enable option

To turn on the `cifs.per_client_stats.enable` option, complete the following step.

Step	Action
1	Enter the following command: **`options cifs.per_client_stats.enable on`**

Displaying CIFS client activity

You use the `cifs top` command to display CIFS client activity, such as which clients are generating large amounts of load. Since this command relies on collected per-client statistics, the `cifs.per_client_stats.enable` option must be turned on.

The default output is a sorted list of clients, showing the number of I/Os, number and size of read and write requests, the number of "suspicious" events, and the IP address and user account of each client. A single client may have more than one entry if it is multiplexing multiple users on a single connection.

Syntax of the cifs top command

The syntax of the `cifs top` command is as follows:

`cifs top [-s sort] [-n maxclients] [-a avg] [-v]`

Options for the cifs top command

The following table describes the options for the `cifs top` command.

Option	Description
`-s` *sort*	Specifies how the client statistics are sorted. Possible values for *sort* are the following, which can be abbreviated to their first character: *ops* (*o*)—Sorts by the number of operations per second of any type *reads* (*r*)—Sorts by kilobytes per second of data sent in response to read requests *writes* (*w*)—Sorts by kilobytes per second of data written to the filer *ios* (*i*)—Sorts by the combined total of reads plus writes for each client *suspicious* (*s*)—Sorts by the number of *suspicious* events sent per second by each client
`-n` *maxclients*	Specifies the maximum number of top clients to display. The default is 20.

Option	Description
`-a` *avg*	Specifies how the statistics are averaged for display. Possible values for `avg` are the following, which can be abbreviated to their first character: *smooth* (*s*)—Uses a smooth average weighted towards recent behavior, but takes into account the client's previous history. *now* (*n*)—Uses a one-second sample, taken immediately. *total* (*t*)—Uses the total count of each statistic divided by the total time since sampling started. If the `-v` option is also used, the total count is displayed without division by the total sampling time.
`-v`	Displays detailed statistics.

What suspicious events are

The following events are examples of "suspicious" events:

```
ACCESS_DENIED returned for FindFirst
ACCESS_DENIED returned for Open/CreateFile
ACCESS_DENIED returned for DeleteFile
SUCCESS returned for DeleteFile
SUCCESS returned for TruncateFile
```

Example of the cifs top command

The following lines show the output of the `cifs top` command when you specify `-n` and `-s`.

```
cifs top -n 3 -s w

ops/s   reads       writes      suspect/s   IP           Name
        (n, KB/s)   (n, KB/s)

263     29   215    137   627        0     10.56.10.120 eng\jose

248     27   190    126   619        1     10.56.10.120 eng\jill
```

ops/s	reads (n, KB/s)	writes (n, KB/s)	suspect/s	IP	Name
246	26 195	125 616	19	10.56.12.118	mktg\bob

DAFS statistics

Command for gathering DAFS statistics

You use the `dafs stats` command to gather statistics about DAFS operations on your filer.

For more information on the `dafs stats` command, see the chapter about file access using DAFS in the *File Access Management Guide*.

Command for displaying DAFS status

You use the `dafs status` command to display information about the DAFS service, NICs, and sessions.

For more information on the `dafs status` command, see the chapter about file access using DAFS in the *File Access Management Guide*.

Improving filer performance

About this section

This section describes configuration procedures that might improve your filer's performance.

Balancing NFS traffic on network interfaces

Attach multiple interfaces on the filer to the same physical network to balance network traffic among different interfaces. For example, if two Ethernet interfaces on the filer named toaster are attached to the same network where four NFS 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.

Ensuring reliable NFS traffic by using TCP

With faster NICs and switches, NetApp recommends supporting NFSv2 or NFS v3 protocol over TCP rather than over UDP. NFS v4 is supported over TCP only.

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. Set the `no_atime_update` option to On for each volume to prevent updates. For example:

```
vol options vol2 no_atime_update 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 set at its default, Off, so that the access time is updated.

Improving read-ahead performance

If the file access patterns of your clients are random (nonsequential) and the cache age is less than 3, setting minimal read-ahead to On might improve performance. By default, the filer uses aggressive read-ahead, which enhances sequential access, and is more commonly used by UNIX clients and applications. Use the `sysstat` command to determine the cache age. Set the `minra` option to On for each volume to specify minimal read-ahead. By default, the option is set to Off and the filer does very aggressive read-ahead.

Adding disks to a disk-bound volume

If you have a single-volume filer, use the `sysstat -u` command to determine the fraction of time that the busiest disk is active. If the fraction is greater than 80%, add disks to the volume through the `vol add` command.

Making optimal use of Slot 8

This method of improving filer performance applies to F700-series filers only. The fastest-available PCI slot in an F700-series filer that is not clustered is Slot 8. If your filer uses a 64-bit network interface card, put it into Slot 8 to take advantage of the speed. If the filer does not use a 64-bit network interface card, put the main FC-AL adapter into this slot.

Putting cards into the correct slots

At boot time or when you use the `sysconfig -c` command, you might see messages indicating that expansion cards must be in certain slots. To improve performance, follow the recommendations in the message. See the *System Configuration Guide* at http://now.netapp.com/NOW/knowledge/docs/hardware/syscfg/ for recommended card placement.

Maintaining adequate free blocks and free inodes

If the percentage of free blocks or free inodes falls to less than 10% on any volume, the performance of writes and creates can suffer. Check free blocks and inodes by using the `df` command and `df -i` command, respectively. If the percentage of used blocks is greater than 90%, increase blocks by adding disks or deleting snapshots. If the percentage of free inodes is less than 10%, increase inodes by deleting files or using the `maxfiles` command.

Using oplocks for CIFS filers

Oplocks allow CIFS clients to cache more data locally, reducing traffic to the filer. On a filer running CIFS that is not in a database environment, set oplocks to On by using the `options cifs.oplocks.enable on` command. If the filer is in a database environment, ensure that the oplocks are not set.

Increasing the TCP window size for CIFS

Increasing the TCP receive window size to its maximum setting on both the filer and a CIFS client can improve performance for large transfers, provided that packet loss is not taking place and the client's send window is large.

Caution

Network Appliance recommends that you call Network Appliance Technical Support before changing this value.

The default is 17,520 bytes. The number of bytes must be between 2,920 and 64,240, in multiples of 1,460 bytes. The TCP window size controls the number of TCP messages that can be transmitted between the filer and the CIFS client at one time.

Caution

If you are using DOS, enter an NT4 value of 8,760.

Use the `options cifs.tcp_window_size 64240` command to maximize the TCP window size on a filer running CIFS. Use the `options nfs.tcp.recvwindowsize 64240` command to maximize the TCP window size on a filer running NFS. Change the window size in the Windows registry on a Windows NT client by adding the DWORD value `\\HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters\TcpWindowSize` and set it to 64,240 (0xFAF0 in hexadecimal).

Note

The `cifs.tcp_window_size` and `nfs.tcp.recvwindowsize` options are invisible until you set them. After you set these invisible options, you can view them by entering the `options cifs` or the `options nfs` command.

Backing up by using qtrees

If your filer has multiple tape drives and a volume that has two to four qtrees, you can improve backup rates by running multiple dump commands in parallel, each reading a different qtree and writing to a different tape drive.

Filer Problem Messages A

About this appendix

This appendix presents filer messages that indicate problems, their causes, and what to do about them. Only a small fraction of the filer problem messages is presented in this appendix.

Note

If you receive any problem message not in this appendix, contact Network Appliance Technical Support.

Topics in this appendix

This appendix discusses the following topics:

- "Serious error messages" on page 428
- "Network error codes" on page 429
- "Disk and FC-AL messages from the Storage Health Monitor" on page 434
- "Cluster messages" on page 439

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: version_no
```

Message component	Description
system hung (NS0)	Indicates the panic class of the message and is significant. The actual text of the message depends on the circumstance.
volname	Is the name of the volume.
version_no	Is the version number.

Recording information and getting assistance after a panic message

To record information and get assistance if your system issues a panic message, complete the following steps.

Step	Action
1	Write down the panic message.
2	Call Network Appliance Technical Support immediately.

Network error codes

Error code descriptions and recommended actions

The following table lists some network error codes, descriptions of problems that the error codes point to, and actions that you can take to fix the problems. For additional information, see the *Network Management Guide*.

Error code	Description	Recommended actions
200	Unspecified hardware problem.	Reseat the network interface card. If the problem persists, contact Network Appliance Technical Support.
201	Link not detected.	Complete the following steps until you detect a link: **1.** Ensure that a cable is connecting the network interface card to the network. **2.** Check the cable. **3.** Check the interface card and corresponding switch port setting for a speed mismatch. **4.** Check for a switch or hub port problem. **5.** Reseat the network interface card. **6.** If the problem persists, contact Network Appliance Technical Support.
202	A link on an Alteon Gigabit adapter is not detected because of an autonegotiation mismatch.	Change the interface configuration or peer switch port configuration so that the settings match.

Error code	Description	Recommended actions
203	No link is detected because of a speed mismatch.	Change the interface configuration or peer switch port configuration to match the speed.
204	The interface is not configured up.	Configure the interface up.
205	Duplex mismatch.	Change the interface or peer switch port duplex setting so they match.
206	Link capacity problem.	Upgrade to a faster interface.
207	The interface is not transmitting or receiving.	Complete the following steps: **1.** Pull the network cable out from the network interface card. **2.** Reinsert the cable. **3.** Change the media type using the `ifconfig` command, wait five seconds, and then restore the media type to the previous setting. **4.** If the problem persists, contact Network Appliance Technical Support.
208	Excessive I/O errors.	Complete the following steps: **1.** Reseat the interface card. **2.** Check the cables. **3.** If the problem persists, contact Network Appliance Technical Support.
209	Excessive unsupported protocol packets are being sent to the filer.	The problem is not with your filer. Contact your network administrator to resolve the problem.

Error code	Description	Recommended actions
301	The IP address and the netmask are inconsistent with the assigned broadcast address.	Change the configuration using the `ifconfig` command.
302	The broadcast address reaches a larger set of hosts than the standard broadcast computed from the IP address and netmask.	If this behavior is erroneous, change the configuration.
303	There are excessive IP reassembly errors.	Switch from NFS over UDP to NFS over TCP.
401	The TCP window advertised by the client is too small.	The problem is not with your filer. Reconfigure the client.
402	There is excessive packet loss on the sending side.	The problem is not with your filer. Examine the network and the client for congestion.
403	There is excessive packet loss on the receiving side.	The problem is not with your filer. Examine the network and the client for congestion.
404	The average TCP packet size is poor on the receiving side because the network, client, or both are not enabled to support jumbo frames.	The problem is not with your filer. Enable support for jumbo frames in network devices and the client.
405	The average TCP packet size is poor on the receiving side because of a problem with the network, client, or both.	The problem is not with your filer. Examine the network and client for configured MTUs.

Error code	Description	Recommended actions
406	The average TCP packet size is poor on the receiving side because of a client application problem.	The problem is not with your filer. Examine the client application data transmission strategy.
407	Excessive TCP listen socket drops because the system is overloaded or under security attack.	Contact your network administrator to resolve the problem.
408	There are excessive filtered TCP port drops because the system is under security attack.	Check your network. Contact your network administrator to resolve the problem.
409	There are excessive embryonic TCP connection drops because the system is under security attack or because a client has a bug.	Contact your network administrator to resolve the problem. A packet trace might assist in locating the problem.
410	Excessive TCP checksum errors because of a client or a network that loses packets when loaded.	The problem is not with your filer. ◆ Check your client system for bugs. ◆ Your network might be overloaded. ◆ Contact your network administrator to resolve the problem.
411	There are packets because of a client. Your system might be under a security attack.	The problem is not with your filer. ◆ Check your client system for bugs. ◆ Check for a security attack.
451	There are xcessive UDP checksum errors.	Switch from NFS over UDP to NFS over TCP.
601	The DNS server is not reachable.	Examine the DNS server and the path to the DNS server.

Error code	Description	Recommended actions
602	The NIS server is not reachable.	Examine the NIS server and the path to the NIS server.

Disk and FC-AL messages from the Storage Health Monitor

About Storage Health Monitor messages

The Storage Health Monitor examines FC-AL, disk shelf, and disk events and provides two kinds of syslog messages about the events: urgent and nonurgent. Each message has a corresponding message that is sent through `autosupport` and is available through SNMP. One set of message strings is provided for each event. There is also an informational message.

Responding to a predicted failure event message

A syslog message about a predicted failure event is an urgent message and should be acted on immediately.

Syslog message: `shm: disk has reported a predicted failure (PFA) event: disk` *`XX`*`, serial_number` *`XXXX`*

Autosupport and SNMP message:

`DISK_FAILURE_PREDICTED!!!`

The Autosupport priority is `alert`.

Description of the problem: The disk has reported a Predicted Failure (PFA) Event. The disk's internal error processing and logging algorithm computation results are exceeding an internally set threshold. The disk is bound to fail in a matter of hours.

Action: If you experience a predicted failure event, complete the following step.

Step	Action
1	Replace the disk, as described in Chapter 5, "Disk Management," on page 201.

Responding to an FC-AL link failure message

A syslog message about an FC-AL link failure is an urgent message and should be acted on immediately.

Syslog message: `shm: link failure detected, upstream from disk: id` *`XX`*`, serial_number` *`XXXXX`*

Autosupport and SNMP message:

`FC-AL LINK_FAILURE`

The Autosupport priority is `err`.

Description of the problem: A Fibre Channel disk or cable, if disks are in different disk shelves, might be malfunctioning and might have caused an open loop condition. This results in a loss of sync over 100 milliseconds for a downstream disk that reported it as a link failure.

Action: If you experience an FC-AL link failure, complete the following steps.

Step	Action
Note The following steps can take several hours to complete.	
1	Use the `halt` command to halt the filer.
2	Check the upstream neighbor device of this disk and the cable between the two if they are on different disk shelves.
3	Take out two adjacent disks and any connecting cable.
4	Replace the cable, if any, to complete the loop.
5	Bring the filer up.
6	Use the `disk scrub` command. **Result:** If this link is bad, the syslog message reappears.
7	Replace the cable.
8	Put in one disk.
9	Use the `disk scrub` command. **Result:** If this disk is bad, the syslog message reappears.
10	Put the other disk back in.
11	Use the `disk scrub` command. **Result:** If this disk is bad, the syslog message reappears.
12	Replace the malfunctioning disk.

Responding to an "FC-AL unreliable" message

An "FC-AL unreliable" message is a nonurgent message and does not require immediate action. However, you should take the actions listed in the Action column of the following table. If you need assistance in performing the actions, contact Network Appliance Technical Support.

Syslog message: `shm: possible link errors on disk: id XX, serial number XXXXX.`

Autosupport and SNMP message:

`FC-AL UNRELIABLE`

The Autosupport priority is `warning`.

Description of the problem: One of a group of four Fibre Channel disks in a disk shelf or any connecting cable might be malfunctioning.

This results in a large number of invalid_CRC frames and data underruns on the loop. The invalid_CRC and underrun count has crossed its set threshold several times.

Action: If you receive an "FC-AL unreliable message," complete the following steps.

Step	Action
Note The following steps can take several hours to complete.	
1	Use the `halt` command to halt the filer.
2	Take out the four disks in the group and any connecting cables that are reported as problems in the syslog messages.
3	Replace the cable, if any, to complete the loop.
4	Bring the filer up.
5	Use the `disk scrub` command. **Result:** If this link is bad, the syslog message reappears.
6	Put in one disk.
7	Use the `disk scrub` command. **Result:** If this disk is bad, the syslog message reappears.
8	Put the disks and cables back in, one at a time, performing Steps 3 through 7 to find the malfunctioning item.
9	Replace the malfunctioning item.

Responding to a disk degraded message

A disk degraded message is a nonurgent message and does not require immediate action. However, you should take the action listed in the Action column of the following table. If you need assistance in performing the action, contact Network Appliance Technical Support.

Syslog message: `shm: disk I/O Completion times too long: disk XX, serial number XXXXX.`

Autosupport and SNMP message:

`DISK_DEGRADED-I/O`

The Autosupport priority is `warning`.

Description of the problem: Either the disk is old and slow, or it is internally recovering errors, thereby taking too long to complete an I/O. The message also indicates that there are too many I/O timeouts and retries on a disk. The disk might also be returning a Command Aborted status quite often. All these can result in a low data-throughput rate for that disk, and lowered overall system performance.

Action: If you receive a disk degraded message, complete the following step.

Step	Action
1	The disk is failure-prone; you should replace it as described in Chapter 5, “Disk Management,” on page 201.

Responding to a disk recovered errors message

A disk recovered errors message is a nonurgent message and does not require immediate action. However, you should take the action listed in the Action column of the following table. If you need assistance in performing the actions, contact Network Appliance Technical Support.

Syslog message: `shm: disk returns excessive recovered errors: disk XX, serial number XXXXX.`

Autosupport and SNMP message:

`DISK_RECOVERED_ERRORS`

The Autosupport priority is `warning`.

Description of the problem: Either the disk has found medium or hardware errors (unrecovered errors), or it has internally recovered numerous errors. The disk might also be returning a Command Aborted status. Quite often the errors returned exceed by many times the disk-vendor-specified bit error rate.

Action: If you receive a disk degraded message, complete the following step.

Step	Action
1	The disk is failure-prone; you should replace it as described in Chapter 5, “Disk Management,” on page 201.

Responding to a loop intermittent instability message

A loop intermittent instability message is an informational message and does not require any action.

Syslog message: `shm: intermittent instability on the loop that is attached to Fibre channel adapter: id` *`XXX`*`, name` *`XXXXX`*`. NOTE: no immediate action needed.`

Autosupport and SNMP message: None

Description of the problem: A Fibre Channel adapter, attached disk shelf, disk, cable, or connectors might cause instability on the FC-AL loop. It results in I/O completion rates that are below a set threshold.

Cluster messages

What to do when you see a cluster message that indicates trouble

If a filer or a cluster ceases to function properly, save or record as many messages as you can. This section lists messages that you can do something to address. If you get a message that is not listed in this section, contact Network Appliance Technical Support.

Detailed information

This section lists the following kinds of cluster messages, what they mean, and what you can do about them:

- "LCD messages" on page 440
- "Syslog messages" on page 442
- "Console messages" on page 445

LCD messages

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

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.	Complete 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.
`Giveback`	The clustered filer has given back control of its partner's disk shelves and is waiting for the partner to boot.	Boot the partner.
`Iconnect down!`	The cluster interconnect link is broken.	Complete 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.

Message	Explanation	What you do
`Iconnect error`	One of the cluster interconnect link cables is broken.	Complete these steps: **1.** Check the interconnect cables to make sure that they are 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, and then the `cf takeover` command.
`takeover`	The filer has taken over its partner.	Perform a giveback.
`takeover in` *`nnn`*	The filer is counting down the seconds until it takes over its partner.	Unless you want takeover to occur, issue the `cf disable` command.

Failed filer LCD message

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.	Complete these steps: **1.** Perform maintenance on the filer. **2.** Perform a giveback from the takeover filer.

Syslog messages

Where syslog messages appear

Syslog messages appear by default in the /etc/messages file. For additional information about syslog messages, see "Configuring message and event logging" on page 90.

Troubleshooting syslog messages

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.	Complete these steps: **1.** Install the Java .zip files from the distribution CD-ROM onto the /etc/java directory of your filer. **2.** Reboot.
`Interconnect is down`	The cluster interconnect is not working properly. Takeover capability is disabled.	Complete 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, and then contact Network Appliance Technical Support.
`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.	Complete 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.

Message	Explanation	What you do
`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 connected properly.
`Volume volume is missing a data disk and has dirty parity.`	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. Complete these steps: **1.** Wait until one of these events happen: ❖ Parity is clean. ❖ The data disk is reconstructed. **2.** Use the `sysconfig -r` command to verify that one of the events in Step 1 has happened. **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.	Complete 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 appears frequently, gather system utilization data (`sysstat`, profiling data, and so on), and then contact Network Appliance Technical Support 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 merits 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 merits investigation. **Note** This circumstance can also occur temporarily during the booting of a clustered filer.	Check the Y (bottom) cable connections to make sure that the cables are seated properly.

Console messages

About console messages

Some messages appear on the console but are not written to the /etc/messages file. Most of them are self-explanatory.

Troubleshooting a 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 its partner appears to be 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`	Complete 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 of ensuring 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. ❖ Enter **y** to continue and override the takeover and allow the local filer to reboot. **Caution** If you enter **y** and override the takeover, client applications might lose data that was written before the reboot.

Booting the Filer B

About this appendix

This appendix describes how to boot a filer from the system boot diskette set, from the PC CompactFlash card on F87, FAS270, and FAS270c filers, and from remote images on supported filers.

Topics in this appendix

This chapter discusses the following topics:

- "Booting from the system boot diskette set" on page 448
- "Booting the F87 filer" on page 451
- "Booting the FAS200 series appliances" on page 453
- "Using filers as netboot servers" on page 458

Booting from the system boot diskette set

Reasons to boot from diskette

You might need to reboot the filer from the system boot diskette under the following circumstances:

- To correct configuration problems
- To recover from a lost password
- To correct certain disk configuration problems

Restrictions on booting from diskette

Do not boot the filer with pre-6.1 system boot diskettes. The attempt to boot will fail because pre-6.1 software cannot recognize system data from a filer running Data ONTAP 6.1 or later.

Booting from a diskette

To boot from a diskette, complete the following steps.

Step	Action	
1	Insert the diskette labeled System Boot Diskette 1 into the filer's diskette drive. **Note** Data ONTAP system boot diskettes are specific to the filer model. 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	**If....**	**Then...**
	Your filer is turned on	From the system console, enter the following command: `reboot` **Result:** The filer begins the boot process.
	Your filer is turned off	Turn on your filer. **Result:** The filer begins the boot process.

Step	Action
3	When the filer's LCD prompts you, remove the diskette and insert the next diskettes into the filer's diskette drive when you are prompted.

4	Press Enter.
	
	Result: The filer boots and, if it can, displays the following boot menu:
	
	`1) Normal Boot`
	`2) Boot without /etc/rc`
	`3) Change Password`
	`4) Initialize all disks`
	`5) Maintenance mode boot`
	
	`Selection (1-5)?`

<table>
<tr><th>Step</th><th>Action</th></tr>
<tr><td>5</td><td>Choose one of the boot types shown in the following table by entering the corresponding number.
<table>
<tr><th>Task</th><th>Boot type</th></tr>
<tr><td>Run the filer normally, but from a diskette.</td><td>Normal boot (1)</td></tr>
<tr><td>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 <code>syslog</code>.</td><td>Boot without /etc/rc (2)</td></tr>
<tr><td>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 <code>halt</code> command.</td><td>Maintenance mode boot (5)</td></tr>
</table></td></tr>
</table>

Booting the F87 filer

Boot options for the F87 filer

The F87 filer boots Data ONTAP from a 32 MB PC CompactFlash card. It does not boot from a set of floppy disks. When you turn on or reboot the F87, the filer console displays a basic boot menu that enables you to perform the following tasks:

- Boot Data ONTAP

  This is the Data ONTAP software release stored in CompactFlash card. By default, the filer automatically boots this release if you do not select an option from the basic menu.

- Use advanced service options

  Advanced service options enable you to perform the following tasks:

  - Run the Network Appliance Diagnostics program
  - Boot the filer from one of its SCSI disks

Note

You boot the filer from a SCSI disk only when you install a new NVRAM adapter in the filer. Do not use this option for normal operation of the filer. For detailed information, see the *F87 Hardware and Service Guide*.

Booting Data ONTAP

To boot the current release of Data ONTAP or the backup release, complete the following steps.

Step	Action

1	Turn on or reboot the filer.
	
	Result : The filer console displays power on self test (POST) messages. After POST successfully executes, the system console displays the following boot menu.
	
	`Data ONTAP`
	
	`Service menu`

Step	Action
2	Select the Data ONTAP option. **Note** If you do not select an option within five seconds, the filer boots the current release of Data ONTAP.

Booting the FAS200 series appliances

Boot options for the FAS200 series appliances

The FAS200 series (FAS250, FAS270, and FAS270c) appliances automatically boot Data ONTAP from a PC CompactFlash card. They do not boot from a set of floppy diskettes or from disk. The CompactFlash card ships with the most current Data ONTAP release and a diagnostic kernel in a single FAT partition. The CompactFlash card contains sufficient space for an upgrade kernel.

Note

FAS200 series appliances can be upgraded to the most recent Data ONTAP release. When you install new software, the `download` command copies a boot kernel to the CompactFlash card, not to a boot block on disk, as in other filer models. For more information, see the *Data ONTAP Upgrade Guide*.

You can also enter one of the following boot options at the CFE (Common Firmware Environment) prompt:

- `boot_ontap`

  Boots the current Data ONTAP software release stored on the CompactFlash card. By default, the filer automatically boots this release if you do not select another option from the basic menu.
- `boot_primary`

  Boots the Data ONTAP release stored on the CompactFlash card as the primary kernel. This option overrides the firmware AUTOBOOT_FROM environment variable if it is set to a value other than PRIMARY. By default, the `boot_ontap` and `boot_primary` commands load the same kernel.
- `boot_backup`

  Boots the backup Data ONTAP release from the CompactFlash card. The backup release is created during the first software upgrade to preserve the kernel that shipped with the filer. It provides a “known good” release from which you can boot the filer if it fails to automatically boot the primary image.
- `netboot`

  Boots from a Data ONTAP version stored on a remote HTTP or TFTP server. Netboot enables you to

  - Boot an alternative kernel if the CompactFlash card becomes damaged
  - Upgrade the boot kernel for several devices from a single server

  To enable Netboot, you must configure networking for the FAS2x0 series appliance (using DHCP or static IP address) and place the boot image on a configured server.

Note

Network Appliance recommends that to protect against data loss in the event of PC CompactFlash card corruption you enable Netboot immediately for your FAS200 system by placing a boot image on a local server. For more information, see "Preparing an FAS200 series filer for Netboot startup" on page 455.

- boot_diags

  Boots a Data ONTAP diagnostic kernel. For more information, see the *NetApp Hardware Diagnostics Guide*.

Note

Other boot options should be used only under the direction of Network Appliance technical staff.

Booting Data ONTAP at the firmware prompt

To boot the current release of Data ONTAP or the backup release at the firmware prompt, complete the following steps.

Step	Action
1	At the filer console, enter the following command: `halt` **Result:** The filer console displays the CFE> firmware prompt.
2	Enter one of the following commands: ◆ **`boot_ontap`**—Boots the current release of Data ONTAP. ◆ **`boot_primary`**—Boots the Data ONTAP primary kernel. ◆ **`boot_secondary`**—Boots the Data ONTAP backup kernel.

Note

For more information about CFE commands, enter **`help`** at the firmware prompt for a list of commands or **`help`** ***`command`*** for command details.

Checking available Data ONTAP versions

By default, an FAS200 series filer boots the current Data ONTAP release from the primary kernel. You might need to check the current booted kernel and other kernels available on the CompactFlash card if the system was netboot started from an unfamiliar system, if an upgrade was unsuccessful, or if you need to run kernel diagnostics.

To determine...	At the filer console, enter...
The current booted Data ONTAP version	`version` **Result:** The console displays the version number of Data ONTAP that is currently running.
Data ONTAP versions available on the CompactFlash card	`version -b` **Result:** The console displays the contents of the CompactFlash FAT file system, including name and version information for the primary, secondary (if present), and diagnostic kernels, and the firmware.

For more information, see the na_version(1) manual page.

Preparing an FAS200 series filer for Netboot startup

To use the Netboot option to start your FAS200 series filer, you must have

- An HTTP or TFTP server available on your network.
- A boot image on a server. The boot image can be copied from the filer boot directory, /etc/boot/netapp-mips, or downloaded from the NOW site.
- Networking configured in the FAS200 series firmware environment. You can use a DHCP server to obtain an IP address, or you can configure the network connection manually.

Note

Network interfaces configured in the firmware environment are not persistent across system halts and reboots. You must configure the network interface each time you use the Netboot option.

Two network interfaces, e0a and e0b, are available in the firmware environment. Only one can be configured at a time.

Starting the FAS200 series filers through the netboot option

To boot Data ONTAP from a remote image, complete the following steps.

Step	Action	
1	At the filer console, enter the following command: **`halt`**	
2	**If you are...**	**Then...**
	Configuring DHCP	Enter the following command at the CFE> prompt: **`ifconfig e0a -auto`**
	Configuring manual connections	Enter the following command at the CFE> prompt: **`ifconfig e0a -addr=`*`filer_addr`* **`-mask=`*`netmask`* **`-gw=`*`gateway`* **`-dns=`*`dns_addr`* **`-domain=`*`dns_domain`* *filer_addr* is the IP address of the filer. *netmask* is the network mask of the filer. *gateway* is the gateway for the filer. *dns_addr* is the IP address of a name server on your network. *dns_domain* is the DNS domain name (that is, if you use this optional parameter, you do not need a fully qualified domain name in the netboot server URL; you need only the server's host name). **Note** Other parameters might be necessary for your interface. Enter **`help ifconfig`** at the firmware prompt for details.

Step	Action

3	At the CFE> prompt, enter the following command:
	`netboot` ***`URL`***
	URL is the location of the remote boot image. It can be either an HTTP or a TFTP network path.
	Examples:
	`netboot http://myserver/bootimages/netapp/netapp-mips`
	`netboot tftp://myserver/bootimages/netapp/netapp-mips`
	Note
	The location and availability of boot images depend on the correct configuration of your netboot server.

Using filers as netboot servers

Netboot server requirements

You can configure a filer to serve boot images to other Network Appliance devices that support netboot. To do so, you must configure

- HTTP services, TFTP services, or both on the filer.
- The rest of your netboot using environment to access the filer as the netboot source. For example, you might configure BOOTP, DHCP, bootparamd, and/or rarpd, depending on the specific procedure you are using.

Serving boot images from the filer /etc/http directory: You can also serve netboot images from your filer by placing them in the /etc/http directory. The contents of this directory are served by default to enable the FilerView graphical management interface. For example, if you create an /etc/http/boot directory and place a boot image called netapp-mips in that directory, the boot image will be available for netboot startup at the following URL:

http://your.filer.com/na_admin/boot/netapp-mips

Configuring HTTP services

To configure a filer as an HTTP netboot server, complete the following steps.

Step	Action
1	Place boot programs in the filer /etc/http directory.
2	At the command line of the source filer, enter the following command: `options httpd.enable on`

Configuring TFTP services

To configure a filer as a TFTP netboot server, complete the following steps.

Step	Action
1	Place boot programs in the filer /etc/tftpboot directory. **Note** Boot programs are not required to be in the /etc/tftpboot directory, but they must be in the directory that is set to be `tftpd.rootdir`.

Step	Action
2	On the command line of the source filer, enter the following command: `options tftpd.enable on`

Specifying the TFTP root directory

Any path name specified on the TFTP command line is considered to be relative to the TFTP root directory. TFTP access using absolute path names succeeds only if the specified files are located in the file system under this directory. The default value of the `tftpd.rootdir` option is /etc/tftpboot.

To specify a TFTP root directory other than /etc/tftpboot, complete the following step.

Step	Action
1	On the command line of the source filer, enter the following command: `options tftpd.rootdir` ***`pathname`*** **Note** A valid value for *pathname* is a fully qualified path name to a valid, existing directory on any volume on the filer.

Enabling console logging of TFTP accessed files

To enable console logging of files accessed with TFTP, complete the following step.

Step	Action
1	On the command line of the source filer, enter the following command: `options tftpd.logging on`

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 mail host for sending Autosupport e-mail messages. At any time after you run the `setup` program, you can configure the filer to work with other clients in the same way it does with the administration host.

agent A Data ONTAP process that gathers status and diagnostic information and forwards it to NMSs.

appliance A device that performs a well-defined function and is easy to install and operate.

ATM Asynchronous transfer mode. 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 e-mail messages from the customer site to Network Appliance or another specified e-mail 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.

cluster A pair of filers connected so that one filer can detect when the other is not working and, if so, can serve the failed filer 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 through the `cf` command.

community A name used as a password by the SNMP manager to communicate with the filer agent.

console A terminal that is attached to a filer's serial port and is used to monitor and manage filer operation.

continuous media scrub A background process that continuously scans for and scrubs media errors on the filer disks.

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 a RAID 4 array, when one or two disks are missing from a RAID DP array, or when the batteries on the NVRAM card are low.

disk ID number A number assigned by a filer to each disk when it probes the disks at boot time.

disk sanitization A multiple write process for physically obliterating existing data on specified disks in such a manner that the obliterated data is no longer recoverable by known means of data recovery.

disk shelf A shelf that contains disk drives and is attached to a filer.

emulated filer A software copy of a failed filer that is hosted by its takeover filer. The emulated filer appears to users and administrators to be a functional version of the failed filer. For example, it has the same name as the failed 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 into which you insert expansion cards.

failed filer A physical filer that has ceased operating. In a clustered configuration, it remains the failed filer until giveback succeeds.

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.

giveback The return of identity from the virtual filer to the failed filer, resulting in a return to normal operation; the reverse of takeover.

group A group of users defined in the filer's /etc/group file.

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.

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. LANE enables ATM-connected end systems to communicate with other LAN-based systems.

local filer The filer you are logged in to.

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.

mailbox disk One of a set of disks owned by each filer that is used to store the cluster state information of a 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 file system disks.

mail host The client host responsible for sending automatic e-mail 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.

MIB Management information base. An ASCII file that describes the information that the agent forwards to NMSs.

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 start the application that is appropriate for the data received.

MultiStore An optional software product that enables you to partition the storage and network resources of a single filer so that it appears as multiple filers on the network.

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 ATM adapter.

network management station See NMS.

NMS Network Management Station. A host on a network that uses a third-party network management application (SNMP manager) to process status and diagnostic information about a filer.

normal mode The state of a filer when there is no takeover in the cluster.

NVRAM cache Nonvolatile RAM in a 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 An adapter card that contains the filer's NVRAM cache.

NVRAM mirror A synchronously updated copy of the contents of the filer NVRAM (nonvolatile random access memory) contents kept on the partner filer.

panic A serious error condition causing the filer to halt. Similar to a software crash in the Windows system environment.

parity disk The disk on which parity information is stored for a RAID 4 disk drive array. In RAID groups using RAID DP protection, two parity disks store the parity and double-parity information. Used to reconstruct data in failed disk blocks or on a failed disk.

partner From the point of view of a local filer, the other filer in a cluster.

partner mode The method you use to communicate through the command-line interface with a virtual filer during a takeover.

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 a 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 special subdirectory of the root of a volume that acts as a virtual subvolume with special attributes.

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 an array. NetApp filers use either RAID Level 4, which stores all parity information on a single disk, or RAID DP, which stores all parity information on two disks.

RAID disk scrubbing The process in which a 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 SCSI disk drives and tape drives.

SCSI address The full address of a disk, consisting of the disk's SCSI adapter number and the disk's SCSI ID, such as 9a.1.

SCSI ID The number of a disk drive on a SCSI chain (0 to 6).

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 online, read-only copy of an entire file system that protects against accidental deletions or modifications of files without duplicating file contents. Snapshots enable users to restore files and 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 a filer's CPU, expansion bus slots, and system memory.

takeover The emulation of the failed filer identity by the takeover filer in a cluster; the opposite of *giveback*.

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 disk shelves and network connections. The takeover filer maintains its own identity and the virtual filer maintains the failed filer identity.

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 takeover mode.

trap An asynchronous, unsolicited message sent by an SNMP agent to an SNMP manager indicating that an event has occurred on the filer.

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.

vFiler A virtual filer you create using MultiStore, which enables you to partition the storage and network resources of a single filer so that it appears as multiple filers on the network.

VGA adapter Expansion card for attaching a VGA terminal as the console.

volume A file system.

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

(useradmin userdel), deletion of administrative user account 30
(useradmin userlist), list of administrative user accounts 27
/etc directory, configuration files in 60
/etc/group file
    accessing from UNIX/NFS clients 260
    accessing from Windows client 259
    default entry 258
    editing 264
    purpose of 256
/etc/hosts.equiv file
    rules for entries 23
    specifing hosts 23
    specifying hosts in 23
/etc/log/auditlog file 94
/etc/messages file 90
    disk failure activity log 169
    disk failures contained in 171
    restart scheduling of 90
/etc/netgroup file, purpose of 256
/etc/nsswitch.conf file 268
/etc/passwd file
    accessing from UNIX/NFS clients 260
    accessing from Windows client 259
    adding entries to 261
    default entry 257
    editing 263
    editing restrictions 261
    purpose of 256
/etc/quotas file
    character coding 348
    description of 334
    Disk field specified in 350
    entries for mapping users 356
    file format 346
    Files field specified in 351
    Soft Disk field specified in 353
    Soft Files field specified in 353
    target specified in 349
    Threshold field specified in 352
    Type field specified in 350
/etc/rc file
    default contents of 64
    startup commands in 64
/etc/shadow file, purpose of 257
/etc/syslog file, messages in 442, 444
/etc/syslog.conf file 90
    configuring message logging in 93
    file format and parameters of 91, 92
/etc/usermap.cfg file, character coding of 61
/home file, contents of 59
/vol/vol0, default volume 58

A

access management, data 106
access, native file 141
account (useradmin userdel), deleting of administrative user 30
accounts (useradmin userlist), listing of administrative user 27
address, disk 202
administration hosts
    adding 24
    definition of 22
    deletion of 24
    reasons for using 22
    removing 24
    use of 22
    when you can add 22
    where they are specified 23
administration tasks, periodic 16
administrative level commands 45
administrative user accounts
    changing the password of (passwd) 29, 30
    creating (useradmin useradd) 27
    deleting (useradmin userdel) 30
    listing (useradmin userlist) 27
    management of 25
advanced level commands 45
audit log file, /etc/log/auditlog 94
automatic technical support 47

Autosupport
configuring 53, 56
contents of email 50
defined 47, 48
events that trigger e-mail 49
mail host support for 48
Network Appliance Technical Support and 48
options autosupport.option (configures Autosupport) 56
requirements for 48
testing (options autosupport.doit) 57
Autosupport messages
DDISK_RECOVERED_ERRORS 437
DISK_DEGRADED-I/O 437
DISK_FAILURE_PREDICTED!!! 434
FC-AL LINK_FAILURE 434
FC-AL UNRELIABLE 436

B

backup
for data protection 4
information in other NetApp guides 20
using qtrees 136
block checksums 112, 202
boot
from diskette 448, 450
options for the F87 filer 451, 453
reasons to 448

C

cf giveback command 326
cf-config-check.cgi utility 332
change privileges, file ownership 100
checksum types 112, 202
checksum upgrade process 399
checksums, block 112, 202
checksums, zoned 112, 202
CIFS
accessing UNIX files with 141
client, requirements to manage filer 23
oplocks
changing the settings (options cifs.oplocks.enable) 145
definition of 145
setting for volumes 115
CIFS commands, options cifs.oplocks.enable (enables and disables oplocks) 145
clients
CIFS, requirements 23
editing configuration file from 63
NFS, requirements 23
platforms supported by FilerView 39
cluster commands
cf (enables and disables takeover) 311
cf nfo (enables and disables negotiated failover) 316
cf nfo status (determines status) 317
cf partner (displays partner's name) 311
cf status (monitors status) 308
halt (halts filer without takeover) 313
partner (accesses emulated filer) 320
partner restore (performs restore on emulated filer) 325
sysconfig/vol status (displays disk information) 311
takeover (description of all takeover commands) 314
cluster configuration checker 332
clusters
after a takeover, description of 281
characteristics of 279
checking configuration through a utility 332
command exceptions when accessing emulated filer 322
configuration information for 311
configuring dedicated and standby interfaces 294
configuring dual-MAC interfaces for takeover 297
configuring interfaces for takeover 294
considerations 280
console messages about 445
description of 278
displaying disk information (sysconfig/vol status) 311
displaying partner's name (cf partner) 311
emulated filer
accessing from takeover filer (partner) 320

accessing remotely 322
description of 320
dumps and restores 325
performing restores (partner restore) 325
emulated LAN considerations 297
enabling and disabling takeover (cf) 311
FORE/IP and LANE interface configurations 292
giveback
description of 281
executing 326
mode 289
halting filer without takeover (halt) 313
hardware components described 279
interface types and configurations 292
LCD messages on takeover filer 440
licensing and unlicensing cluster system 313
managing in normal mode 308
messages on failed filer 441
monitoring cluster status (cf status) 308
negotiated failover
description of 316
determining status (cf nfo status) 317
enabling and disabling (cf nfo) 316
negotiated takeover, messages indicating a 287
network interfaces, role of 291
normal mode, description of 283
preparing to configure 280
status messages, descriptions of 309
successful giveback, messages indicating a 290
successful takeover, messages indicating a 286
system failure 281
takeover
configuring with dedicated and hot standby interfaces 296
description of 278
interface configurations 292
mode, description of 285
of interfaces, description of 291
statistics, description of 318
with SNMP 319
trouble indications, description of 439
types of operation modes 282
command line editing
with editing mechanism 43
with history mechanism 43
commands
administrative level 45
advanced level 45
autosupport.option (sets autosupport options) 56
boot (boots the filer) 73
date (sets system date and time) 83
disk assign 249
halt (halts the filer) 71, 72
license
disabling license code 86
displaying current codes 86
enabling license code 85
options autosupport.doit (tests autosupport) 57
options command 68
options raid.reconstruct.perf_impact (modifies RAID data reconstruction speed) 181
options raid.reconstruct_speed (modifies RAID data reconstruction speed) 182
options raid.resync.perf_impact (modifies RAID plex resynchronization speed) 184
options raid.scrub.duration (sets duration for disk scrubbing) 188
options raid.scrub.enable (enables and disables disk scrubbing) 188
options raid.verify.perf_impact (modifies RAID mirror verification speed) 185
options timed (synchronizes system time) 78
options wafl.root_only_chown (sets file ownerships changes) 100
options, overview of options command 68
passwd (changes administrative user password) 29, 30
passwd (changes filer system password) 26
priv 45
privilege levels 45
rdate (synchronizes system time) 79, 80
reboot (reboots the filer) 73
rsh command list 37
savecore, what it does 87
See also CIFS commands, Cluster commands, disk commands, HTTP commands,

NFS commands, quotas commands, RAID commands, stats commands, storage commands
systat (determines CPU load) 179
timed (synchronizes system time with SNTP) 80
to set filer options 67
useradmin useradd (adds administrative users) 27
useradmin userdel (deletes administrative user account) 30
useradmin userlist (lists administrative accounts) 27
vol options 69
common tasks, list of 70
CompactFlash card
booting an F87 filer 451
booting the FAS200 series filers 453
checking the Data ONTAP version of 454, 458
configuration
display, using sysconfig 386
of autosupport (options Autosupport.option) 56
of dedicated and standby interfaces in a cluster 294
of dual-MAC interfaces for a cluster 297
of message and event logging 90
of Telnet 95
configuration files
/etc 59
/etc/group. *See* /etc/group file in the "Symbols" section of this index
/etc/netgroup. *See* /etc/netgroup file in the "Symbols" section of this index
/etc/passwd. *See* /etc/passwd file in the "Symbols" section of this index
/etc/shadow. *See* /etc/shadow file in the "Symbols" section of this index
about backup and cloning 102
accessing
alternative methods 6
from UNIX/NFS clients 260
from Windows client 259
backing up 102
cloning 103
comparing backups 104
editing from CIFS client 63
editing from NFS client 62
editing from NFS client-setup 62
restoring 103
type, location, and purpose 256
within /etc directory 60
continuous media scrub
description 195
disabling 195
specifying length of a cycle 195
core files
description of 87
saving of 87
where they are written 87
CPU load statistics (systat) 179

D

DAFS
dafs stats command, description of 423
dafs status command, description of 423
data access management 14, 106
data disks
description of 202
removing 229
data reconstruction
after disk failure 171
description of 181
data resource, management of 106
date, setting filer time and 83
default directories
accessing 75
of default volume 59
permissions of 59
default volume
description of 58
directories contained within 59
degraded message, disk 437
degraded mode 171, 230
deletion of administration hosts 24
deletion of administrative user account (useradmin userdel) 30
directory structures
default directory permissions 59

default volumes and directories 59
disk assign command
modifying 250
use on the FAS270 and 270c appliances 249
disk commands
df (determines free disk space) 231
df (filer command for accurate disk space data) 232
df (reports discrepancies) 231
df (used with qtrees) 231
disk scrub (starts and stops disk scrubbing) 186, 189
disk show (for viewing disk ownership on FAS270 and FAS270c filers) 247
storage 208
sysconfig -d 202
sysconfig -r (determines number of hot spare disks) 222
disk degraded message 437
disk failures
data reconstruction after 171
RAID reconstruction after 169
without hot spare, results of 171
disk failures, data reconstruction after 171
disk ownership assignment
description 247
modifying 250
viewing ownership 247
disk recovered errors message 437
disk sanitization
description 233
executing 235
licensing 233, 234
limitations 233
log files 244
selective data sanitization 238
stopping 238
stopping the process 235
disk scrubbing
description of 182
enabling and disabling (options raid.scrub.enable) 188
modifying speed of 182
scheduling 186
setting duration (options raid.scrub.duration) 188
starting/stopping (disk scrub) 186, 189
toggling on and off 188
disk shelves 10
disk space, report of discrepancies (df) 231
DISK_DEGRADED-I/O Autosupport and SNMP message 437
DISK_FAILURE_PREDICTED!!! Autosupport and SNMP message 434
DISK_RECOVERED_ERRORS Autosupport and SNMP message 437
disks
adding new to a volume 224
adding new to filer 223
adding to a RAID group other than the last RAID group 226
adding to an unmirrored volume spanning adapters 226
adding to filer 221, 223
addresses 202
assigning ownership of of FAS270 and FAS 270c filers 247
available space on new 203
data and parity 109, 160
data, description of 12
data, removing 229
description of 201, 427
determining number of hot spares (sysconfig) 222
displaying statistical disk information 386
failed, removing 227
formats 202
hot spare, removing 228
kinds of 202
ownership of on FAS270 and FAS270c filers 247
parity 12, 109, 160
reasons to remove 227
rules for adding disks to volume 222
sanitizing 233
spare 12
types supported 202
viewing information about 216
when to add 221
with FC-AL, description of urgent messages

434
double disk failure
avoiding with media error thresholds 198
RAID DP protection against 161

E

editing mechanism, editing command line with 43
effects of oplocks 145
emulated filer
accessing from takeover filer 320
accessing remotely 322
description of 320
dumps and restores 325
performing restores (partner restore) 325
emulated LANs, considerations when in a cluster 297
environmental adapter, filer main unit 9
error codes, network 429
error messages
DISK_DEGRADED-I/O 437
DISK_FAILURE_PREDICTED!!! 434
DISK_RECOVERED_ERRORS 437
FC-AL link failure 434
FC-AL LINK_FAILURE 434
FC-AL UNRELIABLE 436
FC-AL unreliable 435
loop intermittent instability 438
network error codes 429
panic messages 428
predicted failure event 434
serious 428
shm 434
Storage Health Monitor 434
errors message, disk recovered 437
event message, predicted failure 434

F

F87 filer, booting 451
failed disk, removing 227
failover
negotiated, determining status in a cluster 317
negotiated, enabling and disabling of 316
negotiated, in a cluster 316
failover paths
about 300
about enabling 301
configuring addresses for 303
enabling 305
how the client identifies the 303
failure event message, predicted 434
failure message, FC-AL link 434
failure, data reconstruction after disk 171
FAS200 series filers
booting 453
preparing for Netboot startup 455
starting through the netboot option 456
FAS250 filer
booting 453
default RAID 4 group size 176
default RAID DP group size 176
FAS270 filer
assigning disks to 249
booting 453
FAS270c filer
assigning disks to 249
booting 453
FC-AL LINK_FAILURE message 434
FC-AL UNRELIABLE Autosupport and SNMP message 436
Fibre Channel, multipath I/O 205
file access 141
file grouping, using qtrees 136
file ownership change privileges 100
file, /etc/log/auditlog 94
filer access 22
/etc directory, accessing from CIFS client 75
/etc directory, accessing from NFS client 75
/home directory, accessing from CIFS client 76
/home directory, accessing from NFS client 76
default directories, accessing 75
from the console 33
inaccessiblility, reasons for 77
methods 32
terminating a Telnet session 35
using FilerView 39
using rsh command from a PC client 37
using rsh command from a UNIX client 37
with Remote Shell (rsh) 36
with Telnet 34

filers
adding disks to 221, 223
booting the system (boot) 73
common tasks with 70
console or Telnet session to 33
determining number of hot spare disks in (sysconfig) 222
disabling licenses for (license) 86
displaying current license codes for (license) 86
displaying overall statistics for (sysconfig) 387
displaying version of 386
editing boot configuration file in 74
enabling licenses for (license) 85
halting (halt) 72
improving performance in 424
large file sizes in 89
main unit components 8
methods to access 32
panic messages, description of 428
performance improvements
avoiding inodes access time update 424
backup rate 426
balancing NFS traffic on interfaces 424
caching client data to reduce traffic 425
disk-bound volume 425
large transfer 426
maintain adequate free blocks and inodes 425
make use of fastest slot (F700 series only) 425
read-ahead 425
rebooting the system (reboot) 73
rsh (Remote Shell) access to 36
running in degraded mode 171
savecore (saves the core file) 87
security
controlling file ownership changes (options wafl.root_only_chown) 100
controlling mount privileges (options nfs.mount_rootonly) 99
limiting Telnet access 97
options 96
setting date and time (date) 83
setting filer date while running SnapMirror 84
synchronizing filer system time (options timed) 78
synchronizing system time (rdate) 79, 80
synchronizing system time with SNTP (timed) 80
when to add disks 221
FilerView
accessing filer through 39
description 39
Help system defined 41
use for filer administration 7
using the interface 40
Web browser requirements 39
files, configuration 59
filestats command 392

G

giveback
cluster mode, description of 289
description of, in a cluster 281
forcing 327
processes that prevent giveback 326
group file, editing 264

H

history mechanism, editing command line with 43
host adapter
2202 and 2212 205
changing state of 209
Multipath I/O 205
storage command 208
viewing information about 212
host support for Autosupport, mail 48
hosts
definition of 22
deletion of administration 24
use of administration 22
when you can add 22
hot spare disks
description of 202
removing 228
hub, viewing information about 214

I

inode, definition of 152
instability message, loop intermittent 438
interface, use of FilerView 40
intermittent instability message, loop 438

L

large files on the filer
    description of 89
    software requirements 89
LCD
    cluster messages, description of 440
    filer main unit 9
LDAP
    authorizing UNIX client access to NTFS or mixed file systems 275
    authorizing Windows client access to UNIX files 275
    editing /etc/nsswitch.conf file 273
    enabling and disabling (options ldap.enable) 272
    servers 271
    specify base distinguished name (options ldap.base) 271
    specify preferred servers (options ldap.servers.preferred) 272
    specify servers (options ldap.servers) 272
    specifying admin password (options ldap.passwd) 274
    specifying admin user name (options ldap.name) 273
    specifying LDAP port (options ldap.port) 274
    support of Windows client authentication 275
LDAP commands
    options ldap.base (specifies base distinguished name) 271
    options ldap.enable (enables and disables LDAP) 272
    options ldap.name (specifies admin user name) 273
    options ldap.passwd (specifies admin password) 274
    options ldap.port (specifies LDAP port) 274
    options ldap.servers (specifies servers) 272
    options ldap.servers.preferred (specifies preferred servers) 272
level commands
    administrative 45
    advanced 45
licenses
    disabling 86
    displaying current codes 86
    enabling 85
    managing of 85
limitations of qtrees 137
LINK_FAILURE Autosupport and SNMP message, FC-AL 434
list of administrative user accounts (useradmin userlist) 27
load statistics (systat), CPU 179
loop intermittent instability message 438
LUNs, defined 12

M

mail host support for Autosupport 48
management
    of access to a filer 22
    of clusters, in normal mode 308
    of clusters, in takeover mode 318
    of data access 106
    of licenses 85
management, data access 106
maximum number
    of files per volume 152
    of volumes allowed 107
media error failure thresholds 198
message and event logging
    configuring 90
    descriptions of 90
messages
    DISK_DEGRADED-I/O 437
    DISK_FAILURE_PREDICTED!!! 434
    DISK_RECOVERED_ERRORS 437
    FC-AL link failure 434
    FC-AL LINK_FAILURE 434
    FC-AL UNRELIABLE Autosupport and SNMP 436
    loop intermittent instability 438

panic 428
predicted failure event 434
serious error 428
methods to access a filer 32
mirror verification, description of 185
mirrored volumes, definition of 110
mixed security style, description of 139
mode, degraded 171, 230
mount privileges, controlling of (options nfs.mount_rootonly) 99
Multipath I/O
advantages 205
enabling requirements 205
host adapters 205
setup, example 206
understanding 205

N

native file access 141
negotiated failover
determining status (cf nfo status) 317
in a cluster, description of 316
in a cluster, enabling and disabling) 316
Netboot startup
preparing FAS200 series filers 455
starting FAS200 series filers 456
using filers as Netboot servers 458
network error codes 429
Network Information Services (NIS). *See* NIS
network interfaces
cluster configuration for takeover with dedicated and hot standby 296
cluster configurations for 292
configuring dual-MAC for a cluster 297
configuring for takeover, in a cluster 294
emulated LAN considerations in a cluster 297
FORE/IP and LANE in cluster systems 292
role in a cluster system 291
types and configurations in a cluster 292
NFS protocol
accessing Windows files 141
client
editing configuration files from 62
requirements to manage filer 23
NIS
advantages of 266
cached group information, displaying and updating 267
caching group information 266
NIS options. *See* options
ypgroup (displays cached group information) 267
nis.group_update_schedule option 266
normal mode, description of 283
NTFS security style, description of 139
NVRAM
definition 8
halt command to save data to disk 71

O

operation modes, in a cluster 282
oplocks
definition of 145
effects when enabled 145
enabling and disabling (options cifs.oplocks.enable) 145
setting for volumes 115
options command, syntax 68
ownership change privileges, file 100

P

panic messages 428
parity disks 109, 160
description of 202
size of disks 223
passwords
changing administrative user account (passwd) 29, 30
changing system (passwd) 26
management 25
permissions, of default directories (/etc and /home) 59
plex
definition of 110
name of 110
synchronization 184
predicted failure event message 434
privilege levels for Data ONTAP commands 45

privileges, file ownership change 100
process, checksum upgrade 399

Q

qtree commands
    qtree create (creates a qtree) 138
    qtree security (changes security style) 143
qtrees
    administrative actions 13
    changing security style (qtree security) 143
    creation of (qtree create) 138
    definition of 107
    deleting 151
    description 13
    effects of deleting or renaming on quotas 370
    grouping files 136
    limitations of 137
    reasons for using in backups 136
    using df to view accurate free disk space 232
    using df to view free space (misleading data) 231
    when creating is appropriate 108, 136
qtrees and volumes
    changing security style in 143
    comparison of 108
    data access in 106
    definitions of 107
    security styles available for 139
quota reports
    contents 374
    ID and Quota Specifier fields 376
    soft limits in 380
    thresholds in 379
    types 373
QUOTA_PERFORM_USER_MAPPING in /etc/quotas 357
QUOTA_TARGET_DOMAIN in /etc/quotas 356
quotas
    /etc/quotas file. *See* /etc/quotas file in the "Symbols" section of this index
    applying to multiple IDs 343
    changing 365
    creating quota for target quotas (quotas on) 365
    default and derived, description of 338
    default user and group from default tree quotas 340
    default, where applied 338
    deleting 369
    deleting restrictions, description of 369
    derived, description of 335
    description of 334
    disabling (quota off) 362
    displaying report for (quota report) 381
    effects of deleting or renaming a qtree 370
    enabling (quota on) 361
    enabling and disabling, description of 361
    etc/quotas file. *See* /etc/quotas file in the "Symbols" section of this index
    exceeding, effects of 382
    explicit, description of 335
    getting status of initialization 365
    hard versus soft 335
    initialization, description of 337
    message logging
        display settings (quota logmsg) 364
        turning on or off (quota logmsg) 363
    overriding default 338
    prerequisite for working 337
    quotas file, format of 346
    reinitializing (quota off-quota on) 365
    resizing (quota resize) 367
    specifying default, advantages of 341
    targets, description of 334
    thresholds, description of 352
    types of reports available, description of 373
    types, description of 334
    UNIX IDs in 342
    user and group, rules for 347
    Windows group IDs in 343
    Windows IDs in 342
quotas commands
    quota logmsg (displays message logging settings) 364
    quota logmsg (turns quota message logging on or off) 363
    quota off (disables quotas) 362
    quota on (enables quotas) 361
    quota report (displays report for quotas) 381
    quota resize (resizes quota) 367

R

RAID
- automatic group creation 161
- changing from RAID 4 to RAID DP 172
- changing from RAID DP to RAID 4 174
- changing RAID type 172
- changing the group size option 177
- data reconstruction speed, modifying (options raid.reconstruct.perf_impact) 181
- data reconstruction speed, modifying (options raid.reconstruct_speed) 182
- data reconstruction, description of 181
- description of 159
- determining CPU load (sysstat) 179
- displaying statistics (vol status) 391
- displaying statistics of 386
- group size
  - changing (vol volume) 172, 177
  - comparison of larger versus smaller groups 164
  - default size 165
  - maximum size 176
  - specifying at creation (vol create) 165
- group size changes associated with RAID 4 to RAID DP changes 173
- group size changes associated with RAID DP to RAID 4 changes 174
- handling disk failures 168
- handling media errors during reconstruction 194
- maximum and default group sizes for RAID 4 176
- maximum and default group sizes for RAID DP 176
- maximum group sizes 176
- mirror verification speed, modifying (options raid.verify.perf_impact) 185
- modifying group size 176
- parity checksums 202
- plex resynchronization speed, modifying (options raid.resync.perf_impact) 184
- RAID groups described 109, 160
- reconstruction of disk failure 169
- throttling data reconstruction 181
- types 109
- verifying RAID type 175
- verifying the group size option 178

RAID 4
- description 109
- maximum and default group sizes 176

RAID commands
- vol create (specifies RAID group size) 165
- vol volume (changes RAID group size) 172, 177

RAID DP
- description 109
- maximum and default group sizes 176

RAID groups
- adding disks a RAID group other than the last RAID group 226
- definition of 12
- name of 110

RAID operations
- effects on performance 179
- types you can control 179

RAID type
- changing 172
- verifying 175

reasons to remove a disk 227
reconstruction after disk failure, data 171
recovered errors message, disk 437
Remote Shell access to filer (rsh) 36
remove a disk, reasons to 227
requirements
- for large files on a filer 89
- for Web browser 39
- to manage filer on NFS clients 23

root volume
- changing 130
- definition of 107

rsh (Remote Shell) access to filer 36
rsh commands
- accessing filer from a PC client 37
- accessing filer from a UNIX client 37
- format used with user name and password 36
- list of 37
- privilege levels 45
- use with user names and passwords 36

S

search order, changing (nsswitch.conf file) 268
security
  controlling file ownership changes (options wafl.root_only_chown) 100
  controlling mount privileges (options nfs.mount_rootonly) 99
  limiting Remote Shell access 98
  limiting Telnet access 97
  maintainence options 96
security styles
  changing for volume or qtree (qtree security) 143
  changing of for volumes or qtrees 143
  for volumes and qtrees 139
  mixed 139
  NTFS 139
  setting for volumes 115
  types available for qtrees and volumes 139
  UNIX 139
serial ports, filer main unit 9
serious error messages 428
session to filer, console or Telnet 33
shadow file, purpose 257
shm error messages 434
slots, filer main unit 9
SNMP messages
  DISK_DEGRADED-I/O 437
  DISK_FAILURE_PREDICTED!!! 434
  DISK_RECOVERED_ERRORS 437
  FC-AL LINK_FAILURE 434, 436
  FC-AL UNRELIABLE 436
SNMP protocol, in clustered takeover mode 319
startup commands in /etc/rc, operational actions 64
statistics (systat), CPU load 179
stats
  buffer cache 401
  checking expansion cards 386
  checksum upgrade process 399
  CPU utilization 401
  dafs stats command, description of 423
  dafs status command, description of 423
  displaying adapter information 389
  displaying all statistics 402
  displaying backplane environment status 405
  displaying chassis environment status 405
  displaying CPU environment status 405
  displaying Data ONTAP version 386
  displaying disk information 386, 389
  displaying disk information (vol status) 391
  displaying extended utilization 402
  displaying Fibre Channel driver statistics 407
  displaying link statistics 407
  displaying media changer information 389
  displaying overall filer information 387
  displaying power supply status 405
  displaying RAID and checksum information 391
  displaying RAID information 386
  displaying relative environment information 405
  displaying relative physical drive position 407
  displaying shelf environment status 405
  displaying statistics summary 402
  displaying tape drive information 387
  displaying tape information 390
  displaying tape library information 387
  displaying temperature status 405
  displaying volume information 387
  displaying volume language (vol status) 391
  displaying volume state statistics 391
  enabling client statistics (options nsf.per_client_stats.enable) 409
  environment, description of 404
  Fibre Channel statistics, description of 407
  filer statistics, description of 401
  network traffic 401
  nfsstat command, description of 408
  protocol operations 401
  read/write operations 401
  storage command, description of 389
  sysconfig command, description of 386
  sysstat and uptime commands 401
  sysstat command, description of 401
  terminating sysstat command 402
  uptime command, description of 403
  vol status command, description of 391
  when to enable per-client statistics 409
status commands
  environment chassis (displays shelf

environment information) 405
environment chassis backplane (displays backplane environment information) 405
environment chassis cpu (displays CPU environment information) 405
environment chassis power (displays power supply information) 405
environment chassis temp (displays temperature information) 405
environment command, description of 404
environment status (displays all filer environment information) 405
environment status shelf (displays shelf environment information) 405
fcstat device_map (displays relative physical drive position) 407
fcstat fcal_stats (displays Fibre Channel driver statistics) 407
fcstat link_stats (displays link statistics) 407
filestats command, description of 392
options nfs.per_client_stats.enable (enables clients statistics) 409
storage command, description of 389
storage hub (displays hub information) 389
storage show adapter (displays adapter information) 389
storage show disk (displays disk information) 389
storage show mc (displays media changer information) 389
storage show port (displays switch port information) 389
storage show switch (displays switch information) 389
storage show tape (displays tape information) 390
storage stats tape (displays tape statistics) 390
sysconfig -c (checks expansion cards) 386
sysconfig -d (displays disk information) 386
sysconfig -m (displays tape drive information) 387
sysconfig -m (displays tape library statistics) 387
sysconfig -r (displays RAID information) 386
sysconfig -v (displays overall filer statistics) 387
sysconfig -V (displays volume statistics) 387
sysstat -s (displays statistics summary) 402
sysstat -u (displaying extended utilization) 402
sysstat -x (displays all statistics) 402
ups (displays UPS environment information) 405
vol status (displays volume state) 391
vol status -c (displays checksum statistics) 399
vol status -d (displays disk statistics) 391
vol status -l (displays volume language) 391
vol status -r (displays RAID statistics) 391
storage commands
changes state of host adapter 209
disable 209, 210
enable 209, 210
reset tape drive statistics 220
show 211
syntax 208
view information about disks 216
view information about host adapters 212
view information about hubs 214
view information about media changers 218
view information about primary and secondary paths 217
view information about storage subsystem 211
view information about supported tape drives 219
view information about switch ports 218
view information about switches 218
view information about tape drives 219
view tape drive statistics 219
Storage Health Monitor 434
support for Autosupport, mail host 48
support, automatic technical 47
supported, types of disks 202
syslog messages 434
system board, filer main unit 8
system failures on a clustered system, description of 281
system memory, filer main unit 8
system password, changing 26

T

takeover
cluster mode, description of 285
configuring dedicated and hot standby interfaces 296
description of 278
description of mode 285
description of, in a cluster 278
interface configurations 292
interfaces, description of 291
mode with SNMP 319
statistics in 318
tape drives, displaying tape drive statistics (sysconfig) 387
tapes, displaying tape library statistics (sysconfig) 387
technical support, automatic 47
Telnet
access to filer 34
configuration of 95
filer access 33
limiting access to 97
session to filer console 33
termination of session with the filer 35
time
setting filer date and 83
synchronizing filer 78
troubleshooting
cluster problems 439
console messages 445
disk and FC-AL problems 434
filer panic messages 428
LCD messages on clustered takeover filer 440
messages on failed clustered filer 441
reasons to boot from diskette 448
syslog messages 442, 444
types of disks supported 202
types, checksum 202

U

Uninterruptible Power Supply (UPS)
adding a device to be monitored 406
enabling or disabling monitoring of 406
UNIX security style, description of 139
unmirrored volumes, definition of 110
UNRELIABLE Autosupport and SNMP message, FC-AL 436
unreliable message, FC-AL 435
upgrade process, checksum 399
use of administration hosts 22
use of FilerView interface 40
user account (useradmin userdel), deletion of administrative 30
user accounts (useradmin userlist), list of administrative 27
user information
/etc/group. *See* /etc/group file in the "Symbols" section of this index
/etc/netgroup 256
/etc/passwd. *See* /etc/passwd file in the "Symbols" section of this index
/etc/shadow. *See* /etc/shadow file in the "Symbols" section of this index
reason and location for maintaining 254
search order
description of 254
determination and changing 268

V

vol options command 69
volume commands
maxfiles (displays number of files in a volume) 153
maxfiles (increases number of files) 152, 155, 156, 157
vol create (creates a volume) 113
vol create (specifies RAID group size) 165
vol destroy (destroys an off-line volume) 131
vol lang (changes volume language) 121
vol offline (takes a volume off-line) 123
vol online (brings volume on-line) 124
vol options (changes RAID group type) 172
vol rename (renames a volume) 128
vol restrict (puts volume in restricted state) 124
vol status (displays volume language) 120
vol volume (changes RAID group size) 177
volume, root
changing 130

definition of 107
volumes
adding disks to 224
administrative actions 13
block statistics, displaying (vol status) 391
bringing on-line 124
changing the RAID type of 172
conventions of 112
creating (vol create) 107, 112, 113
default volumes and directories 59
definition of 107
destroying (vol destroy) 131
destroying, reasons for 131
disk statistics, displaying (vol status) 391
increasing number of files (maxfiles) 152, 155, 156, 157
language
changing (vol lang) 121
choosing of 117
displaying of (vol status) 120
setting of 119
when to change 120
languages supported 117
maximum number of files 152
mirrored, definition of 110
moving from one filer to another 125
number of files, displaying (maxfiles) 153
performance improvements for disk-bound volumes 425
plex, definition of 110
plex, name of 110
post-creation changes 115
RAID groups, definition of 12
RAID groups, name of 110
renaming a volume (vol rename) 128
renaming, reasons for 128
rules for adding disks to volume 222
security style 115
specifying RAID group size (vol create) 165
taking off-line (vol offline) 123
unmirrored, definition of 110
vol status command, description of 391
volume language, displaying (vol status) 391
volume state, definition of 122
volume state, determining 123, 147
volume state, displaying (vol status) 391
volume statistics, displaying (sysconfig) 387
when creating is appropriate 107
when to put in restricted state 124
volumes and qtrees
changing security style 143
comparison of 108
data access 106
definitions of 107
security styles available 139

W

Windows NT commands, performing Windows filer tasks 6

Y

ypgroup command 267

Z

zoned checksums 112, 202