Data ONTAP™ 7.0.1
System Administration 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 210-01110-A0
April 2005

Copyright and Trademark Information

Copyright information

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Portions Copyright © 2001, Sitraka Inc.

Portions Copyright © 2001, iAnywhere Solutions

Portions Copyright © 2001, i-net software GmbH

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

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

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

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

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

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

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

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

THIS SOFTWARE AND DOCUMENTATION IS PROVIDED "AS IS," AND COPYRIGHT HOLDERS MAKE NO REPRESENTATIONS OR WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO, WARRANTIES OF MERCHANTABILITY OR FITNESS FOR ANY PARTICULAR PURPOSE OR THAT THE USE OF THE SOFTWARE OR DOCUMENTATION WILL NOT INFRINGE ANY THIRD PARTY PATENTS, COPYRIGHTS, TRADEMARKS OR OTHER RIGHTS.

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

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

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

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

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

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

Trademark information

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

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

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

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

Network Appliance NetCache is certified RealSystem compatible.

Table of Contents

Preface ix

Chapter 1 **Introduction to Administering NetApp Systems** 1

About NetApp storage systems 2

About system components 4

About Data ONTAP 7

Chapter 2 **Interfacing with Data ONTAP** 15

How you administer a NetApp system 16

Using the command line interface 18

Using Data ONTAP commands at different privilege levels 21

Chapter 3 **Accessing the NetApp System** 23

Accessing the system 24

Accessing a system from the console 26

Accessing a system using a Telnet session 28

Accessing a system using a Remote Shell connection 32

Accessing a system using the HTTP/Web interface, FilerView 35

Managing access from administration hosts 40

Controlling system access 43

Chapter 4 **Understanding the Root Volume** 49

Default directories in the root volume 50

Contents of the /etc directory 51

Accessing the default directories on the NetApp system 54

Editing configuration files 56

Changing the root volume 58

Chapter 5 **Starting and Stopping the System** 61

Booting the system 62
Booting from the system boot diskette set 64
Booting the F87 system 67
Booting the FAS200 series systems 69
Booting the F825, F880, and FAS900 series systems 74

Using NetApp systems as netboot servers 77

Rebooting the system 80

Halting the system 81

Chapter 6 **Managing Administrator Access** 83

Using administrator accounts to control administrative access 84
About managing administrator access 85
Managing users 88
Managing groups 93
Managing roles 96
Listing and deleting users, groups and roles 102
Administrative user creation examples 108

Managing passwords 110

Chapter 7 **Performing General System Maintenance** 115

Aggregate snapshot management 116

Managing licenses 119

Setting the system date and time 121

Synchronizing the system time 123

Understanding core files 126

Configuring message logging 128

Configuring audit logging 132

Configuring NetApp system startup 134
About the /etc/rc file 135
Editing the NetApp system's boot configuration file 137
Recovering from /etc/rc errors 138

NetApp system configuration backup and cloning 139

UPS management 142

Chapter 8 **Using AutoSupport** 145

Learning about AutoSupport 146

Configuring AutoSupport. 148

Troubleshooting AutoSupport 154

Interpreting AutoSupport messages 157
About AutoSupport events 158
Contents of AutoSupport messages 159

Chapter 9 **Using SecureAdmin** 163

Learning about SecureAdmin. 164
How SecureAdmin uses SSH 165
How SecureAdmin uses SSL. 168

Managing SSH for SecureAdmin. 169
Setting up and starting SSH 170
Reinitializing SSH 172
Stopping or starting SSH service. 174
Setting up public key-based authentication 175

Managing SSL for SecureAdmin 179

General SecureAdmin administration 184

Chapter 10 **Cluster Management** 185

About NetApp clusters 186

How a clustered system works 190

Cluster operation modes 192
Normal operation mode 193
Takeover operation mode 195
Giveback operation mode 198

Understanding interface configurations for takeover. 200

Enabling cluster failover in a DAFS configuration. 202
Designating failover addresses in /etc/hosts 204
Enabling failover paths on the filer 207

Managing a cluster 209
Managing a cluster in normal mode 210
Enabling negotiated failover between filers 218
Managing a cluster in takeover mode 220
Managing an emulated filer 222

Backing up and restoring from the emulated filer 227
Performing a giveback 228
Enabling automatic giveback. 231
Configuring automatic giveback 232

Checking cluster configuration 234

Cluster LCD messages 235

Chapter 11 **System Information and Performance** 239

NetApp system configuration information 240

Storage information. 243

Aggregate information 245

Volume information 247

File statistics for volumes. 249
About the filestats command 250
Options for the filestats command 253

Environmental status information 256

Fibre Channel information 259

Getting NetApp system information using the stats command. 260

Improving NetApp system performance 270

Chapter 12 **Troubleshooting Tools** 275

NetApp system panics 276

Error messages 277

Other troubleshooting resources 279

Glossary 281

Index 291

Preface

Introduction

This guide describes how to configure, operate, and manage Network Appliance™ storage systems that run Data ONTAP™ 7.0 software. It covers F700 and F800 series filers, FAS appliances and NearStore® systems.

Audience

This guide is for system administrators who are familiar with operating systems, such as the UNIX®, Windows NT®, Windows 2000®, Windows Server 2003®, or Windows XP® operating systems, that run on the appliance's clients. It also assumes that you are familiar with how to configure the appliance 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.

Terminology

NetApp®'s storage products (filers, FAS appliances, and NearStore® systems) are all storage systems—also sometimes called filers or 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 Data ONTAP commands on the system console or from any client computer that can access the NetApp system through a session using Telnet, rsh, HTTP, SSH, or PuTTY.

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 NetApp 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 Administering NetApp Systems 1

About this chapter

This chapter provides a brief overview of what a NetApp® storage system is and what features it provides. It includes concepts that are described in detail in other guides in the Data ONTAP 7.0.1 library, such as an aggregate, a FlexVol™ volume (a flexible volume), a FlexClone™ volume (a flexible volume clone), and a traditional volume. For detailed information about these topics, see the *Storage Management Guide*.

Topics in this chapter

This chapter discusses the following topics:

- "About NetApp storage systems" on page 2
- "About system components" on page 4
- "About Data ONTAP" on page 7

About NetApp storage systems

What a NetApp storage system is

A NetApp storage system is a hardware- and software-based data storage and retrieval system. It responds to network requests from clients and fulfills them by writing data to or retrieving data from its disk array. NetApp storage systems provide a modular hardware architecture running the Data ONTAP™ operating system and Write Anywhere File Layout (WAFL®) software. For information about all of the models of NetApp storage systems, see http://www.netapp.com/products/.

Note

Backend storage subsystems such as IBM®, Hitachi Data Systems®, and HP® for gFiler™ gateways provide storage for data. NetApp gFiler gateways fulfill client requests from Logical Unit Numbers (LUNs) on the backend storage subsystems. For more information about gFiler gateways, see the *gFiler Planning Guide*.

Data ONTAP is the operating system for all NetApp systems. It provides a complete set of storage management tools through its command-line interface, through the FilerView® interface, through the DataFabric® Manager interface (which requires a license). For more information about Data ONTAP, see “About Data ONTAP” on page 7.

Components of a NetApp system

A NetApp system consists of the following components.

Component	Function
System main unit, or chassis, also known as the storage engine	The hardware device that receives and sends data. This unit also houses the system components and detects and gathers information about the hardware and the hardware configuration, the system components, operational status, hardware failures, and error conditions. For information about how to view this information, see “Environmental status information” on page 256. For information about environmental error codes, see the *Diagnostics Guide* on the NetApp on the Web (NOW) site at http://now.netapp.com.

Component	Function
Disk shelves	Containers, or *device carriers*, that hold disks and associated hardware (such as power supplies, connectivity, and cabling) that are connected to the system main unit of NetApp storage appliances. For more information, see “Disk shelves and disks” on page 6. **Note** For gFiler gateways, see the documentation for your storage subsystem for information about disks.

About system components

Types of components

The system has the following components:

- Internal components that enable the system to function
- Slots and ports that connect the system to networks
- Disk shelves that contain the disks

Internal components

The following internal components enable the system to function.

Component	Description
System board	The system board is also referred to as the main board of the system. 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.
CompactFlash™ card (not available on all models)	The system 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 system in an emergency.

Component	Description
Diskette drive (not available on all systems)	In an emergency or for major upgrades, Data ONTAP uses the diskette drive to boot from diskettes. **Note** Not applicable to gFiler gateways.
LCD and LEDs	The system displays status information on the LCD and LEDs.
Environmental adapter	The environmental adapter performs the following functions: ◆ Monitors the system's temperature and fans ◆ Sends critical information to the system's LCD ◆ Logs information ◆ Shuts down the system if its temperature is beyond a critical range or the fans cease operating
Remote Management Controller (RMC) (not available with all systems)	The RMC provides enhanced AutoSupport, such as "down filer" notification. For information on how to configure RMC, see the *Software Setup Guide*.

Slots and ports

The system has slots for external connections and ports for a console and diagnostic hardware. For information on how to configure host adapters for your NetApp system, see the *System Configuration Guide*.

Component	Description
Slots	The system contains expansion slots for the following host adapters: ◆ Network interface cards (NICs) ◆ Disk shelf adapters ◆ Tape drive adapters
Serial ports	The two serial ports are as follows: ◆ The console port connects the system 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.

Disk shelves and disks

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. For detailed information about disk shelves see the appropriate hardware service guide for your specific disk shelf. For detailed information about managing disks, see the *Storage Management Guide*.

Note

For information about disk shelves connected to gFiler gateways, see the appropriate gFiler integration guide, the *gFiler Planning Guide*, and disk shelf guide.

About Data ONTAP

Data ONTAP features

Data ONTAP provides the following features:

- Network file service
- Multiprotocol file and block sharing
- Data storage management
- Data organization management
- Data access management
- Data migration management
- Data protection
- System management
- AutoSupport

Network file service: Data ONTAP enables users on client workstations (or hosts) to create, delete, modify, and access files or blocks stored on the NetApp system.

NetApp systems can be deployed in network attached storage (NAS) and storage area network (SAN) environments for accessing a full range of enterprise data for users on a variety of platforms. NetApp systems can be fabric-attached, network-attached, or direct-attached to support NFS, CIFS, HTTP, FTP, and DAFS for file access, and iSCSI for block-storage access, all over TCP/IP, as well as SCSI over FCP for block-storage access, depending on your specific data storage and data management needs.

Client workstations are connected to the system through direct-attached or TCP/IP network-attached connections, or through FCP, fabric-attached connections. For information about configuring a NetApp system in a NAS network, see the *System Configuration Guide* and the *Network Management Guide*. For information about configuring a NetApp system in a SAN fabric, see the *Compatibility and Configuration Guide for NetApp's FCP and iSCSI Products*.

Multiprotocol file and block sharing: Clients can use the following protocols to access data on the NetApp appliance:

- NFS (Network File System)—used by UNIX systems
- (PC)NFS (Personal Computer NFS)—used by PCs to access NFS
- CIFS (Common Internet File System)—used by Windows clients

- FTP (File Transfer Protocol)—used for file access and retrieval
- HTTP (HyperText Transmission Protocol)—used by the World Wide Web and corporate intranets
- WebDAV (Web-based Distributed Authoring and Versioning)— used by HTTP clients for distributed web content authoring operations
- DAFS (Direct Access File System)—used by DAFS clients
- FCP (Fibre Channel Protocol)—used for block access in storage area networks
- iSCSI (Internet Small Computer System Interface)—used for block access in storage area networks

Files written using one protocol are accessible to clients of any protocol, provided that system 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 using one protocol can also be accessed by clients using the other protocol.

For information about NAS file access protocols, see the *File Access and Protocols Management Guide*.

For information about SAN block access protocols, see the *Block Access Management Guide for FCP* or the *Block Access Management Guide for iSCSI*.

Data storage management: Data ONTAP stores data on disks in disk shelves connected to NetApp appliances. Disks are organized into redundant array of independent disks (RAID) groups. RAID groups are organized into plexes, and plexes are organized into aggregates. These topics are explained in detail in the *Storage Management Guide*.

Note

For information about managing disks connected to gFiler gateways, see the *gFiler Planning Guide*.

Data organization management: Data ONTAP organizes the data in user- and system-files, and directories, in file systems called volumes, optionally in qtrees, and optionally, in LUNs (Logical Unit Numbers) in SAN environments. Aggregates provide the physical storage to contain volumes. These topics are explained in detail in the *Storage Management Guide*. LUNs are described in detail in the *Block Access Management Guide for FCP* and the *Block Access Management Guide for iSCSI*.

When Data ONTAP is installed on an appliance at the factory, a root volume is configured as /vol/vol0, which contains system files in the /etc directory. For more information about the root volume, see “Understanding the Root Volume” on page 49.

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.

  These topics are explained in detail in the *File Access and Protocols Management Guide*.
- 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.

  Snapshots are explained in detail in the *Data Protection Online Backup and Recovery Guide*.

Data migration management: Data ONTAP manages data migration by means of the following features:

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

Data protection: NetApp systems provide a wide range of data protection features, as described in the following table.

Feature	Description
Snapshot™ software	Manual or automatically scheduled multiple backups (or snapshots) of data using a minimal amount of additional disk space at no performance cost. For information about how Data ONTAP organizes and manages data, see the chapters on aggregate management and volume management in the *Storage Management Guide*. For information about snapshots, see the *Data Protection Online Backup and Recovery Guide*.

Feature	Description
SnapRestore® software (license required)	The SnapRestore feature performs fast restoration of backed-up data on request from snapshots on an entire volume. For information about SnapRestore, see the *Data Protection Online Backup and Recovery Guide*.
SnapMirror® software (license required)	System-to-system snapshot mirroring enables you to mirror snapshots on one system to a partner system. Should the original system be disabled, this ensures quick restoration of data from the point of the last snapshot. For information about SnapMirror, see the *Data Protection Online Backup and Recovery Guide*.
SnapVault™ software (license required)	SnapVault combines snapshot schedules and Qtree SnapMirror (QSM) to provide disk-based data protection for NetApp systems. You can also install the Open Systems SnapVault agent on non-NetApp systems. This allows SnapVault to back up and restore data to those systems also. Using SnapVault, you can periodically replicate selected snapshots from multiple client NetApp systems to a common snapshot on the SnapVault server. The snapshots on the server become the backups. You decide when to dump data from the SnapVault server to tape. As a result, you avoid the bandwidth limitations of tape drives, you restore data faster, and you don't need to perform full dumps from primary storage, so you don't need to schedule a backup window. For information about SnapVault, see the *Data Protection Online Backup and Recovery Guide*.

Feature	Description
Tape backup and restore	Tape backup `dump` and `restore` commands enable you to back up system or SnapVault snapshots to tape. Because the snapshot rather than the active file system is backed up to tape, the system can continue its normal functions while the tape backup is occurring. For information about tape backup, see the *Data Protection Tape Backup and Recovery Guide*.
`aggr copy`	This is fast block copy of data stored in aggregates; it enables you to quickly copy blocks of stored system data from one aggregate to another. For information about aggregates and `aggr copy`, see the *Storage Management Guide.*
`vol copy`	This is fast block copy of data stored in volumes; it enables you to quickly copy blocks of stored system data from one volume to another. For information about volumes and `vol copy`, see the *Data Protection Online Backup and Recovery Guide*.
SyncMirror® (cluster configuration required)	The SyncMirror software performs real-time RAID-level (RAID 4 or RAID-DP™) mirroring of data to two separate plexes that are physically connected to the same system head. If there is an unrecoverable disk error on one plex, the system automatically switches access to the mirrored plex. For information about supported RAID levels and plexes, see the *Storage Management Guide*. For information about SyncMirror, see the *Data Protection Online Backup and Recovery Guide*.
MetroCluster	MetroCluster enhances SyncMirror functionality for disaster recovery by providing continuous volume mirroring over 500-meter to 30-kilometer distances. For information about disaster protection using MetroCluster, see the *Data Protection Online Backup and Recovery Guide*

Feature	Description
NDMP (Network Data Management Protocol)	NDMP support enables third-party applications that use NDMP to manage tape backup operations of system data. The `ndmpcopy` command carries out NDMP-compliant backups and restores. Security login restricts access to NDMP operations. For information about NDMP, see the *Data Protection Tape Backup and Recovery Guide*.
NVFAIL	The `nvfail` option provides protection against data corruption by nonvolatile RAM (NVRAM) failures. For information about NVFAIL, see the *Data Protection Online Backup and Recovery Guide.*
virus scan support	Data ONTAP provides support for third-party-scanning software for files accessed by CIFS clients. For information about virus protection for CIFS, see the *Data Protection Online Backup and Recovery Guide.*

System management: Data ONTAP provides a full suite of system management commands that allows you to monitor system activities and performance. You can use Data ONTAP to

- Manage network connections
- Manage adapters
- Manage protocols
- Configure pairs of NetApp systems into clusters for failover
- Manage storage
- Dump data to tape and restore it to the NetApp system
- Mirror volumes (synchronously and asynchronously)
- Create vFiler™ units. For information about vFiler units, see the *MultiStore Management Guide*

For information about all Data ONTAP commands, see the *Commands: Manual Page Reference, Volume 1* and *Commands: Manual Page Reference, Volume 2*.

AutoSupport: Data ONTAP provides the AutoSupport feature, which automatically sends AutoSupport Mail notifications about system problems to NetApp technical support and up to five designated recipients. For more information about AutoSupport, see “Using AutoSupport” on page 145.

Interfacing with Data ONTAP 2

About this chapter

This chapter describes all of the ways you can execute Data ONTAP commands, including the Manage ONTAP™ Developer SDK software.

Topics in this chapter

This chapter discusses the following topics:

- "How you administer a NetApp system" on page 16
- "Using the command line interface" on page 18
- "Using Data ONTAP commands at different privilege levels" on page 21

How you administer a NetApp system

Ways to administer a system

You can use the following methods of administering a system with Data ONTAP:

- Command execution through the system's command line interface (CLI)
- Command execution through Windows network management tools
- Configuration file editing
- Command execution through the FilerView interface
- Command execution through DataFabric Manager software
- Programmatic execution through Manage ONTAP Developer SDK software

Command execution through the system's CLI

You can use the system's CLI to execute all Data ONTAP administrative commands, with the exception of some Windows server administrative commands.

You can access the system's command line from

- A serial terminal connected to the console port of the system
- A Telnet session to the system
- A remote shell program, such as the UNIX `rsh` utility (provides access for a limited set of commands)
- A secure shell application program, such as SSH, OpenSSH for UNIX

Command execution through Windows

You can use Windows commands to perform system administrative tasks related to Windows network operations. You can also use a secure shell application program, such as PuTTY.

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

Configuration file editing

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

You can access configuration files by mounting the root directory of the NetApp system on a UNIX client or by mapping the administrative share (C$) to a drive on a Windows client, then editing the file from the client. For more information about the root volume, /vol/vol0, and the configuration file in its root directory, /etc, see “Default directories in the root volume” on page 50.

Command execution through FilerView

You use FilerView to perform most administrative tasks from a Web-based interface. You can use FilerView whether or not you purchased a license for the HTTP protocol.

DataFabric Manager software

DataFabric Manager (DFM) is a simple, centralized administration tool that enables comprehensive management of enterprise storage and content delivery infrastructure. This suite of tools, which runs on a management server, consolidates tasks that would otherwise require separate steps and allows for a set of optional modules that provides specific additional functionality.

You must purchase the DataFabric Manager license to use this product. For more information about DataFabric Manager, see the DataFabric Manager Information Library at
http://now.netapp.com/NOW/knowledge/docs/DFM_win/dfm_index.shtml

Manage ONTAP Developer SDK software

The *Manage ONTAP SDK* contains resources necessary to develop third-party applications which monitor and manage NetApp systems. The kit contains libraries, code samples and bindings in Java, C, and Perl for the new ONTAPI™ programming interface set. A NetApp system simulator which runs on Linux or Solaris, that simulates the NetApp system to a very low level, is also available as a separate distribution. For more information, read about the Network Appliance Advantage Developer Program at
http://www.netapp.com/solutions/dfms/advantage.html

Using the command line interface

Features for editing commands

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

- History
- Command-line editor
- Online command-line help
- Online Data ONTAP command-line manual (man) pages

Using the history feature

The history feature enables you to scroll through recently entered commands, as described in the following table.

If you want to...	With the cursor keys	With a key-combination
Scroll back through commands	Press the Up arrow key.	Press Ctrl-P.
Scroll forward through commands	Press the Down arrow key.	Press Ctrl-N.

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, as shown in the following table.

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.

If you want to...	Then...
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.

Online command-line help

When using the command line, you can get command-line syntax help from the command line by entering the name of the command followed by `help` or the question mark (`?`). The fonts or symbols used in syntax help are as follows:

`keyword` specifies the name of a command or an option that must be entered as shown.

< > (less than, greater than symbols) specify that you must replace the variable identified inside the symbols with a value.

| (pipe) indicates you must choose one of elements on either side of the pipe.

`[ ]` (brackets) indicate that the element inside the brackets is optional.

{ } (braces) indicate that the element inside the braces is required.

The following example show the result of entering the `environment help` command at the command line:

```
filer> environment help
Usage: environment status |
                [status] [shelf [<adapter> ] ] |
                [status] [chassis [all | Fans | Power |
Temperature | PS1 | PS2 | RTC battery ]]
```

You can also type the question mark at the command line for a list of all the commands that are available at the current level of administration (administrative or advanced). For information about administration levels, see “Using Data ONTAP commands at different privilege levels” on page 21.

Command line man pages

NetApp provides online Data ONTAP manual (man) pages for the following types of information. They are grouped into sections according to standard UNIX naming conventions.

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

Man pages can be viewed in the following ways:

- At the NetApp system command line, by entering

  `man` *`command_or_file_name`*
- From the FilerView main navigational page
- In the following documents:
  - *Commands: Manual Page Reference, Volume 1*
  - *Commands: Manual Page Reference, Volume 2*

Note

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

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

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 system. At the advanced level, in addition to all the 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 NetApp technical support.

Caution

Using advanced commands can be dangerous. Using some advanced commands without consulting NetApp technical support might result in data loss.

How different privilege settings apply to different sessions

Sessions opened through the console, Telnet, and secure shell applications 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 an administrator who is connected to the system 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 for 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 reset to 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.
	Note
	If no argument is given, the default, `admin`, is applied.
	Example: Assuming the name of the NetApp system is *toaster*, the system prompt is *toaster*>, as shown in the following example.
	`toaster> priv set advanced`
	Result: The following message is displayed, followed by the advanced mode system prompt.
	`Warning: These advanced commands are potentially dangerous; use them only when directed to do so by Network Appliance personnel.`
	`toaster*>`

Accessing the NetApp System 3

About this chapter

This chapter describes the methods you can use to access the NetApp system, what you need to know about those methods, and how you manage NetApp access from administration hosts.

Topics in this chapter

This chapter discusses the following topics:

- "Accessing the system" on page 24
- "Accessing a system from the console" on page 26
- "Accessing a system using a Telnet session" on page 28
- "Accessing a system using a Remote Shell connection" on page 32
- "Accessing a system using the HTTP/Web interface, FilerView" on page 35
- "Managing access from administration hosts" on page 40
- "Controlling system access" on page 43

Accessing the system

What type of access is needed

The ways to access the system depend on the tasks you intend to perform. You can

- Administer the system—To access the system to administer it, you only need network connectivity to the system and authentication privileges. No licenses are required to access the system to administer it.
- Store data on the system—To store and receive data on the system, you must have an NFS or a CIFS license installed to access the system, mount system directories, and store data and retrieve data.

Accessing the system to administer it (no licenses are required)

You can directly access a system to administer it using one of the following methods:

- From a console that is attached by a cable to the system's serial port
- From the Ethernet network interface card (NIC) that is preinstalled in the NetApp system. NetApp highly recommends you use this card to connect to a TCP/IP network to administer the system:
  - From any client by using a Telnet session
  - From any client by using a Remote Shell connection
  - From any client by using a Web browser and the FilerView interface
  - From any client by using a secure shell client application, such as SSH, OpenSSH for UNIX hosts, or PuTTY for Windows hosts

Accessing the system to store and retrieve data (licenses are required)

You can directly access a system to administer it, and to store and retrieve data, using one of the following methods:

- From a console that is attached by a cable to the system's serial port
- From the Ethernet network interface card (NIC) that is preinstalled in the NetApp system. NetApp highly recommends you use this card to connect to a TCP/IP network to administer the system:
  - 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

- From an NFS or CFS client by using a secure shell client application, such as SSH, OpenSSH for UNIX hosts or PuTTY for Windows hosts
- From an NFS client or CIFS client by using a Web browser and the DataFabric Manager (DFM) interface (a DFM license is also required)
- From a CIFS client to provide support for the SnapDrive™ feature in a Windows environment
- From an NFS client or CIFS client to manage Fibre Channel switches (in a SAN environment)

Sharing console and Telnet sessions

If you open a console session and a Telnet session at the same time, the default behavior of a NetApp system causes those sessions to share a single session. As a result, the following actions occur:

- Commands typed at either location are echoed to the other location.
- Messages are displayed at both locations.

Note

You can abort commands entered at the console or through a Telnet session by pressing Ctrl-C.
If the default is in effect when 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, or you can change the default behavior.

You can change this default behavior by entering the following command:

```
options telnet.distinct.enable on
```

This causes the console session to have a distinct user environment from the Telnet session. As a result,

- Commands that are typed at one location are *not* echoed to the other location.
- Messages are *not* displayed at both locations.
- User privileges are *not* shared between console and Telnet sessions.
- Auditlog entries identify console and Telnet access by including the phrase “console shell” or “telnet shell,” as shown in the following example:

  ```
  Fri Feb 18 12:51:13 GMT [toaster: rc:debug]: root:IN:console
  shell:df
  ```

Accessing a system from the console

Using the serial port

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

Step	Action
1	At the console, press Enter. **Result:** The system responds with the login or password prompt.
2	If the system displays the login prompt, do one of the following: ◆ To access the system with the system account, enter the following account name: `root` ◆ To access the system with an alternative administrative user account, enter the following: ***username*** *username* is the administrative user account. **Result:** The system 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.

4	When you see the NetApp system prompt followed by a system message, press Enter to get to the system prompt.
	Example:
	`toaster> Thu Aug 5 15:19:39 PDI [filer: telnet_0:info]: root logged in from host: unix_host12.xxx.yyy.com`
	Press Enter
	`toaster>`
	Note
	You can abort commands entered at the console by pressing Ctrl-C.

When options command values take effect

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

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

For more information about these options, see the online na_options page or the *Commands: Manual Reference Page, Volume 1*.

Accessing a system using a Telnet session

Requirements for using a Telnet session

The following requirements must be met before you can connect to a NetApp system using a Telnet session.

- The `telnet.enable` option must be set to On, which is the default setting. You verify the option is on by entering the `options telnet` command. You set the option to On by entering the `options telnet.enable on` command. For more information, see the Options man page.
- The `telnet.access` option must be set so that the protocol access control defined for the NetApp system allows Telnet access. For more information, see the Options and Protocolaccess man pages.

Number of active Telnet sessions allowed

Only one Telnet session can be active at a time. You can, however, open a console session at the same time a Telnet session is open. For more information, see "Sharing console and Telnet sessions" on page 25.

Using a Telnet session

To access a system from a client through a Telnet session, complete the following steps.

Step	Action
1	Open a Telnet session on a client.
2	Connect to the NetApp system using its name.
3	If the NetApp system displays the login prompt, do one of the following: ◆ To access the system with the system account, enter the following account name: **`root`** ◆ To access the system with an alternative administrative user account, enter the following: ***`username`*** *username* is the administrative user account. **Result:** The NetApp system responds with the password prompt.

Step	Action
4	Enter the password for the root or administrative user account. **Note** If no password is defined for the account, press Enter.

5	When you see the NetApp system prompt followed by a system message, press Return to get to the system prompt.
	Example:
	`toaster> Thu Aug 5 15:19:39 PDI [toaster: telnet_0:info]: root logged in from host: unix_host12.xxx.yyy.com`
	Press Enter
	`toaster>`
	Note
	You can abort commands entered through a Telnet session by pressing Ctrl-C.

Terminating a Telnet session

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

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

What you can configure

You can configure the following behavior for Telnet sessions:

- A banner message appearing above the Telnet login prompt

- A timeout period

Configuring a banner message

You can configure a banner message to appear at the beginning of a Telnet session to a NetApp system by creating a file called *issue* in the /etc directory of the administration host's root volume. The message only appears at the beginning of the session. It is not repeated if there are multiple failures when attempting to log in. The following example shows how the message in /etc/issue appears, assuming the contents of the issue file is "This system is for demonstrations only."

```
admin_host% telnet toaster
Trying 172.15.95.97.49...
Connected to toaster.xyz.com
Escape character is '^]'.

This system is for demonstrations only.

Data ONTAP <toaster.xyz.com>
Login:
```

Configuring the timeout period

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. The range of minutes is 1 to 35,791. The maximum number is equal to approximately 596 hours, or slightly less than 25 days.

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

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

For more information about these options, see the online na_options(1) man page or the *Commands: Manual Reference Page, Volume 1*.

Accessing a system using a Remote Shell connection

Remote Shell with trusted remote hosts supported

The system supports a Remote Shell with trusted remote hosts. Trusted remote hosts are 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 system. The following table describes the circumstances under which a user name and password are required.

If you are logged in ...	Then...
As *root* on a UNIX host that is listed in the system's /etc/hosts.equiv file	You do not need to supply a user name or a password.
As a *user* other than root on a UNIX host that is listed in the system'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 system.

Note

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

Maximum number of rsh sessions allowed

You can have up to 24 concurrent rsh sessions running on a system, and you can have up to 4 concurrent rsh sessions running on each vfiler.

Format for rsh commands with user name and password

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

`rsh -l` *`username`* `[:`*`password`*`]` *`host_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 in this manner is a security risk, especially for UNIX clients. On many UNIX clients, this command can be visible to other users on the system who run the `ps` program at the same time the command is executed.
On any client, the password is visible in plaintext over the network. Any program that is capturing network traffic when the password is sent will record the password. To avoid exposing the password, log in as root on a client listed in the NetApp system's /etc/hosts.equiv file and issue the command.

Commands not accepted from rsh

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

```
arp
orouted
ping
routed
savecore
setup
traceroute
```

Accessing a system from a UNIX client by using an rsh command

To access a system from a UNIX 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 system from a Windows client by using a Remote Shell application

To access a system from a Windows 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.
3	Enter the `rsh` command, replacing *command* and *arguments* with the desired command and arguments: **`rsh filername command arguments`** **Example:** `rsh filer12 sysstat 1`

Resetting options to default values from rsh

If you want to reset options to their default values from rsh, you must precede the quotation characters (") with the escape character, which is the back slash (\). For example, if you want to reset the CIFS home directory path from a Windows host using a console session, you would enter the following command:

```
c:\> toaster options cifs.home_dir ""
```

However, from an rsh session, you must enter the following command:

```
c:\> rsh toaster options cifs.home_dir \"\"
```

Accessing a system using the HTTP/Web interface, FilerView

What FilerView is

FilerView is an HTTP/Web-based graphical management interface that enables you to manage most NetApp system functions from a Web browser rather than by entering commands at the console, through a Telnet session, the `rsh` command, or by using scripts or configuration files.

You can also use FilerView to view information about the system, its physical storage units, such as adapters, disks and RAID groups, and its data storage units, such as aggregates, volumes, and LUNs. You can also view statistics about network traffic.

FilerView is easy to use, and you can access online Help, which explains Data ONTAP features and how to use them.

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 also be compatible with FilerView.

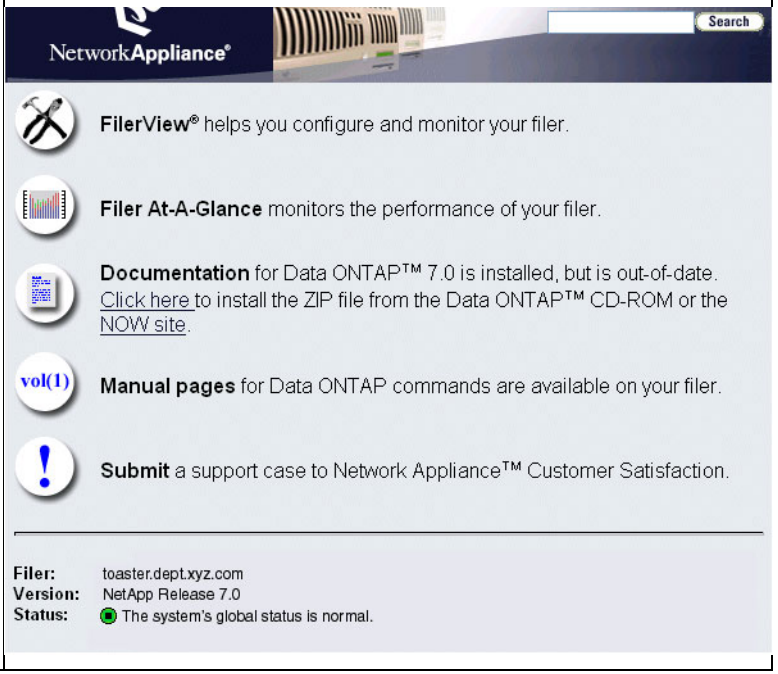
Client platforms that support FilerView

The following platforms support FilerView:

- Windows 95
- Windows 98
- Windows 2000
- Windows Server 2003
- Windows NT
- Windows XP
- UNIX
- Solaris
- Linux
- HPUX

Accessing a system from a client by using FilerView

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

Steps	Action
1	Start your Web browser.
2	Enter the following URL, replacing *filername* with the name of your system: **`http://filername/na_admin`** *filername* is either the fully qualified name or the short name of the system or the IP address of the system. **Result:** The Network Appliance Online administrative window appears. 

Steps	Action
3	Click FilerView. **Result:** ◆ If the system is password protected, you are prompted for a user name and password. ◆ Otherwise, FilerView is launched, and a screen appears with a list of categories in the left panel and the System Status information in the main panel. **Note** To access FilerView securely, set up the SecureAdmin feature and enable it for SSL by using the `secureadmin` command. Then access the Web site using https://filername/na_admin. For more information about SecureAdmin, see Chapter 9, "Using SecureAdmin," on page 163 and the na_secureadmin(1) man page.

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 topics:

- Most of the categories represent management functions.
- The Real Time Status category contains choices that launch separate tools that monitor system performance.
- The Wizards category contains choices that launch separate wizards for system setup, CIFS setup, and vFiler setup.

Selecting a function or wizard

To select a function or wizard, complete the following step.

Step	Action
1	Click the function name or wizard. **Result:** ◆ If you click a function name, the category expands and you can select a specific task. ◆ If you click a wizard, you are prompted to enter data or make selections.

Viewing the real-time displays

To view the real-time displays, complete the following steps.

Step	Action
1	Click Real Time Status.
2	Click the display you want to view. **Note** If you select Health Monitor, a Java applet is launched, which may take several seconds to load.

Right frame: If you select the add function from the left frame for volumes or aggregates, a wizard is launched. If you select manage, configure or report functions from the left frame, the right frame changes to display forms that provide information about the system configuration. You can change the system configuration by entering data in the fields or by making selections from lists.

To change the system configuration, complete the following steps.

Step	Action
1	Change information that is displayed in the right frame.
2	Save the information by pressing the Add button.

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 a category. **Result:** A two-frame Help screen appears. The left frame displays an expandable table of contents, with additional tabs at the top labeled Index and Search. As you select topics from the left frame, the right frame displays a page with tabs at the top labeled Concepts, Procedures, and More Information, which relate to the topic you selected. The tabs only appear when the topic in left frame specifically cover procedures, and More Information is not available for all topics.
2	To view tasks, click the Procedures tab when it appears at the top of the right frame.
3	To view additional information about a topic, when available, click the More Information tab when it appears at the top of the right frame.
4	To view descriptions of fields, click the Help icon next to the appropriate field in the left FilerView frame.

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 system remotely than to go to the system's physical location.

What an administration host is

The administration host is a remote workstation from which you can administer the system. The remote workstation can be an NFS client or a CIFS client. When you designate an administration host, the system'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 system. After booting the system, you can add as many clients as you want to be administration hosts.

Administration host privileges

The system 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 system root directory and edit configuration files from the administration host. ◆ Enter Data ONTAP commands by using a Remote Shell connection.
A CIFS client	Edit configuration files from any CIFS client as long as you connect to the system as root or "Administrator."

Controlling access for administrative hosts

You can make sure that only specific hosts are used as administrative hosts to increase security for your system.

Requirements for using an NFS client

If you plan to use an NFS client to manage the system, 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 system, 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*]

or

hostname ["user name"]

Example:

```
client1 peter
```

or

```
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 system 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 Data ONTAP 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.

Controlling system access

About this section

Data ONTAP enables you to control how administrators can access the system. By limiting how. and from where. administrations can log on, you can increase the security of your NetApp system.

This section contains the following topics:

- "Controlling Telnet access using host names" on page 43
- "Controlling Telnet access using host names, IP addresses, and network interface names" on page 44
- "Controlling Remote Shell access" on page 44
- "Controlling mount privilege" on page 45
- "Controlling file ownership change privileges" on page 46
- "Controlling anonymous CIFS share lookups" on page 46
- "Options that help maintain security" on page 47

Controlling Telnet access using host names

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

Step	Action	
1	Access the NetApp system 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 NetApp system 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 NetApp system 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 NetApp system from a single host, complete the following steps.

Step	Action
1	Access the NetApp system 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 NetApp system 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 NetApp system'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 NetApp system 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 user without root privileges changes the owner of a file, the set-user-id and set-group-id bits are cleared.
- If a user without root privileges 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 NetApp system 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 NetApp system, complete the following steps.

Step	Action
1	Access the NetApp system 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`

Options that help maintain security

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

Option	Description
`trusted.hosts`	Controls which hosts can access the NetApp system through a Telnet session for administrative purposes. You can disable Telnet access completely, allow Telnet access only from up to five hosts, or allow Telnet access from all hosts.
`telnet.access`	Controls which hosts can access the NetApp system through a Telnet session for administrative purposes. You can restrict Telnet access to the NetApp system by specifying host names, IP addresses, or network interface names.
`rsh.access`	Controls which hosts can access the NetApp system through a Remote Shell session for administrative purposes. You can restrict Remote Shell access to the NetApp system by specifying host names, IP addresses, or network interface names.

Option	Description
nfs.mount_rootonly	Controls whether the NetApp system'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.
cifs.restrict_anonymous.enable	Controls whether anonymous CIFS users can look up CIFS shares, users, or groups on a NetApp system. 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.

Understanding the Root Volume 4

About this chapter

This chapter introduces you to the NetApp system's root volume, which contains special directories and configuration files that help you administer your NetApp system.

About the root volume

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

For NetApp systems that have Data ONTAP 7.0 or later installed at the factory, the root volume is a FlexVol. Systems installed with earlier versions of Data ONTAP have a traditional root volume.

For more information about traditional and FlexVol volumes, see the *Storage Management Guide*.

Topics in this chapter

This chapter discusses the following topics:

- "Default directories in the root volume" on page 50
- "Contents of the /etc directory" on page 51
- "Accessing the default directories on the NetApp system" on page 54
- "Editing configuration files" on page 56
- "Changing the root volume" on page 58

Default directories in the root volume

What the default directories are

The root volume contains the following directories that were created when the NetApp system was set up:

- The /etc directory, which contains configuration files that the NetApp system 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 NetApp system 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 contents

The root directory contains an /etc directory, which stores NetApp system configuration files, executables required to boot the system, and some log files. You can modify the configuration files from the administration host.

You can also affect how the system boots by modifying the /etc/rc file. For more information about this file, see “Configuring NetApp system startup” on page 134.

Caution

Do not delete any directories from the /etc directory unless instructed to do so by NetApp personnel.

Descriptions of configuration files

The following table describes the configuration files in the /etc directory. Some of the configuration files might not exist on your NetApp system if you use the system 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 NetApp system contacts name services
quotas	Stored user and group quota specifications
rc	Script of commands to be executed when the NetApp system 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 NetApp system is a member of a domain)

File name	Contents
cifsconfig.cfg	CIFS commands that the NetApp system used for configuration **Caution** The NetApp system maintains this file automatically whenever you enter a `cifs` command, which can cause changes you make manually to be lost.
group	CIFS group names, GIDs (group identification numbers), and members' names Not used if you use Network Information Service (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

About /etc/messages

By default, all system messages of level INFO and higher are sent to the console and to the /etc/messages file. This way, you can see a record of events on your NetApp system, as well as use scripts to parse for particular events if needed.

The /etc/messages file is rotated once a week, and six weeks of messages are retained.

You can use the `logger` command to create and send a system message explicitly. For more information about the `logger` command, see the na_logger(1) man page.

If you would like to change the level of messages that are sent to /etc/messages, you can edit /etc/syslog.conf. For more information about message levels and the /etc/syslog.conf file, see the na_syslog.conf(5) man file.

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. Standard Generalized Markup Language (SGML) entities are allowed only in the root volume UNIX encoding.

Accessing the default directories on the NetApp system

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 NetApp system. **Result:** You have access to the NetApp system 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 NetApp system. **Result:** You have access to the NetApp system 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 NetApp system.

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 NetApp system. **Note** You can also browse the Network Neighborhood to locate the NetApp system and the HOME directory.

Accessing log files using HTTP

To access your log files using HTTP, complete the following step.

Step	Action
1	Point your browser to the following location: **`http://<system_name>/na_admin/logs/`** *system_name* is the name of your NetApp system.

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 NetApp system. You must use an editor from an NFS client or a CIFS client to edit NetApp system 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 NetApp system 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 NetApp system console, a Telnet client, or the `rsh` command to issue the following command to the NetApp system: **`exportfs`**
4	Mount the NetApp system 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 NetApp system as Administrator. The NetApp system root directory was shared automatically as C$. The administrator has read, write, and execute rights to the share.

Note

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 NetApp system as Administrator.
2	Display the contents of the NetApp system's C$ share, and select a file to edit.

Configuration files that require a carriage return after the last entry

Data ONTAP requires that the following configuration files be terminated with a carriage return:

- /etc/passwd
- /etc/group
- /etc/netgroup
- /etc/shadow

When you edit these files, be sure to insert a carriage return after the last entry.

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.*

Changing the root volume

About changing the root volume

You might want to change the NetApp system's root volume, perhaps to move your root volume from a traditional volume to a FlexVol. Every NetApp system must have a root volume; therefore, you must always have one volume designated as the root volume.

Changing the root volume of a NetApp system

You can change which volume on your NetApp system is used as the root volume.

Note

To change your root volume from a traditional volume to a FlexVol or from a FlexVol to a traditional volume, use the procedure outlined in the chapter about volumes in the *Storage Management Guide*.

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

Step	Action
1	Identify an existing volume to use as the new root volume, or create the new root volume using the `vol create` command. For more information about creating volumes, see the *Storage Administration Guide*. **Note** The required minimum size for root volumes can vary. NetApp recommends that your root volume be at least 90 GB.
2	Using the appropriate file access protocol for your installation, copy the /etc directory and all of its subdirectories from the current root volume to the new root volume. Or, you can use `ndmpcopy`. For more information about `ndmpcopy`, see the *Data Protection Tape Backup and Recovery Guide*.

Step	Action
3	Enter the following command: **`vol options vol_name root`** *vol_name* is the name of the new root volume. **Note** After a volume is designated to become the root volume, it cannot be brought offline or restricted.
4	Enter the following command to reboot the NetApp system: **`reboot`** **Result:** When the NetApp system finishes rebooting, the root volume is changed to the specified volume.

About the root option for aggregates

Besides the *volume* `root` option that you use to determine which volume will be the root volume after the next NetApp system reboot, there is also an *aggregate* `root` option. The aggregate `root` option is used only when, for some reason, the NetApp system cannot determine which volume to use as the root volume.

Note

NetApp recommends that you not change the value of the `root` option for any aggregate unless instructed to do so by NetApp technical support.

Starting and Stopping the System 5

About this chapter

This chapter describes how to boot a NetApp from the system boot diskette set, from the PC CompactFlash card, and from remote images on supported systems.

Topics in this chapter

This chapter discusses the following topics:

- "Booting the system" on page 62
- "Using NetApp systems as netboot servers" on page 77
- "Rebooting the system" on page 80
- "Halting the system" on page 81

Booting the system

Where the system boots from

NetApp systems boot from a PC CompactFlash Card, from disk, or from a set of floppy diskettes, depending on the system. If your system has a PC CompactFlash Card, it is shipped with the most current Data ONTAP release and a diagnostic kernel in a single File Allocation Table (FAT) partition. The CompactFlash card contains sufficient space for an upgrade kernel.

Note

Single partitions were introduced in Data ONTAP 6.5. Prior to that, CompactFlash cards had four partitions. If you are upgrading your system from a version of Data ONTAP earlier than 7.0, your system will repartition the CompactFlash card. This allows the existing 32-MB CompactFlash to accommodate the size of the new kernel. Support for single partitions also includes support for larger 256-MB cards. For information about how to upgrade the CompactFlash card in your system, see the *Upgrade Guide*.

Booting the system from the system or firmware prompt

You can boot the system if it is on and displaying one of the following prompts:

- The system prompt (`toaster>`); if so, see "Rebooting the system" on page 80."
- The firmware prompt (`ok`)

To boot the system, complete the following step.

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

Platform-specific booting information

The following sections describe how to boot specific NetApp systems. For information about how your system boots, see the hardware and service guide for your system.

- "Booting from the system boot diskette set" on page 64
- "Booting the F87 system" on page 67
- "Booting the FAS200 series systems" on page 69
- "Booting the F825, F880, and FAS900 series systems" on page 74

Booting from the system boot diskette set

Systems that support floppy diskettes

Only older NetApp systems have floppy diskette drives. You can visually inspect your system to determine whether it has a floppy diskette drive, or you can look at the Software Download page on NOW to see whether there are floppy diskette images to download.

Reasons to boot from diskette

You might need to boot the system 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 system with an earlier version of system boot diskettes. The attempt to boot will fail because earlier versions of Data ONTAP software cannot recognize system data from a system running Data ONTAP 7.0 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 system's diskette drive. **Note** Data ONTAP system boot diskettes are specific to the system model. You received diskettes appropriate to the model you are updating. If you have more than one model, be sure to check the diskette labels to make sure you're using the correct diskettes for the system.

<table>
<tr><th>Step</th><th colspan="2">Action</th></tr>
<tr><td rowspan="3">2</td><td>If....</td><td>Then...</td></tr>
<tr><td>Your system is turned on</td><td>From the system console, enter the following command:
<code>reboot</code>
Result: The system begins the boot process.</td></tr>
<tr><td>Your system is turned off</td><td>Turn on your system.
Result: The system begins the boot process.</td></tr>
<tr><td>3</td><td colspan="2">When the system’s LCD prompts you, remove the diskette and insert the next diskettes into the system’s diskette drive when you are prompted.</td></tr>
<tr><td>4</td><td colspan="2">Press Enter.
Result: The system boots and, if it can, displays the following boot menu:

<code>1) Normal Boot
2) Boot without /etc/rc
3) Change Password
4) Initialize all disks
4a)Same as option 4, but create a flexible root volume.
5) Maintenance mode boot

Selection (1-5)?</code></td></tr>
</table>

<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 system 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 system 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>Change the password of the system</td><td>Change password (3)</td></tr>
<tr><td>Initialize all the disks and create a traditional root volume</td><td>Initialize all disks (4)</td></tr>
<tr><td>Initialize all the disks and create a flexible root volume</td><td>Same as option 4, but create a flexible root volume. (4a)</td></tr>
<tr><td>Perform some aggregate and disk operations and get detailed aggregate 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 aggregate/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 system

Boot options for the F87 system

The F87 system boots Data ONTAP from a 32-MB PC CompactFlash card. It does not boot from a set of diskettes. When you turn on or reboot the F87, the system 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 system 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 system from one of its SCSI disks

Note

You boot the system from a SCSI disk only when you install a new NVRAM adapter in the system. Do not use this option for normal operation of the system. 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 system.
	Result : The system 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 system boots the current release of Data ONTAP.

Booting the FAS200 series systems

Boot options

The FAS200 series (FAS250, FAS270, and FAS270c) systems automatically boot Data ONTAP from a PC CompactFlash card. They do not boot from a set of 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 systems 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 system 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 system 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 system. It provides a “known good” release from which you can boot the system if it fails to automatically boot the primary image.

- `netboot`

  Boots from a Data ONTAP version stored on a remote HTTP or TFTP (Trivial File Transfer Protocol) 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 appliance (using DHCP or static IP address) and place the boot image on a configured server.

Note

To protect against data loss in the event of PC CompactFlash card corruption, NetApp recommends that you enable netboot immediately for your system by placing a boot image on a local server. For more information, see "Preparing a system for netboot startup" on page 71.

- `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

For FAS2000 series appliances, you can boot the current release of Data ONTAP or the backup release at the firmware prompt, by completing the following steps.

Step	Action
1	At the system console, enter the following command: **`halt`** **Result:** The system 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 system 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 started via netboot from an unfamiliar system, if an upgrade was unsuccessful, or if you need to run kernel diagnostics.

To determine...	At the system 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 a system for netboot startup

To use the netboot option to start your system, 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 system boot directory, /etc/boot/netapp-mips, or downloaded from the NetApp on the Web (NOW) site at http://now.netapp.com/.
- 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 appliances through the netboot option

To boot Data ONTAP from a remote image, complete the following steps.

Step	Action	
1	At the system 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 system. *netmask* is the network mask of the system. *gateway* is the gateway for the system. *dns_addr* is the IP address of a name server on your network. *dns_domain* is the Domain Name Service (DNS) domain name. 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.

Booting the F825, F880, and FAS900 series systems

Booting F800 and FAS900 systems with CompactFlash cards

F800 series and FAS900 series systems are configured to boot from 32-MB or 256-MB PC CompactFlash cards. If you have upgraded your F800 and FAS900 systems to Data ONTAP 6.5 or higher, either your 32-MB CompactFlash boot cards have been repartitioned as described in the *Data ONTAP Upgrade Guide*, or you have upgraded to 256-MB CompactFlash boot cards.

If you reboot your system, the system will reboot by default in normal mode. You can also invoke a boot menu that allows you to reboot in alternative modes for the following reasons:

- To correct configuration problems
- To recover from a lost password
- To correct certain disk configuration problems

To boot or reboot an upgraded F825 or F880 system or an FAS900 series system, complete the following steps.

Step	Action	
1	**If....**	**Then...**
	Your system is turned on	From the system console, enter the following command: `reboot` **Result:** The system begins the boot process.
	Your system is turned off	Turn on your system. **Result:** The system begins the boot process.

<table>
<tr><th>Step</th><th colspan="2">Action</th></tr>
<tr><td rowspan="3">2</td><td>If....</td><td>Then...</td></tr>
<tr><td>You want the system to boot automatically in normal mode</td><td>Allow the system to reboot uninterrupted until the following prompt appears:
<code>root logged in from console</code>
You are done.</td></tr>
<tr><td>You want to select from a menu of alternative boot modes</td><td>When prompted to do so, press Ctrl-C to display the boot menu.
Result: The system displays the following boot menu:
<pre>1) Normal Boot
2) Boot without /etc/rc
3) Change Password
4) Initialize all disks
4a) Same as option 4, but create a
flexible root volume.

5) Maintenance mode boot

Selection (1-5)?</pre></td></tr>
</table>

<table>
<tr><th>Step</th><th>Action</th></tr>
<tr><td>3</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>Boot the system normally.</td><td>Normal boot (1)</td></tr>
<tr><td>Troubleshoot and repair configuration problems.
Note
Booting without /etc/rc causes the system 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>Change the password of the system</td><td>Change password (3)</td></tr>
<tr><td>Initialize all the disks and create a traditional root volume</td><td>Initialize all disks (4)</td></tr>
<tr><td>Initialize all the disks and create a flexible root volume</td><td>Same as option 4, but create a flexible root volume. (4a)</td></tr>
<tr><td>Perform some aggregate and disk operations and get detailed aggregate 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 aggregate/volume problems.
◆ You exit Maintenance mode with the halt command. To reboot the appliance, enter <code>boot</code> after the firmware <code>ok</code> prompt.</td><td>Maintenance mode boot (5)</td></tr>
</table></td></tr>
</table>

Using NetApp systems as netboot servers

Netboot server requirements

You can configure a system to serve boot images to other NetApp devices that support netboot. To do so, you must configure

- HTTP services, TFTP services, or both on the system.
- The rest of your netboot using environment to access the system 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 system /etc/http directory: You can also serve netboot images from your system 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 system as an HTTP netboot server, complete the following steps.

Step	Action
1	Place boot programs in the system's /etc/http directory.
2	At the command line of the source system, enter the following command: `options httpd.enable on`

Configuring TFTP services

To configure a system as a TFTP netboot server, complete the following steps.

Step	Action
1	Place boot programs in the system /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`.
2	On the command line of the source system, 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 system, 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 system.

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 system, enter the following command: `options tftpd.logging on`

Rebooting the system

About rebooting

Rebooting the NetApp system is equivalent to halting and booting the system. During a reboot, the contents of the system's NVRAM are flushed to disk, and the system sends a warning message to CIFS clients.

Rebooting the system from the system console

You can reboot the system if the system console is displaying the command prompt. To reboot the system, complete the following step.

Step	Action
1	At the NetApp system prompt, enter the following command: `toaster> reboot`

Halting the system

Command to use

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

Reasons to use the halt command

The system stores requests it receives in nonvolatile random-access memory (NVRAM). For the following reasons, you should always execute the `halt` command before turning the system 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 system, 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 system without giving advance warning to CIFS users. Halting the CIFS service without giving CIFS users enough time to save their changes can cause data loss.

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 system is built into Windows NT and Windows XP.

Halting the system

To halt the system, complete the following step.

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

Managing Administrator Access 6

About this chapter

This chapter describes how to control administrator access to your system to provide increased security and auditing capability.

Topics in this chapter

This chapter discusses the following topics:

- "Using administrator accounts to control administrative access" on page 84
- "Managing passwords" on page 110

Using administrator accounts to control administrative access

About this section

This section contains the following topics:

- "About managing administrator access" on page 85
- "Managing users" on page 88
- "Managing groups" on page 93
- "Managing roles" on page 96
- "Listing and deleting users, groups and roles" on page 102
- "Administrative user creation examples" on page 108

About managing administrator access

Reasons for creating administrator accounts

You can use the default system administration account, or root, for managing a NetApp system. You can also create additional administrator user accounts, using the `useradmin` command. You might want to create administrator accounts for the following reasons:

- You can specify administrators and groups of administrators to have differing degrees of administrative access to your NetApp appliances.
- You can limit an administrator's access to specific NetApp systems by giving him or her an administrative account on only those systems.
- Having different administrative users allows you to display information about who is performing what commands on the system.

  The auditlog file keeps a record of all administrator operations performed on the system and the administrator who performed it, as well as any operations that failed due to insufficient capabilities.
- You assign each administrator to one or more groups whose assigned roles (sets of capabilities) determine what operations that administrator is authorized to carry out on the appliance.
- If an appliance running CIFS is a member of a domain or a Windows workgroup, domainuser accounts authenticated on the Windows domain can access the Data ONTAP Application Program Interface (API) using Windows RPC calls.

  For more information about authenticating users using Windows domains, see the section on user accounts in the CIFS chapter of the *File Access and Protocols Management Guide*.

What users, groups, roles, and capabilities are

You can grant different levels of administrative access to users of a NetApp appliance. Terms used in this section include the following:

user: An account that is authenticated on the NetApp system. Users can be placed into NetApp system groups to grant them capabilities on the NetApp system.

domainuser: A nonlocal user who belongs to a Windows domain and is authenticated by the domain. This type of user can only be put into NetApp system groups, thereby being granted capabilities on the NetApp system. This only works if CIFS has been set up on the appliance.

group: A collection of users and domainusers that can be granted one or more roles. Groups can be predefined, created, or modified. When CIFS is enabled, groups act as Windows groups.

role: A set of capabilities that can be assigned to a group. Roles can be predefined, created, or modified.

capability: The privilege granted to a role to execute commands or take other specified actions. Types of capabilities include:

- Login rights
- Data ONTAP CLI (command-line interface) rights
- Data ONTAP API (application programming interface) rights
- Security rights

How users are assigned capabilities

You cannot assign administrative roles or capabilities directly to administrative users or domainusers. Instead, you assign users to groups whose assigned roles match the capabilities that you want those users to be able to exercise.

- You can assign a set of capabilities to a role, then assign that role to a group. You then add an administrative user to the group that has the administrative role and capabilities that you want that user to have.
- You can also assign users and domainusers to some NetApp-predefined groups whose default roles match the roles that you want the users in question to exercise.

Naming requirements for users, groups and roles

When you name your users, groups and roles, you must meet these naming requirements:

- Names are case insensitive.
- Names can contain any alphanumeric character, a space, or a symbol that is not one of the following characters:

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

- If the name contains spaces or special characters, enclose the name in double quotes (“”) when you use it in a command.
- Do not give a user and a group the same name if you are using the NetApp system in a CIFS environment; doing so might cause some Windows functions to fail.

Windows special groups

Windows has some special groups it uses for security and administration purposes. Do not create administrative groups on your NetApp system with the same name as a Windows special group. The special Windows group names include the following names:

- System
- Everyone
- Interactive
- Network
- Creator/Owner
- Creator Group
- Anonymous Logon
- Authenticated Users
- Batch
- Dialup
- Service
- Terminal User

Changing capabilities of other groups and roles

If you are an administrator assigned to a group with capabilities that are equal to or greater than another group, you can make the following changes to that other group:

- Change the capabilities of the other group
- Change the capabilities of the roles within the other group
- Change the membership of the other group

Managing users

Creating users and assigning them to groups

You can use the `useradmin user add` or `useradmin user modify` commands to create or modify a user and assign that user to one or more predefined or customized groups. This gives that user the roles and capabilities associated with those groups.

Note

When you use the `useradmin user modify` command to modify the groups an existing user is assigned to, whatever groups the user was previously assigned to are replaced with the group or groups you supply in the command.

In Data ONTAP 7.0, user names are case insensitive. This means that you cannot create a user named “fred” if you already have a user named “Fred”.

To create a new user or assign an existing user to a custom group or predefined group, complete the following steps.

Step	Action
1	Use `useradmin user add` to create a new user and assign that user to a predefined group. Or, use `useradmin user modify` to assign an existing user to different groups. **`useradmin user {add|modify} user_name [-c comments] [-n full_name ] [-p password] -g {custom_group|Administrators|"Backup Operators"|Guests|"Power Users"|Users}[,...]`** *user_name* is the user whose name you want to assign to a customized or predefined group. The user name is case insensitive and can be up to 32 characters long. For more information about naming requirements, see "Naming requirements for users, groups and roles" on page 86. `[ -c` *comments* `]` specifies a maximum 128-character comment which can be viewed through the `useradmin list` command. Comments cannot contain a colon character (:). `[ -n` *full_name* `]` specifies the full name for the user. `[ -p` *password* `]` is the password required of the specified administrative user (used only for rsh access). This password must conform to the rules in the `security.passwd.rules.enable` option. If this option is set to Off, the restrictions are not enforced. *custom_group* is a customized group with roles assigned through the `useradmin group` command. For more information about creating custom groups, see "Assigning roles to groups" on page 93. `Administrators | "Backup Operators" | Guests | "Power Users" | Users` are groups predefined by Data ONTAP with default roles and capabilities. For a description of the predefined groups see "Predefined groups" on page 93.
2	To verify the success of your operation, enter the following command. **`useradmin user list user_name`** The specified user is listed along with the groups, roles, and capabilities that the user has inherited.

Example: The following command uses the predefined Administrators group and role definitions to create the user mollymulberry and grant her rights to invoke every type of administrative capability (login, CLI, API, and security).

```
useradmin user add molly -c "Molly Mulberry" -g Administrators
```

Granting access to Windows domain users

You can specify non-local administrative users to have administrative access to the appliance after authentication by a Windows Domain Controller, rather than by the appliance itself.

Note

Non-local users, or users added with the `useradmin domainuser add` command, cannot log in to a NetApp system using the login protocols (telnet, console, rsh, ssh or http-admin). Non-local users can log in only by using Windows RPC calls, and can use only the API capability once they are logged in. This is true regardless of the capabilities associated with any NetApp system group the non-local user has been placed in.

To grant a Windows domain user administrative access to an appliance, complete the following steps.

Step	Action
1	Use the `useradmin domainuser add` to assign a Windows domain user to a custom or predefinedgroup. **`useradmin domainuser add win_user_name -g {custom_group|Administrators|"Backup Operators"|Guests|"Power Users"|Users}[,...]`** *win_user_name* is the Windows domain user whose name or Security ID (SID) you want to assign to a customized or predefined group. This value can be in one of the following formats: ◆ `name` **Note** If you do not specify the domain name, the domain is the NetApp system, and the user is considered distinct from any user in the Windows domain with the same user name. ◆ `domain\name` ◆ `textual_sid_S-x-y-z` For more information about these formats, see the na_cifs_lookup(1) man page. *custom_group* is a customized group with roles assigned through the `useradmin group` command. For more information about creating customized groups, see "Predefined groups" on page 93. `Administrators | "Backup Operators" | Guests | "Power Users" | Users` are groups predefined by Data ONTAP with default roles and capabilities. For a description of the predefined groups see "Using administrator accounts to control administrative access" on page 88. **Example:** The following command line adds the user userjoe in the MyDomain domain to the Power Users group and effectively grants MyDomain\userjoe all administrator capabilities that are granted to the Power Users group through the roles that have been assigned to it. **`useradmin domainuser add MyDomain\userjoe -g "Power Users"`**

Step	Action
2	To verify the success of your operation, enter the following command. **useradmin domainuser list -g {*custom_group*|Administrators|"Backup Operators"|Guests|"Power Users"|Users}** The SID of the user in question is among those listed in the output of this command.

Changing another user's capabilities

You must be an administrator and your user account must be assigned to a group that has greater capabilities than the group the user is assigned to if you want to perform any of the following tasks:

- Change the capabilities of a user
- Change the comment about a user
- Change the full name of a user
- Change the name of a group

Note

You cannot create or change a group, a user, or a role, to have more capabilities than you have.

If you want to change the password of another user, your account must also be assigned to a group that has the `security-password-change-others` capability.

Managing groups

Predefined groups

If you want to, you can assign a user or domainuser to a predefined set of groups and roles provided by Data ONTAP. The predefined groups and roles are described in the following table:

Predefined Group	Default roles	Default privileges
Administrators	`admin`	Grants all CLI, API, login, and security capabilities.
Power Users	`power`	Grants the ability to ◆ Invoke all `cifs`, `exportfs`, `nfs`, `useradmin` CLI commands ◆ Make all `cifs` and `nfs` API calls ◆ Log in to Telnet, HTTP, rsh, and ssh sessions
Backup Operators	`none`	None
Users	`audit`	Grants the ability to make `snmp-get` and `snmp-get-next` API calls.
Guests	`none`	None
Everyone	`none`	None

Assigning roles to groups

To assign a group one or more predefined or customized roles and effectively give that group the capabilities associated with those roles, complete the following steps.

Note

When you use the `useradmin group modify` command to modify an existing group, whatever roles were previously assigned to that group are replaced with the roles you supply in the command.

Step	Action
1	Use the `useradmin group add` command to create a new group or the `useradmin group modify` command modify a group. **`useradmin group {add|modify} group_name [-c comments] [-r {custom_role|root|admin|power|audit}[,...]]`** *group_name* is the group that you want to create or to which you want to assign one or more roles. Group names are case insensitive and can be up to 256 characters. For more information about naming requirements, see "Naming requirements for users, groups and roles" on page 86. **Note** Do not create groups with the same name as any of the Windows special groups. For a list of the Windows special groups, see "Windows special groups" on page 87. *custom_role* is a customized role with capabilities assigned through the `useradmin role add` command. `root` | `admin` | `power` | `audit` are roles predefined by Data ONTAP with default capabilities. For a description of the predefined roles see "Predefined roles" on page 96. **Example:** The following command gives the group "admin users" capabilities associated with the admin role, and removes any roles previously assigned to the admin_users group. **`useradmin group modify "admin users" -r admin`**
2	Enter the following command to verify the success of your operation. **`useradmin group list group_name`** The roles and capabilities assigned to the group in question are listed in the output of this command.

Renaming a group

You can rename a group using the `useradmin group modify` command. To rename a group, complete the following step.

Step	Action
1	Enter the following command: **`useradmin group modify`** ***`group_name`*** **`-g`** ***`new_group_name`*** *group_name* is the name of the group you want to change. *new_group_name* is the name you want the group to have after the change. **Note** Do not attempt to rename a group with the same name as any of the Windows special groups. For a list of the Windows special groups, see “Windows special groups” on page 87.

Managing roles

Predefined roles

Data ONTAP provides the following predefined roles.

Role	Default capability assignments	Summary of default granted capabilities
root	`-a *`	Grants all possible capabilities.
admin	`-a cli-*`, `api-*`, `login-*`, `security-*`	Grants all CLI, API, login, and security capabilities.
power	`-a cli-cifs*`, `cli-exportfs*`, `cli-nfs*`, `cli-useradmin*`, `api-cifs-*`, `api-nfs-*`, `login-telnet`, `login-http-admin`, `login-rsh`, `login-ssh`	Grants the ability to ◆ Invoke all `cifs`, `exportfs`, `nfs`, `useradmin` CLI commands ◆ Make all `cifs` and `nfs` API calls ◆ Log in using telnet, HTTP, rsh, and ssh sessions
audit	`-a api-snmp-get`, `api-snmp-get-next`	Grants the ability to make `snmp-get` and `snmp-get-next` API calls.
none	None	Grants no administrative capabilities.

Assigning capabilities to roles

You use the `useradmin role add` or `useradmin role modify` commands to define and modify the capabilities of roles that can be assigned to a group.

Supported capability types

Data ONTAP supports four capability types, as described in the following table.

Capability Type	Description
`login`	Grants the specified role `telnet`, `console`, `rsh`, `ssh`, or `http-admin` login capabilities. `login-*` gives the specified role the ability to log in through all supported protocols. `login-`*protocol* gives the specified role capability to log in through a specified protocol. Supported protocols include: ◆ `login-telnet`—gives the specified role the ability to log in to the system using Telnet. ◆ `login-console`—gives the specified role the ability to log in to the system using the console. ◆ `login-rsh`—gives the specified role the ability to log in to the system using `rsh`. ◆ `login-ssh`—gives the specified role the ability to log in to the system using SSH. ◆ `login-http-admin`—gives the specified role the ability to log in to the system using HTTP.

Capability Type	Description
`cli`	Grants the specified role the ability to execute one or more Data ONTAP command line interface (CLI) commands. `cli-*` grants the specified role the capability to execute all supported CLI commands. `cli-`*cmd** gives the specified role the capability to execute all commands associated with the CLI command *cmd*. For example, the following command gives the specified role the capability to execute all `vol` commands: **`useradmin role modify status_gatherer -a cli-vol*`** **Note** Users with `cli` capability also require at least one `login` capability to execute CLI commands.
`security`	Grants the specified role security-related capabilities, such as the ability to change other users' passwords or to invoke the CLI `priv set advanced` command. `security-*` grants the specified role all security capabilities. `security-`*capability* grants the specified role one of the following specific security capabilities: ◆ `security-passwd-change-others` gives the specified role the capability to change the passwords of all users with equal or less capabilities. For more information about comparing capabilities, see "Changing another user's capabilities" on page 92. ◆ `security-priv-set-advanced` gives the specified role the capability to access the advanced CLI commands.

Capability Type	Description
`api`	Grants the specified role the capability to execute Data ONTAP API calls. `api-*` grants the specified role all api capabilities. `api-`*api_call_family* `-*` grants the specified role the capability to call all API routine in the family *api_call_family*. `api-`*api_call* grants the specified role the capability to call the API routine *api_call*. **Note** You have more fine-grained control of the command set with the `api` capabilities because you can give subcommand capabilities as well. Users with `api` capability also require the `login-http-admin` capability to execute API calls.

Creating a new role

To create a new role, complete the following steps:

Step	Action
1	Enter the following command: **`useradmin role add `*`role_name`*` [-c `*`comments`*`] -a capability1[,capability2...]`** *role_name* is the name of the role you want to create. Role names are case insensitive and can be 1-32 characters. For more information about naming requirements, see “Naming requirements for users, groups and roles” on page 86. *comments* is a short string you can use to document this role The *capability* parameters are the capabilities you want to grant to this new role.

Step	Action
2	To verify the success of the operation, enter the following command. **`useradmin role list role_name`** **Result:** The capabilities allowed for the specified role are listed.

Modifying an existing role

To modify the capabilities of an existing role, you use the `useradmin role modify` command.

Note

When you use the `useradmin role modify` command to modify an existing role, whatever capabilities were previously assigned to that role are replaced with the capabilities you supply in the command.

To modify an existing role, complete the following steps.

Step	Description
1	Enter the following command: **`useradmin role modify role_name [-c comments] -a capability1[,capability2...]`** *role_name* is the name of the role that you want to modify. *comments* is a short string you can use to document this role The *capability* parameters are the capabilities you want to grant to this role. **Example:** The following command line assigns the role "class2loginrights" telnet capabilities, console login capabilities, and all CLI capabilities, while removing any other capabilities that the role was granted previously. **`useradmin role modify class2loginrights -c "This role is for telnet and console logins" -a login-telnet,login-console,cli-*`**
2	To verify the success of the operation, enter the following command. **`useradmin role list role_name`** **Result:** The capabilities allowed for the specified role are listed.

Example of granting API command families

You can grant API capabilities for API command families. For example, to grant the myrole role only the capability to run CIFS commands, you use the following command:

```
useradmin role add myrole -a api-cifs-*
```

Listing and deleting users, groups and roles

Listing users, domainusers, groups, or roles

Use the following `useradmin` commands to display information for users, domainusers, groups, or roles.

Command	Description

`useradmin user list`	Lists all administrative users configured for this appliance. Each user entry includes the user name, comment information, a Data ONTAP-generated user ID number, and groups that each user belongs to. For example:
	`> useradmin user list`
	`Name: root`
	`Info: Default system administrator.`
	`Rid: 0`
	`Groups:`
	
	`Name: administrator`
	`Info: Built-in account for administering the filer`
	`Rid: 500`
	`Groups: Administrators`
	
	`Name: fred`
	`Info: This is a comment for fred.`
	`Rid: 131343`
	`Groups: Users`
	`...`

Command	Description
`useradmin user list` *user_name*	Lists the extended information for a specific administrator. The extended information includes the user name, comment information, the groups that the user belongs to, a Windows-based name if the user has one, a Data ONTAP-generated user ID number, and effective allowed capabilities. For example: **>useradmin user list fred** Name: fred Info: This is a comment for fred Rid: 131343 Groups: Users Full Name: Allowed Capabilities: login-http-admin,api-snmp-get,api-snmp-get-next
`useradmin user list -g` *grp_name*	Lists information for all users assigned to a specified group. For example: **> useradmin user list -g Admistrators** Name: wilma Info: Rid: 131074 Groups: Administrators Name: marshall Info: Rid: 131454 Groups: Administrators ...

Command	Description

`useradmin domainuser list -g` *`group_name`*	Lists the SIDs of all Windows domain administrative users assigned to a specified group.
	
	To list the user name, comment information, and the groups that each user belongs to, follow up with `cifs lookup` and `useradmin user list` commands. For example:
	
	`>useradmin domainuser list -g administrators`
	`List of SIDS in administrators`
	`S-1-7-24-1214340929-620487827-8395249115-512`
	`S-1-7-24-1838915891-154599588-1081798244-500`
	`For more information about a user, use the 'cifs lookup' and 'useradmin user list' commands.`
	
	`>cifs lookup S-1-7-24-1214340929-620487827-8395249115-512`
	`name = MBS-LAB\Domain Admins`
	
	`>cifs lookup S-1-7-24-1838915891-154599588-1081798244-500`
	`name = ZND\Administrator`
	
	`>useradmin user list Administrator`
	`Name: Administrator`
	`Info: Built-in account for administering the filer`
	`Rid: 500`
	`Groups: Administrators`
	`Full Name:`
	`Allowed Capabilities: login-*,cli-*,api-*,security-*`
	
	Note
	The RID value of 500 for the Administrator user corresponds to the last number in the Administrator user's SID.

Command	Description

`useradmin group list`	Lists all the administrative user groups configured for this appliance. Each group entry includes the group name, comment information, Data ONTAP-generated user ID number, and every role associated with that group. For example:
	`> useradmin group list`
	`Name: Administrators`
	`Info: Members can fully administer the filer`
	`Rid: 544`
	`Roles: admin`
	
	`Name: Backup Operators`
	`Info: Members can bypass file security to backup files`
	`Rid: 551`
	`Roles: none`
	
	...

`useradmin group list` *`group_name`*	Lists the extended details for a specified single group. An extended entry for a single group includes the group name, comment information, roles assigned to that group, and allowed capabilities. For example:
	`>useradmin group list Administrators`
	`Name: Administrators`
	`Info: Members can fully administer the filer.`
	`Roles: admin`
	`Allowed Capabilities: login-*,cli-*,api-*,security-*`

Command	Description

`useradmin role list`	Lists all the roles configured for this appliance. Each role entry lists the role name, comment information, and allowed capabilities. For example:
	`>useradmin role list`
	`Name: admin`
	`Info:`
	`Allowed Capabilities: login-*,cli-*,api-*,security-*`
	
	`Name: audit`
	`Info:`
	`Allowed Capabilities: login-http-admin,api-snmp-get,api-snmp-get-next`
	
	`Name: none`
	`Info:`
	`Allowed Capabilities:`
	
	`...`

`useradmin role list` *`role_name`*	Lists the information for a single specified role name. For example:
	`> useradmin role list admin`
	`Name: admin`
	`Info: Default role for administrator privileges.`
	`Allowed Capabilities: login-*,cli-*,api-*,security-*`

Deleting users, domainusers, groups, or roles

Use the following commands to delete users, domainusers, groups, or roles.

Command	Description
`useradmin user delete` *user_name*	Deletes the specified user from the appliance. The `useradmin user delete` command deletes any local user except for “root” and “snmp”. The “snmp” administrative user can be disabled by modifying the audit role or “User” group. **Note** You cannot delete or modify a user with greater capabilities than you have. For more information about comparing capabilities, see “Changing another user’s capabilities” on page 92.
`useradmin domainuser delete` *win_user_name* `-g` *group1,[group2,...]*	Removes the specified user from the specified group or groups. This command does not delete the user from the domain. **Note** If you want to completely delete a user from the NetApp system, use the `useradmin user delete` command instead.
`useradmin group delete` *group_name*	Deletes the specified group from the appliance. **Note** All users must be removed from a group before the group itself can be deleted.
`useradmin role delete` *role_name*	Deletes the specified role from the system. **Note** A role that is still assigned to a group cannot be deleted.

Administrative user creation examples

Example creation of a user with custom capabilities

The following `useradmin` commands create a user with a limited and specialized set of administrator capabilities. The commands carry out the following operations:

- Create the following roles:
  - "only_ssh" is allowed to log in only via `ssh`
  - "qtree_commands" can run any `qtree` command in the CLI.
- Create the following group:
  - "ssh_qtree_admins" is allowed to log in only via `ssh` and run the `qtree` commands in the CLI, using the two roles created in the previous step.
- Create a user, "wilma" and assign that user to the ssh_qtree_admins group. As a member of the ssh_qtree_admins group, user wilma now inherits the capabilities from the roles assigned to that group.
- Display the details and capabilities inherited by the new user wilma.

```
> useradmin role add only_ssh -a login-ssh
Role <only_ssh> added.
> Thu Apr 22 10:50:05 PDT [tpubs-cf1:
useradmin.added.deleted:info]: The role 'only_ssh' has been added.
>useradmin role add qtree_commands -a cli-qtree*,api-qtree-*
Role <qtree_commands> added.
> Thu Apr 22 10:51:51 PDT [tpubs-cf1:
useradmin.added.deleted:info]: The role 'qtree_commands' has been
added.
>useradmin group add ssh_qtree_admins -r only_ssh,qtree_commands
Group <rsh_qtree_admins> added.
> Thu Apr 22 10:53:07 PDT [tpubs-cf1:
useradmin.added.deleted:info]: The group 'ssh_qtree_admins' has
been added.
>useradmin user add wilma -g ssh_qtree_admins
New password:
Retype new password:
User <wilma> added.
> Thu Apr 22 10:54:43 PDT [tpubs-cf1:
useradmin.added.deleted:info]: The user 'wilma' has been added.
> useradmin user list wilma
Name: wilma
Info:
Rid: 131074
```

```
Groups: ssh_qtree_admins
Full Name:
Allowed Capabilities: login-ssh,cli-qtree*,api-qtree-*
```

Example creation of a user with no administrative capabilities

In a CIFS environment, you might want to create users on the NetApp system that are in local groups but do not have console access or any administrative capabilities on the NetApp system. These users would still have the file access permissions granted by the local groups.

To add a user with no console access or administrative capabilities to a NetApp system, complete the following steps.

Step	Action
1	Enter the following command: **`useradmin user add user_name -g "Guests"`** *user_name* is the user name for the new user
2	Enter the user's password when prompted.
3	To verify that you have created the user with no capabilities, enter the following command: **`useradmin user list user_name`** **Result:** "Allowed Capabilities" should be blank.

Managing passwords

This section describes managing passwords on the NetApp appliance. The following topics are included:

- "Changing the system password" on page 110
- "Changing a local user account password" on page 111
- "Changing password requirements" on page 111

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 NetApp system	1. Enter the following command:
		`passwd`
		2. Enter the system account name:
		`root`
		3. Enter the existing system password (not required if you are logged in as root or have the `security-passwd-change-others` capability).
		4. Enter a new password, and then enter it a second time to confirm it.

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

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 NetApp system	**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 NetApp system	Enter the following command: `rsh` ***`filer_name`*** `passwd` ***`new_password username`***

Changing password requirements

By default Data ONTAP requires passwords be a minimum of 8 characters long, with at least two alphabetic and at least one numeric characters. Data ONTAP provides security options that allow you to change these password requirements and indicate whether the requirements apply to root and Administrator users.

The following set of security options is available:

Password rule option	Description
`options security.passwd.rules. enable {on|off}`	Specifies whether a check for password composition is performed when new passwords are specified. If this option is set to On, passwords are checked against the rules specified in this table, and the password rejected if it doesn't pass the check. If this option is set to Off, the check is not performed. The default value for this option is On. By default, this option does not apply to the users "root" or "Administrator" (the NT Administrator account).
`options security.passwd.rules. everyone {on|off}`	Specifies whether a check for password composition is performed for the "root" and "Administrator" users. If the `security.paswd.rules.enable` option is set to Off, this option does not apply. The default value for this option is Off.
`options security.passwd.rules. maximum` *max_num*	Specifies the maximum number of characters a password can have. **Note** A maximum of 16 characters of the password are used to match the password, even if this option is set to a value greater than 16. Users with passwords longer than 14 characters will not be able to log in via the Windows interfaces, so if you are using Windows, NetApp recommends this option be set to no higher than 14. The default value for this option is 256.
`options security.passwd.rules. minimum` *min_num*	Specifies the minimum number of characters a password must have. The default value for this option is 8.

Password rule option	Description
`options security.passwd.rules.minimum.alphabetic` *min_num*	Specifies the minimum number of alphabetic characters a password must have. The default value for this option is 2.
`options security.passwd.rules.minimum.digit` *min_num*	Specifies the minimum number of digit characters a password must have. These are numbers from 0 to 9. The default value for this option is 1.
`options security.passwd.rules.minimum.symbol` *min_num*	Specifies the minimum number of symbol characters (white space and punctuation characters) a password must have. The default value for this option is 0.

Performing General System Maintenance 7

About this chapter

This chapter describes some general maintenance tasks you might need to perform to manage your NetApp system.

Topics in this chapter

This chapter discusses the following topics:

- "Aggregate snapshot management" on page 116
- "Managing licenses" on page 119
- "Setting the system date and time" on page 121
- "Synchronizing the system time" on page 123
- "Understanding core files" on page 126
- "Configuring message logging" on page 128
- "Configuring audit logging" on page 132
- "Configuring NetApp system startup" on page 134
- "NetApp system configuration backup and cloning" on page 139
- "UPS management" on page 142

Aggregate snapshot management

What an aggregate snapshot is

An aggregate snapshot is a point-in-time, read-only image of an aggregate. It is similar to a volume snapshot, except that it captures the contents of the entire aggregate, rather than any particular volume.

For more information about snapshots, see the *Data Protection Online Backup and Recovery Guide*.

What you can do with aggregate snapshots

You use aggregate snapshots when the contents of an entire aggregate need to be recorded. However, you do not restore data directly from an aggregate snapshot. To restore data, you use a volume snapshot.

You use aggregate snapshots in the following situations:

- If you are using MetroCluster or RAID SyncMirror and you need to break your mirror, you take an aggregate snapshot prior to breaking the mirror to decrease the time it takes to resync the mirror later.
- If you are making a global change to your NetApp system, and you want to be able to restore the entire system state if the change produces unexpected results, you take an aggregate snapshot before making the change.
- If the aggregate file system becomes inconsistent, aggregate snapshots can be used by NetApp technical support to restore the file system to a consistent state.

How you manage aggregate snapshots

The commands you use to manage aggregate snapshots are the same as the ones you use to manage volume snapshots, with the `-A` flag added. However, be sure to read the rest of this section to understand the recommended ways to manage your aggregate snapshots.

How you create aggregate snapshots

Usually, you do not need to create aggregate snapshots manually. A schedule is automatically set up to generate new aggregate snapshots periodically. NetApp recommends that you do not change the aggregate snapshot schedule.

If you do need to create an aggregate snapshot manually, you use the same command as you would for a volume snapshot, except that you add the `-A` flag. For more information on creating snapshots, see the *Data Protection Online Backup and Recovery Guide.*

About the aggregate snapshot reserve

Just as there is space reserved for volume snapshots in their volume (the volume snapshot reserve), there is space reserved for aggregate snapshots in the aggregate. This space is called the aggregate snapshot reserve.

The default size of the aggregate snapshot reserve is 5 percent of the aggregate size. For example, if the size of your aggregate is 500 GB, then 25 GB is set aside for aggregate snapshots.

Note

Unlike volume snapshots, aggregate snapshots cannot consume any space outside of their snapshot reserve.

About aggregate snapshot autodeletion

As more and more data blocks in the aggregate are changed, the aggregate snapshot reserve gradually becomes full. Because aggregate snapshots usually do not need to be preserved for long periods of time (you usually need only the most recent aggregate snapshot), Data ONTAP automatically deletes the oldest aggregate snapshots to recover space in the aggregate snapshot reserve.

When an aggregate snapshot is automatically deleted, a message similar to this one is logged:

```
Sun May 23 15:10:16 EST [wafl.snap.autoDelete:info]: Deleting
snapshot ‘nightly.0’ in aggregate ‘aggr1’ to recover storage
```

How you manage your aggregate snapshot reserve

Usually, the default aggregate snapshot reserve of 5 percent is sufficient. However, if you find that aggregate snapshots are being created and deleted often enough to affect system performance, then you should increase the aggregate snapshot reserve.

Note

NetApp recommends that you do not *decrease* the size of the aggregate snapshot reserve below the default of 5 percent.

How you manage automatic snapshot deletion

In most cases you should leave automatic aggregate snapshot deletion enabled. If this option is turned off for a particular aggregate, then every volume in that aggregate requires up to 2 times its size in order to satisfy a space guarantee of `volume`.

However, in some specific situations, you may need to disable automatic aggregate snapshot deletion temporarily. For example, if you want to create an aggregate snapshot before breaking your RAID SyncMirror, and you want to make sure that snapshot is not automatically deleted. To disable automatic aggregate snapshot deletion, you use the `aggregate option` command.

For example, to turn off aggregate snapshot autodeletion for the aggregate myAggr, you would use the following command:

```
aggr options myAggr snapshot_autodelete off
```

Note

If you do not have sufficient free space in your aggregate to satisfy the new space requirements when you turn off aggregate snapshot autodeletion, then space guarantees will be disabled for one or more of your volumes. For this reason, you should plan to re-enable aggregate snapshot autodeletion as quickly as possible.

About disabling aggregate snapshots

You can turn off aggregate snapshots, using the same `nosnap` option that you would for volume snapshots. However, NetApp recommends that you leave aggregate snapshots enabled in case you need any low-level file system repair.

Managing licenses

About license codes

A license code is a string of characters that is unique to a particular NetApp system (such as ABCDEFG). You receive license codes for every protocol and option that you purchase. Not all purchased license codes are installed on a NetApp system 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 NetApp 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

Adding 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 NetApp.

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 list of possible NetApp services. **Note** To display the complete list of services, use the `license` command with no parameters.

Setting the system 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.

Note

Keeping the system date and time correct is important to ensure that the NetApp system can service requests correctly. To keep your NetApp system's date and time correct automatically, see "Synchronizing the system time" on page 123.

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

Step	Action
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

If you use the `date` or `rdate` command to set a NetApp system's date earlier when SnapMirror is running, snapshots can appear out of sequence. When this occurs, SnapMirror assumes that the older snapshot was created before the newer one, and asks 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 NetApp system completes the changes.
- Change the date prior to the next scheduled SnapMirror transfer.

Synchronizing the system time

About synchronizing system time

You can use the `timed` daemon to automatically keep the system time for your NetApp system synchronized with a time server. NetApp recommends that you use this feature because problems can occur when the NetApp system clock is inaccurate.

About time servers

To automatically keep your NetApp system time synchronized, you need the name of at least one time server. NetApp recommends that you supply the name of more than one time server in case one becomes unavailable.

About the protocols

There are two protocols you can use for time synchronization: SNTP and `rdate`.

SNTP (Simple Network Time Protocol) is more accurate; therefore, it is the preferred protocol. You can get a list of public NTP (Network Time Protocol) time servers (used for SNTP) from the NTP.Servers Web at http://ntp.isc.org/bin/view/Servers/WebHome.

If you can't access an SNTP server, you can use `rdate`. Many Unix servers can function as an `rdate` server; see your system administrator to set up or identify an `rdate` server in your environment.

Synchronizing system time

To set the `timed` daemon to keep the NetApp system time synchronized with the time server, complete the following steps.

Step	Action
1	If the current time for the NetApp system is not fairly close to the actual time, use the date command to set the system time to the correct time. For more information about setting the system clock, see "Setting the system date and time" on page 121.

Step	Action
2	At the command line, set the appropriate timed options using the options command. At a minimum, you must set the `proto` option to use either `sntp` or `rdate` (`sntp` is the preferred protocol), and set the `servers` option to at least one valid time server for the protocol you select. For more information about the protocols, see "About the protocols" on page 123. For more information about the timed options, see the na_options(1) man page or "About the timed options" on page 124.
3	Enter the following command to enable the timed daemon: **`options timed.enable on`**

About the timed options

The following table summarizes the available timed options and their default value. For more detailed information on the timed options, see the na_options(1) man page.

Timed option	Function	Allowable values	Default value
`enable`	Enables time synchronization.	◆ `On` ◆ `Off`	`On`
`log`	Specifies whether time changes should be logged to the console.	◆ `On` ◆ `Off`	`Off`
`max_skew`	Specifies the maximum allowable skew between the system time and the time server time. If the skew exceeds this value, synchronization does not occur.	◆ `ns` ◆ `nm` ◆ `nh`	`30m`

Timed option	Function	Allowable values	Default value
`proto`	Specifies the protocol used to synchronize the time.	◆ `rtc` (internal Real-Time Clock) ◆ `rdate` RFC 868 ◆ `sntp` RFC 2030 (preferred)	`rtc`
`sched`	Specifies the timed synchronization schedule.	◆ `hourly` ◆ `multihourly` ◆ `daily` ◆ `custom`	`1h` (hourly)
`servers`	Specifies up to five time servers used by the timed daemon.	**Example:** `times1,times2.ntap.com,10.15.46.92`	null string
`window`	Specifies a window of time around the synchronization time when the synchronization can occur.	◆ *n*s ◆ *n*m	`0s`

Example clock synchronization

The following example configures `timed` to use the SNTP protocol with the default hourly synchronization schedule.

```
filer1> date
Thu Dec  9 13:49:10 PST 2004
filer1> options timed.proto ntp
filer1> options timed.servers pool.ntp.org,10.15.46.92
filer1> options timed.enable on
```

Understanding core files

About core files

When a hardware or software failure causes the NetApp system to crash, the system creates a core file that NetApp technical support can use to troubleshoot the problem. The NetApp system 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 of any working disk owned by the local NetApp system. 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 NetApp system belongs to a cluster and `options cf.takeover.on_panic` is enabled, a core dump file is written to a spare disk on that system.

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.

Options that control core dump file creation

There are two options you can use to control core dump file creation, as listed in the following table:

Option	Description	Default value
`coredump.dump.attempts`	Controls how many attempts are made to create a core dump file.	`2`
`coredump.metadata_only`	Specifies whether the buffers containing user data are dumped.	`On` (user data buffers are not dumped)

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

Automatic technical support notification

Your NetApp system sends e-mail automatically to NetApp technical support upon each system reboot, if the AutoSupport feature is enabled and configured correctly. NetApp technical support uses the AutoSupport message and the core file to troubleshoot the problem.

Note

If you have disabled AutoSupport e-mail, you should contact NetApp technical support when your system creates a core file.

Configuring message logging

What message logging is

The NetApp system maintains messages in the /etc/messages file on its root volume.

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

Note

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 of important events.

Where messages are sent

Message logging is done by a `syslogd` daemon. The /etc/syslog.conf configuration file on the NetApp system'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.

Accessing your message log files

You can access the /etc/messages files using your NFS or CIFS client, or using HTTP. For more information, see "Accessing the default directories on the NetApp system" on page 54.

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).

The /etc/syslog.conf file

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 NetApp system daemons, such as `rshd`
`kern`	Messages from the NetApp system 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
`*`	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>` *action* For more information about these parameters, see “The /etc/syslog.conf file” on page 129.
3	Save and close the /etc/syslog.conf file. **Result:** The changes you made to the syslog.conf file are read automatically and are reflected in the message logging.

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.

Note

You can also configure auditing specific to your file access protocol. For more information, see the *File Access and Protocols Management Guide*.

Accessing your audit log files

You can access the auditlog files using your NFS or CIFS client, or using HTTP. For more information, see "Accessing the default directories on the NetApp system" on page 54.

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`.

The system saves auditlog files for six weeks, unless any auditlog file reaches the maximum size, in which case the oldest auditlog file is discarded.

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.

Step	Action
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 10,000,000 (about 10 MB).

Configuring NetApp system startup

About this section

This section describes how to customize your system startup by changing the /etc/rc file. It contains the following information:

- "About the /etc/rc file" on page 135
- "Editing the NetApp system's boot configuration file" on page 137
- "Recovering from /etc/rc errors" on page 138

About the /etc/rc file

Startup commands are stored in /etc/rc file

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

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 NetApp system to

- Communicate on your network
- Use the NIS and DNS services
- Save the core dump that might exist if the NetApp system crashed before it was booted

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
savecore
```

Explanation of default /etc/rc contents

This section explains the sample /etc/rc file:

Description: `hostname toaster`

Sets the NetApp system host name to “toaster.”

Description:
`ifconfig e0 `hostname`-0`

```
ifconfig e1 `hostname`-1
ifconfig f0 `hostname`-f0
ifconfig a5 `hostname`-a5
```

Sets the IP addresses for the NetApp system 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 NetApp system 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 set static routes for the NetApp system by adding route commands to the /etc/rc file. The network address for MyRouterBox must be in /etc/hosts on the root volume.

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

Description: `routed on`

Starts the routing daemon.

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

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.

For more information about core files, see “About core files” on page 126.

Editing the NetApp system's boot configuration file

What the NetApp system's boot configuration file contains

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

Editing the NetApp system's boot configuration file

To edit the NetApp system'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. Doing so can cause problems when the NetApp system boots if CIFS is not fully initialized or the commands cause deadlocks between the /etc/rc file and CIFS.
3	Save the edited file.
4	Reboot the NetApp system to test the new configuration. If the new configuration does not work as you want, repeat Step 2 through Step 4.

Recovering from /etc/rc errors

Why the NetApp system becomes inaccessible to the administration host

The NetApp system 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 `ifconfig` command. The NetApp system is inaccessible because it is not on the network.
- You improperly export NetApp system directories to the NFS client that is the administration host. The NetApp system is inaccessible because you cannot mount the system 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 NetApp system root directory to the administration host.
	A CIFS environment	Add a share to the NetApp system root directory.
3	Edit the NetApp system /etc/rc file from the administration host.	
4	Reboot the NetApp system.	
5	If the changes do not correct the problem, repeat Step 1 through Step 4.	

NetApp system configuration backup and cloning

About NetApp system configuration backup and cloning

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

Backing up a NetApp system configuration

When you use the `config` command to back up a NetApp system configuration, the system 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 NetApp system 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 NetApp system-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 NetApp system configuration

To clone the configuration of one NetApp system 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 NetApp system from which you want to clone the configuration. *username* is the login name of an administrative user on the remote NetApp system. *password* is the remote user password. **Example:** The following is an example of the `config clone` command cloning the tpubs-dot configuration to the NetApp system toaster. `toaster> config clone tpubs-dot root:hello`

About restoring a NetApp system configuration

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

Restoring a NetApp system configuration

To restore NetApp system 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 NetApp system-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`

Step	Action
2	Reboot the system to run commands in the /etc/rc file.

Comparing NetApp system configurations and backup configuration files

To compare a NetApp system'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}`
	output_file is the name of the file to contain the differences. If you omit this parameter, the output of the command is printed to the console.
	config_file1 is the name of the first configuration file you want to compare.
	config_file2 is the name of the second configuration file you want to compare.
	Example 1: The following example compares the NetApp system'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 -o diff.txt 11_05_2001 11_15_2001`

UPS management

About UPS management

Data ONTAP enables you to register and monitor the status of UPS (Uninterruptible Power Supply) devices you are using with your NetApp system. In addition, you can configure the timing of certain Data ONTAP events when a power loss occurs.

About registering, enabling and disabling UPS devices

You use the `ups` command to register UPS devices with Data ONTAP, and to review their status. For more information about the `ups` command, see the `na_ups(1)` man page.

About the UPS shutdown options

Data ONTAP provides two configurable values to help you manage your NetApp system in case of a power outage:

- `warningtime`

  The `warningtime` option specifies when Data ONTAP generates a warning SNMP trap, AutoSupport message and log message.

  The default value of the `warningtime` option is 300 seconds (5 minutes).

- `criticaltime`

  The `criticaltime` option specifies when Data ONTAP generates another SNMP trap, AutoSupport message and log message, and then starts shutting down the NetApp system.

  The default value of the `criticaltime` option is 60 seconds (1 minute).

About the shutdown process

When a power loss occurs, the UPS device begins supplying power to your NetApp system from its batteries. The UPS can only supply power as long as its batteries still have enough charge; the UPS is there to give you time to shut down your NetApp system cleanly.

The shutdown process is outlined in the following tasks.

Note

If you do not have AutoSupport enabled, the AutoSupport messages will not be generated.

Step	Task
1	When the power loss occurs, an SNMP trap, AutoSupport message, and log messages are generated alerting you that the power loss has occurred.
2	When the UPS has `warningtime` seconds of battery life remaining, Data ONTAP generates another SNMP trap, AutoSupport message, and log message.
3	When the UPS has `criticaltime` seconds of battery life remaining, Data ONTAP generates another SNMP trap, AutoSupport message, and log message and starts shutting down the NetApp system. **Note** The `criticaltime` notifications may not be sent, depending on system load.

About using the default shutdown event timings

For many environments, you can simply use the default values of 5 minutes for `warningtime` and 1 minute for `criticaltime`. However, NetApp recommends that you make sure these values are set appropriately for your environment to avoid any data loss in case of a power outage. The `warningtime` value should give you enough time to do whatever manual processes you need to do prior to system shutdown, and `criticaltime` should provide enough time for the system to shut down cleanly.

Factors that might influence shutdown event timing for your environment

Factors that can affect the optimal value of these options for your environment include:

- UPS battery availability

  If your UPS cannot support the default timing values, then your NetApp system will not be able to shut down cleanly.

- NetApp system workload

  If you have a large number of users, a large number of CIFS sessions, or any other workload factors that require a longer time to shut down, you need to increase the warning and critical time values to ensure that the system has sufficient time to shut down cleanly.

- Company policies and procedures

  You may need to change the shutdown event timings to adhere to a protocol or requirement in place at your company.

Configuring the shutdown event timings

If, after reviewing this documentation, you decide that you need to change these values, you can do so using the `registry` command.

Note

NetApp strongly recommends that you contact NetApp technical support before changing the shutdown event timing values.

Using AutoSupport 8

About this chapter

This chapter discusses the AutoSupport feature that enables Data ONTAP to automatically send information about your NetApp system to NetApp technical support and to other recipients you specify. This feature provides you with customized real-time support to monitor the performance of your system.

Topics in this chapter

This chapter discusses the following topics:

- "Learning about AutoSupport" on page 146
- "Configuring AutoSupport" on page 148
- "Troubleshooting AutoSupport" on page 154
- "Interpreting AutoSupport messages" on page 157

Learning about AutoSupport

About AutoSupport

The `autosupport` daemon monitors the NetApp system's operations and sends automatic messages to NetApp technical support to alert it to potential system problems. If necessary, NetApp technical support contacts you at the e-mail address that you specify to help resolve a potential system problem.

The following list outlines facts you should know about AutoSupport:

- The `autosupport` daemon is enabled by default on the NetApp system.
- AutoSupport messages are generated
  - When events occur on the NetApp system that require corrective action from the system administrator or NetApp technical support
  - When the NetApp system reboots
  - Once a week, early Sunday morning, at approximately midnight
  - When you initiate a test message using the `autosupport.doit` option
- The system can send AutoSupport messages by SMTP, HTTP, or HTTPS (the Secure Sockets Layer [SSL] Internet transport protocol). HTTPS is the default.
- If an AutoSupport message cannot be sent successfully, an SNMP trap is generated.

Note

For more information about AutoSupport, see the NOW site at http://now.netapp.com/autosupport/.

AutoSupport is enabled by default

AutoSupport is enabled by default when you configure your NetApp system for the first time. After a grace period of 24 hours, AutoSupport messages start being generated. You can disable AutoSupport at any time using the `autosupport.enable` option, but NetApp strongly recommends that you leave it enabled. Enabling AutoSupport can significantly speed problem determination and resolution should a problem occur on your NetApp system.

AutoSupport transport protocols

AutoSupport supports the following type of transport protocols for delivering AutoSupport messages to NetApp:

- HTTPS

- HTTP
- SMTP

The default transport protocol is HTTPS. Because SMTP can introduce limitations on message length and line length, NetApp recommends that you use HTTPS or HTTP for your AutoSupport transport protocol.

AutoSupport messages using HTTP/HTTPS

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

To use HTTP or HTTPS to send AutoSupport messages, you may need to configure an HTTP or HTTPS proxy.

AutoSupport daemon requires an external mail host if you use SMTP

The NetApp system 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 NetApp system uses the mail host's mail server to send periodic e-mail messages automatically to NetApp technical support about the system's status. You can configure AutoSupport to use one or more mail hosts.

Note

Make sure that mailhosts in the AutoSupport email delivery pathway are configured to send and receive the 8-bit MIME encoding.

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 session and share 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 the configured recipient addresses, indicating that the test AutoSupport message was received successfully. If the message is not sent, perform the troubleshooting procedure as outlined in "Troubleshooting AutoSupport" on page 154.
`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. For example, postmaster@mycompany.com.
`autosupport.mailhost host1[, ..., host5]`	Defines up to five mail host names. The host names should be entered as a comma-separated list with no spaces in between. The default is an empty string.

AutoSupport option	Description
`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.noteto` *`address1`*`[, ..., `*`address5`*`]`	Defines the list of recipients for the AutoSupport short note e-mail. Up to five e-mail addresses are allowed. Enter the addresses as a comma-separated list with no spaces in between. The default is an empty list to disable short note e-mails.
`autosupport.retry.count` *`#retries`*	Defines the number of times the NetApp system will try to resend the AutoSupport notification before giving up, if previous attempts have failed. Retries can be between 5 and 4,294,967,295. The default is 15.
`autosupport.retry.interval` *`interval`*	Defines the time to wait before trying to resend a failed 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 24 hours. The default is 4m (4 minutes).
`autosupport.support.enable [ on | off ]`	Enables and disables the AutoSupport notification to NetApp. The default is On.

AutoSupport option	Description
`autosupport.support.proxy`	Allows you to set an HTTP proxy if necessary. This is useful only if `autosupport.support.transport` is HTTP or HTTPS. The default value for this option is the empty string. **Note** The value you use for this option is site-specific; see your IT department for the correct value for your site.
`autosupport.support.to`	Indicates where AutoSupport notifications to NetApp are sent if `autosupport.support.transport` is SMTP. This option is read-only and is shown for informational purposes only.
`autosupport.support.transport [ http | https | smtp ]`	Defines the type of delivery for AutoSupport notifications to NetApp. The default is `https`.
`autosupport.support.url`	Indicates where AutoSupport notifications to NetApp 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 option	Description
`autosupport.to` *`address1`*`[, ..., `*`address5`*`]`	Defines the list of recipients for the AutoSupport e-mail notification. Up to five e-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	Enter the following command, using any of the AutoSupport option commands shown in "AutoSupport options" on page 148: **`options autosupport.option arguments`** *autosupport.option* is one of the options shown in the table. *arguments* refers to any required or optional arguments shown in the 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 receive a return message from NetApp, indicating that the AutoSupport process is working correctly. For more details, see the information on the `autosupport.doit` command in “AutoSupport options” on page 148.

Troubleshooting AutoSupport

Troubleshooting AutoSupport over HTTP or HTTPS

If the AutoSupport test message is not being sent, and you are using HTTP or HTTPS, try to resolve the problem by completing the following steps.

Step	Action
1	Ensure that DNS is enabled and configured correctly on your system by entering the following command on the NetApp system: **`dns info`**
2	Ensure that the system is routing out to the Internet successfully by entering the following command: **`traceroute -p`** *port* **`support.netapp.com`** Generally, *port* is 80 if you are using HTTP, or 443 if you are using HTTPS.

Troubleshooting AutoSupport over SMTP

If the AutoSupport test message is not being sent, and you are using SMTP, try to resolve the problem by completing the following steps.

Step	Action
1	Set debug level in the syslog.conf file by creating the following line in the /etc/syslog.conf file: `*.debug /etc/messages`
2	Initiate AutoSupport by using the `doit` option. **Result:** An AutoSupport error message is displayed.
3	Check that the mail host specified in the options is a host that the NetApp system can talk to by entering the following command on the NetApp system: `ping` *`mailhost_name`* *mailhost_name* is the name of the mail host specified in the AutoSupport options.

Step	Action
4	Log on to the host designated as the mail host and make sure that it can serve SMTP requests by entering the following command (25 is the listener SMTP port number): `netstat -aAn|grep 25` **Result:** A message will appear, similar to the following text: ff64878c tcp 0 0 *.25 *.* LISTEN.

5	Telnet to the SMTP port from some other host by entering the following command:
	`telnet mailhost 25`
	Result: A message will appear, similar to the following text:
	`Trying 192.9.200.16 ...`
	`Connected to filer.`
	`Escape character is '^]'.`
	`220 filer.netapp.com Sendmail 4.1/SMI-4.1 ready at Thu, 30 Nov 95 10:49:04 PST`

6	If you are still experiencing problems, use a local area network (LAN) trace.

If the AutoSupport message is too large

The AutoSupport message contains the /etc/messages file. If that file becomes too large, the size of the AutoSupport message can cause problems, especially if you are using the SMTP transport protocol.

To keep the size of AutoSupport messages down, complete the following steps.

1	Make sure that the /etc/messages file is being rotated on a weekly basis as expected. If necessary, rotate the file manually.
2	Make sure you have your /etc/syslog.conf file to capture only system messages of level WARN or above in the /etc/messages file. For more information about editing the /etc/syslog.conf file, see the na_syslog.conf(5) man page.

3	Consider using HTTP or HTTPS for your AutoSupport transport protocol.
4	If the above steps do not resolve the problem, you can set the `autosupport.content` option to `minimal`. **Note** Using this setting is not recommended, as it may impact the quality of your support from NetApp.

Interpreting AutoSupport messages

About AutoSupport messages

AutoSupport messages can help you understand the status and operations of your NetApp system. This section helps you understand and interpret the AutoSupport messages generated by your system.

Detailed information

This section contains the following topics:

- "About AutoSupport events" on page 158
- "Contents of AutoSupport messages" on page 159

About AutoSupport events

Event and log-level descriptions

The NetApp system sends AutoSupport messages to NetApp technical support about your system after any of several events. The AutoSupport message includes a log level that indicates the priority assignment from NetApp 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 log levels in the subject lines 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 NOW site and select Product Documentation > AutoSupport Features and Functions > AutoSupport Message Matrices.
2	On the Message Matrices page under Select a Release, select your version of Data ONTAP and click View Matrix. **Result:** 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 partially displayed in the `AutoSupport.content Minimal` format.
- Commands marked with a dagger (†) are advanced commands. It you want more information about the output of these commands, call NetApp technical support.
- Commands marked with a double dagger (‡) print as shown if the `verbose` option is *not* set for them.

  Commands that show both 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 NetApp system
- Encrypted software licenses*
- Host name of the NetApp system*
- SNMP contact name and location (if specified)*
- Console encoding type

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

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

- Checksum status
- Error-Correcting Code (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
- FlashCard 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 148. You can also find information about AutoSupport features and functions on the NOW site at http://now.netapp.com/autosupport/.

Using SecureAdmin 9

About this chapter

This chapter describes how to configure the SecureAdmin feature on your NetApp system.

Topics in this chapter

This chapter discusses the following topics:

- "Learning about SecureAdmin" on page 164
- "Managing SSH for SecureAdmin" on page 169
- "Managing SSL for SecureAdmin" on page 179
- "General SecureAdmin administration" on page 184

Learning about SecureAdmin

How SecureAdmin improves security

SecureAdmin makes it very difficult for someone to intercept a NetApp system administrator's password over the network, because the password and all administrative communication are encrypted. SecureAdmin also provides a secure communication channel between a client and the NetApp system by using one or both of the following protocols:

- Secure Shell (SSH) protocol

  SSH provides a secure remote shell and interactive network session. SecureAdmin supports SSH 1.*x* clients and SSH 2.0 clients.
- Secure Sockets Layer (SSL) protocol

  SSL provides secure web access for FilerView and Data ONTAP APIs.

For detailed information

The following sections discuss how the SSH and SSL protocols improve security:

- "How SecureAdmin uses SSH" on page 165
- "How SecureAdmin uses SSL" on page 168

How SecureAdmin uses SSH

How SSH improves security

SSH improves security by providing a means for a NetApp system to authenticate the client and by generating a session key that encrypts data sent between the client and NetApp system. SSH performs public-key encryption using a host key and a server key.

Note

SecureAdmin supports password authentication and public-key-based authentication. SecureAdmin does not support the use of a .rhosts file or the use of a .rhosts file with RSA host authentication.

Supported encryption algorithms

SecureAdmin supports the following encryption algorithms:

- RSA/DSA 1024 bit
- 3DES in CBC mode
- HMAC-SHA1
- HMAC-MD5

Supported SSH clients

SecureAdmin supports the following SSH clients:

- OpenSSH client version 3.4 on UNIX platforms
- SSH Communications Security client version 3.2.0 on Windows and UNIX platforms
- Vandyke SecureCRT version 3.4.6 on Windows platforms
- PuTTY version 0.52 beta on Windows platforms

About keys used to improve security

SSH uses three keys to improve security:

- Host key
- Server key
- Session key

SSH uses the host key to encrypt and decrypt the session key. You determine the size of the host key, and Data ONTAP generates the host key when you configure SecureAdmin.

SSH uses the server key to encrypt and decrypt the session key. You determine the size of the server key when you configure SecureAdmin. If SSH is enabled, Data ONTAP generates the server key when any of the following events occur:

- You start SecureAdmin.
- An hour elapses.
- The NetApp system reboots.

SSH uses the session key to encrypt data sent between the client and NetApp system. The session key is created by the client. To use the session key, the client encrypts the session key using the host and server keys and sends the encrypted session key to the NetApp system, where it is decrypted using the host and server keys. After the session key is decrypted, the client and NetApp system can exchange encrypted data.

How SecureAdmin creates a secure session using SSH

The following table shows how SecureAdmin creates a secure session between the NetApp system and client.

Stage	What the client does	What the NetApp system does
1	The client sends an SSH request to the NetApp system.	The NetApp system receives the SSH request from the client.
2		The NetApp system sends the public portion of the host key, and the server key if SSH 1.*x* is used, to the client.
3	The client stores the public portion of the host key for future host authentication.	
4	The client generates a random session key.	

Stage	What the client does	What the NetApp system does
5	The client encrypts the session key by using the public portion of the host key, and the server key if SSH 1.*x* is used, and sends it to the NetApp system.	
6		The NetApp system decrypts the session key using the private portions of the host key, and the server key if SSH 1.*x* is used.
7	The NetApp system and the client exchange information that they encrypt and decrypt using the session key.	

How SecureAdmin uses SSL

How the SSL protocol improves security

Secure Sockets Layer (SSL) improves security by providing a digital certificate that authenticates NetApp systems and allows encrypted data to pass between the system and a browser. SSL is built into all major browsers; therefore, installing a digital certificate on the NetApp system enables the SSL capabilities between system and browser.

Unlike using FilerView to send the NetApp system password in plain text, using SSL and Secure FilerView improves security by encrypting the administrator's password and all administrative communication when you manage your system from a browser.

Managing SSH for SecureAdmin

Options that affect SSH operation

SecureAdmin uses the following options to enable secure sessions using SSH:

- `options ssh.passwd_auth.enable`—Controls password-based authentication.
- `options ssh.pubkey_auth.enable`—Controls public key authentication.
- `options ssh.access`—Controls access to a NetApp system.

Note

The default value for `ssh.passwd_auth.enable` and `ssh.pubkey_auth.enable` is On. The default value for `ssh.access` allows everyone to access the NetApp system.

Ways to manage SSH

You can manage the SSH portion of SecureAdmin in the following ways:

- "Setting up and starting SSH" on page 170
- "Reinitializing SSH" on page 172
- "Stopping or starting SSH service" on page 174
- "Setting up public key-based authentication" on page 175

Setting up and starting SSH

Guidelines for determining host and server key sizes

The setup process involves creating host and server keys. You can determine the size of the host and server keys by using the following guidelines:

- If you are using the SSH 1.*x* protocol, the size of the host and server keys can range from 384 bits to 2,048 bits.
- If you are using the SSH 2.0 protocol, the size of the host and server keys can range from 768 to 2,048 bits.
- As the size increases, the security increases; however, initiating a new SecureAdmin session takes longer and NetApp system performance might decrease.
- The size of the host key must differ from the size of the server key by at least 128 bits. It does not matter which key is larger.

Files where host keys are saved

If you are using the SSH 1.*x* protocol, the host key is stored in the /etc/sshd/ssh_host_key file.

If you are using the SSH 2.0 protocol, the RSA host key is stored in the /etc/sshd/ssh_host_rsa_key file and the DSA host key is stored in the /etc/sshd/ssh_host_dsa_key file.

Setting up and starting SSH

To set up and start SSH, complete the following steps.

Note

The setup procedure requires you to enter key sizes for the SSH 1.*x* and SSH 2.0 protocols, regardless of the protocol you use. For example, if you plan to use the SSH 2.0 protocol, you still must enter values for the SSH 1.*x* host key and server key sizes. NetApp recommends that you accept the default value for keys that you do not use.

Step	Action
1	Enter the following command: `secureadmin setup ssh`

Step	Action
2	When prompted, enter a size for the host key if you are using the SSH 1.*x* protocol. **Note** The default size for the host key is 768 bits.
3	When prompted, enter a size for the server key if you are using the SSH 1.*x* protocol. **Note** The default size for the server key is 512 bits.
4	When prompted, enter a size for the host keys if you are using the SSH 2.0 protocol. **Note** The default size for the host key is 768 bits.
5	When prompted, confirm the parameters that you specified. **Result:** SecureAdmin generates the host key in the background, and, after a minute or two, the `setup` program sends a `syslog` message announcing that SSH is set up.
6	After the `syslog` message is generated, activate the host and server keys by entering the following command: **`secureadmin enable ssh1|ssh2`** Use `ssh1` to enable SSH service for SSH 1.*x* clients or `ssh2` to enable SSH service for SSH 2.0 clients.

Reinitializing SSH

Changing key sizes by reinitializing SSH

Reinitializing SSH enables you to change the sizes of existing host and server keys. To reinitialize SSH, complete the following steps.

Step	Action
1	Cancel the existing host and server keys by stopping the SSH daemon with the following command: **`secureadmin disable ssh1|ssh2`** Use `ssh1` to disable SSH service for SSH 1.*x* clients or use `ssh2` to disable SSH service for SSH 2.0 clients.
2	Enter the following command: **`secureadmin setup -f ssh`**
3	When prompted, enter a size for the host key if you are using the SSH 1.*x* protocol.
4	When prompted, enter a size for the server key if you are using the SSH 1.*x* protocol.
5	When prompted, enter a size for the host key if you are using the SSH 2.0 protocol.
6	Activate the new host and server key sizes by entering the following command: **`secureadmin enable ssh1|ssh2`** Use `ssh1` to enable SSH service for SSH 1.*x* clients or use `ssh2` to enable SSH service for SSH 2.0 clients.

Result: Clients that have a copy of the old host key give the following warning after they receive a new key from the NetApp system:

```
@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@
@       WARNING: HOST IDENTIFICATION HAS CHANGED!         @
@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@
IT IS POSSIBLE THAT SOMEONE IS DOING SOMETHING NASTY!
Someone could be eavesdropping on you right now (man-in-the-middle
attack)!
It is also possible that the host key has just been changed.
Please contact your system administrator.
Add correct host key in /u/sisa/.ssh/known_hosts to get rid of this
message.
Agent forwarding is disabled to avoid attacks by corrupted servers.
Are you sure you want to continue connecting (yes/no)?
```

Stopping or starting SSH service

Disabling or enabling SSH

You disable or enable SSH to stop or start SSH service, respectively. To disable or enable SSH, complete the following step.

Step	Action
1	Enter the following command: **`secureadmin disable|enable ssh1|ssh2`** Use `disable` to stop SSH service or `enable` to restart SSH service. Use `ssh1` to support SSH 1.*x* clients or `ssh2` to support SSH 2.0 clients. **Example:** The following command enables SSH service for SSH 2.0 clients: `secureadmin enable ssh2`

Setting up public key-based authentication

RSA key pair requirement

Setting up key-based authentication requires an RSA key pair (a private and public key) in addition to the host and server keys. Public-key-based authentication differs between the two versions of SSH; SSH 1.*x* uses an RSA key pair and SSH 2.0 uses a DSA key pair in addition to an RSA key pair. For both versions of SSH, you must generate the key pairs and copy the public key to the NetApp system.

Generating a key pair for SSH 1.*x*

To generate an RSA key pair for SSH 1.*x* and then copy it to the NetApp system, complete the following steps.

Step	Action
1	Using your SSH 1.*x* client, generate an RSA key pair. **Result:** Your client generates the RSA key pair, a public key and a private key, and stores them on the client.
2	Copy the generated public key to the NetApp system root volume and append it to the /etc/sshd/*user_name*/.ssh/authorized_keys file.

Example: The following is an example of generating an RSA key pair with an OpenSSH UNIX client:

```
% ssh-keygen -t rsa1 -b 1024
Generating public/private rsa1 key pair.
Enter file in which to save the key (/u/john/.ssh/identity):
Enter passphrase (empty for no passphrase):
Enter same passphrase again:
Your identification has been saved in /u/john/.ssh/identity
Your public key has been saved in /u/john/.ssh/identity.pub
The key fingerprint is:
6a:c7:93:7c:b5:f4:12:87:81:56:5e:a2:62:40:07:8a john@unix1
```

In this example, the identity.pub file is the public-key file that you copy to the NetApp system root volume.

The following commands append the public key to the /etc/sshd/*user_name*/.ssh/authorized_keys file on NetApp system sys1:

```
% mount sys1:/ /mnt_sys1
% cat identity.pub >> /mnt_sys1/etc/sshd/john/.ssh/authorized_keys
```

Generating key pairs for SSH 2.0

Generating key pairs for SSH 2.0 requires generating an RSA key pair and a DSA key pair. If you use SSH 2.0 clients other than OpenSSH, you might have to edit the public key before you can use it. See "Public keys generated by SecureCRT and ssh.com clients" on page 177 for more information.

To generate key pairs for SSH 2.0, complete the following steps.

Step	Action
1	Using your SSH 2.0 client, generate an RSA key pair. **Result:** Your client generates the RSA key pair, a public key and a private key, and stores them on the client.
2	Using your SSH 2.0 client, generate a DSA key pair. **Result:** Your client generates the DSA key pair, a public key and a private key, and stores them on the client.
3	Copy the generated public key to the NetApp system default directory and append it to the /etc/sshd/*user_name*/.ssh/authorized_keys file.

Example: The following is an example of generating RSA and DSA key pairs with an OpenSSH UNIX client.

```
% ssh-keygen -t rsa -b 1024
Generating public/private rsa key pair.
Enter file in which to save the key (/u/john/.ssh/id_rsa):
Enter passphrase (empty for no passphrase):
Enter same passphrase again:
Your identification has been saved in /u/john/.ssh/id_rsa
Your public key has been saved in /u/john/.ssh/id_rsa.pub

% ssh-keygen -t dsa -b 1024
```

```
Generating public/private dsa key pair.
Enter file in which to save the key (/u/john/.ssh/id_dsa):
Enter passphrase (empty for no passphrase):
Enter same passphrase again:
Your identification has been saved in /u/john/.ssh/id_dsa
Your public key has been saved in /u/john/.ssh/id_dsa.pub
```

In this example, the id_rsa.pub and id_dsa.pub files are the public-key files that you copy to the NetApp system root volume.

The following commands append the public keys to the /etc/sshd/*user_name*/.ssh/authorized_keys file on NetApp system sys1:

```
% mount sys1:/ /mnt_sys1
% cat id_rsa.pub >> /mnt_sys1/etc/sshd/john/.ssh/authorized_keys
% cat id_dsa.pub >> /mnt_sys1/etc/sshd/john/.ssh/authorized_keys
```

Public keys generated by SecureCRT and ssh.com clients

SSH 2.0 public keys generated by SecureCRT and ssh.com clients contain comments and line breaks that make the public keys useless. You must make the following edits to the generated public keys before SecureAdmin can use them:

- Remove any text that is not part of the public key.
- Remove line breaks and spaces to make the public key one continuous string of characters.
- Before the first character of the public key, add ssh-rsa followed by a space.

Example: The following is an example of an SSH 2.0 public key generated by a SecureCRT client. The generated public key contains extra text and line breaks at the end of each line.

```
---- BEGIN SSH2 PUBLIC KEY ----
Subject: john
Comment: "john@johnnt"
AAAAB3NzaC1yc2EAAAADAQABAAAAgQDJhJ6nk+2hm5iZnx737ZqxFgksPl3+OY1cP8
0s
1amXuUrwBp3/MUODEP5E51lzqjO0w5kyJlvPjCiLg9UqS7JeY5yd/6xyGarsde26De
1E
rbVJ1uqnxyAOlV9A1hjBE8TbI+lyYBH+WezT0nySix6VBQTAWhv43r9lSudswYV80Q
==
---- END SSH2 PUBLIC KEY ----
```

The following is the public key after removing text that is not part of the public key, removing line breaks at the end of each line, and adding ssh-rsa at the beginning of the public key.

```
ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAAAgQDJhJ6nk+2hm5iZnx737ZqxFgksPl
3+OY1cP80s1amXuUrwBp3/MUODEP5E51lzqjO0w5kyJlvPjCiLg9UqS7JeY5yd/6xy
Garsde26De1ErbVJ1uqnxyAOlV9A1hjBE8TbI+lyYBH+WezT0nySix6VBQTAWhv43r
9lSudswYV80Q==
```

Managing SSL for SecureAdmin

Ways to manage SSL

You can manage the SSL portion of SecureAdmin in the following ways:

- Set up and start SSL
- Reinitialize SSL
- Disable and enable SSL

About the certificate used by SSL

SSL uses a certificate to provide a secure connection between the NetApp system and a Web browser. SecureAdmin allows two types of certificates:

- Self-signed certificate

  A certificate generated by Data ONTAP. Self-signed certificates can be used as is, but they are less secure than certificate-authority signed certificates, because the browser has no way of verifying the signer of the certificate. This means the system could be spoofed by an unauthorized server.

- Certificate-authority signed certificate

  A certificate-authority signed certificate is a self-signed certificate that is sent to a certificate authority to be signed. The advantage of a certificate-authority signed certificate is that it verifies to the browser that the system is the system to which the client intended to connect.

Setting up and starting SSL

To set up SSL, complete the following steps.

Step	Description
1	Enter the following command: `secureadmin setup ssl`
2	Enter information when Data ONTAP prompts you. **Note** To use the default settings, press Enter at each of the prompts. **Result:** Data ONTAP generates a self-signed certificate and keys, and places them in the /etc/keymgr directory.

<table>
<tr><th>Step</th><th colspan="2">Description</th></tr>
<tr><td>3</td><td colspan="2">Enter y if the information is correct or n if any of the information is incorrect.

Result: If you enter y, Data ONTAP generates two files and saves them in the /etc/keymgr directory:
◆ A self-signed certificate called secureadmin.der
◆ A certificate signing request called secureadmin.pem

If you enter n, Data ONTAP repeats the prompts.</td></tr>
<tr><td rowspan="3">4</td><td>If...</td><td>Then...</td></tr>
<tr><td>You want to use a self-signed certificate</td><td>You have successfully installed your certificate. Perform the steps described in “Testing your certificate” on page 181 to test your certificate.</td></tr>
<tr><td>You want to use a certificate-authority-signed certificate</td><td>Perform the steps in “Installing a certificate-authority-signed certificate” on page 181 to install a certificate-authority-signed certificate.

Note
You can use the self-signed-certificate Data ONTAP has already installed until you can install your certificate-authority-signed certificate.</td></tr>
</table>

Installing a certificate-authority-signed certificate

To install a certificate-authority-signed certificate, complete the following steps.

Step	Action
1	Send the certificate signing request, `secureadmin.pem`, to the certificate authority. This file is found in the `/etc/keymgr/cert` directory on the NetApp system. **Note** This process might take a few days.
2	Back up the `secureadmin.pem` file by making a copy.
3	When the certificate authority returns the signed certificate, copy the signed certificate into a temporary location on the NetApp system.
4	Install the certificate by entering the following command: **`secureadmin addcert ssl`** ***`directory_path`*** *directory_path* is the full path to the certificate. **Example:** The following command installs a certificate called secureadmin.pem, currently located in the tempdir directory, into the /etc/keymgr directory: `secureadmin addcert ssl /etc/tempdir/secureadmin.pem`
5	Disable SSL by entering the following command: **`secureadmin disable ssl`**
6	Enable SSL by entering the following command: **`secureadmin enable ssl`**
7	Perform the steps described in "Testing your certificate" on page 181 to test your certificate.

Testing your certificate

To verify that your certificate is installed correctly, complete the following steps.

Note

These steps can verify either a self-signed certificate or a certificate-authority-signed certificate.

Step	Action
1	Start your Web browser.
2	Enter the following URL: `https://`*`systemname`*`/na_admin` *systemname* is the name of your NetApp system.
3	Click Secure FilerView. Secure FilerView starts up in a new browser window.
4	Check your browser to verify that you have made a secure connection. **Note** Most browsers show a small padlock icon in their status bar when they have successfully made a secure connection to the server. If the padlock icon is *not* displayed, *you do not have a secure connection*, even if you clicked Secure FilerView.

Reason for reinitializing SSL

You should reinitialize SSL if you change the domain name of the NetApp system, because the domain name recorded in the certificate becomes obsolete. The result is that the connection is still encrypted, but the NetApp system is not authenticated. The next time you connect to the system, the browser issues a warning that the domain name of the system does not match the record on the certificate.

Note

Changing the domain name for a NetApp system that is using SSL can cost time and money because you must have the new certificate signed by a certificate authority.

Reinitializing SSL

To reinitialize SSL, complete the following steps.

Step	Description
1	Disable SecureAdmin by entering the following command: **`secureadmin disable ssl`**
2	Run setup to reinitialize SSL by following the steps in “Setting up and starting SSL” on page 179.

Disabling or enabling SSL

To disable or enable SSL, complete the following step.

Note

When you disable SSL, you disallow all administrative requests over HTTPS. Enabling SSL allows administrative requests over HTTPS to succeed.

Step	Action
1	Enter the following command: **`secureadmin disable|enable ssl`** Use `disable` to disable SSL or use `enable` to restart SSL.

General SecureAdmin administration

Enabling or disabling both protocols used by SecureAdmin

To enable or disable SecureAdmin for both SSH and SSL, complete the following step.

Step	Action
1	Enter the following command: **`secureadmin enable | disable all`** Use `enable all` to start SSH and SSL or use `disable all` to stop SSH and SSL.

Determining whether SecureAdmin is running

To determine whether SecureAdmin is running, and, therefore, whether administrative transactions between the NetApp system and a client are being encrypted, complete the following step.

Step	Action

1	Enter the following command:
	`secureadmin status`
	Result: Information similar to the following is displayed:
	`ssh2 - active`
	`ssh1 - inactive`
	`ssl - inactive`

Cluster Management 10

About this chapter

This chapter discusses managing a NetApp system cluster. For more information about configuring a new cluster or changing the configuration of your cluster, see the *Cluster Guide* for your hardware platform.

Topics in this chapter

This chapter discusses the following topics:

- "About NetApp clusters" on page 186
- "How a clustered system works" on page 190
- "Cluster operation modes" on page 192
- "Understanding interface configurations for takeover" on page 200
- "Enabling cluster failover in a DAFS configuration" on page 202
- "Managing a cluster" on page 209
- "Checking cluster configuration" on page 234
- "Cluster LCD messages" on page 235

About NetApp clusters

What a NetApp cluster is

A NetApp cluster is two NetApp filer head units (or nodes) that are connected with an interconnect adapter and cables. The two nodes are configured so they can share access to a set of disks, subnets, and tape drives.

While most NetApp filers have external head units, such as the F800 series and the FAS900 series, there are some appliances, such as the FAS200 series that have internal head units. The FAS270c, for example, is a cluster consisting of two internal head units that reside in and share the same disk shelf.

Benefits of NetApp clusters

Configuring NetApp filers in a cluster provides the following benefits:

- Fault tolerance

  When one node fails, a takeover occurs, and the partner node continues to serve the failed node's data.

- Non-disruptive software upgrades

  When you initiate a takeover, the partner node continues to serve data for the failed node's data while you upgrade the node you halted.

  For more information about non-disruptive upgrades, see the *Upgrade Guide*.

- Non-disruptive appliance and disk maintenance

  When you halt one node and prevent the takeover, the partner node continues to serve data for the failed node's data while you replace or repair hardware in the node you halted.

Note

NetApp clusters do not protect against NIC or network failures. To provide protection against these types of failures, use virtual network interfaces, or VIFs. For more information about VIFs, see the *Network Administration Guide*.

For more information

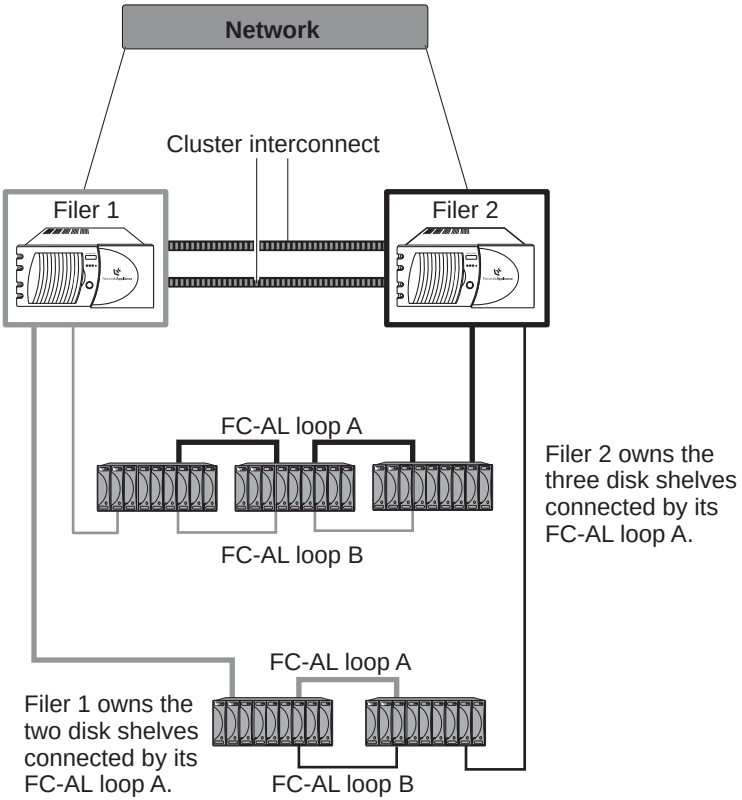
This guide provides information about managing a cluster. For more information about clusters, see the following documentation:

- For information about installing and cabling a cluster, see the cluster guide for the appliance model you purchased, for example, the *F800 and FAS900 Series Appliance Cluster Guide*.

- For information about setting up and configuring a cluster, see the *Software Setup Guide*.

Cluster diagram

The following figure diagrams a common configuration of the cluster hardware components.



Note

The above diagram assumes that each filer owns the disks attached to its own shelves. If you are using a FAS270c (which contains two filer heads) or SnapMover, then you have software disk ownership. In these cases, you need to establish disk ownership explicitly.

For more information about software disk ownership, see the Disk chapter of the Storage Management Guide. For more information about SnapMover, see the Multistore Guide.

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 or more Fibre Channel-Arbitrated Loops (FC-AL loops), in which
  - Each filer manages its own disks.
  - Each filer in takeover mode manages its partner's disks.

Note

For the FAS 270c and systems using the SnapMover® software, software disk ownership is in use, so disk ownership is established explicitly rather than by which shelf the disk is attached to.

For more information about software disk ownership, see the chapter on disk management in the *Storage Management Guide*. For more information about SnapMover, see the *Multistore Management Guide*.

- They own their spare disks and do not share them with the other filer.
- They each have two mailbox disks located on the root volume (four if the root volume is mirrored using SyncMirror). The mailbox disks are used to
  - Maintain consistency between the pair
  - Continually check whether the other filer is running or if it has performed a takeover
  - Store cluster information that is not specific to any particular node
- 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 must be the same type and model.
- Filer capacity

  The number of disks must not exceed the maximum cluster capacity. To determine the maximum capacity, see the System Configuration Guide on the NOW site at Service & Support > Access Product Documentation > System Configuration Guide.

- Disks and disk shelf compatibility

  The disk shelves must have two paths so that each filer can attach to the shelf.

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 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 MultiStore, 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

When takeovers occur

Takeovers are initiated in a clustered pair when one of the following conditions occur:

- A filer is in a cluster that is configured for immediate takeover on panic and undergoes a system failure.
- A filer undergoes a system failure and cannot reboot.
- There is a mismatch between the disks that one filer can see and the disks that the other filer can see, and the `cf.takeover.on_disk_shelf_miscompare` option is set to On.

  In this case, a negotiated failover is initiated. The filer that cannot see its own disks requests that its partner filer initiate a takeover. When the impaired filer shuts down, or after a timeout of approximately three minutes, the partner filer performs the takeover.
- A filer is unable to send heartbeat messages to its partner.
- You halt one of the clustered filers without using the `-f` flag.
- You initiate a takeover manually.

What happens during takeover

When a takeover occurs, the unimpaired 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.

Note

When a takeover occurs, existing CIFS sessions are terminated. A graceful shutdown of the CIFS sessions is not possible, and some data loss could occur for CIFS users.

What happens after takeover

After 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 in effect when there is no takeover or giveback in progress.
- Takeover mode is in effect when one filer has taken over its partner.
- Giveback mode is in effect as the cluster 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 193
- "Takeover operation mode" on page 195
- "Giveback operation mode" on page 198

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 background process 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 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's configuration files to duplicate the failed filer's 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:

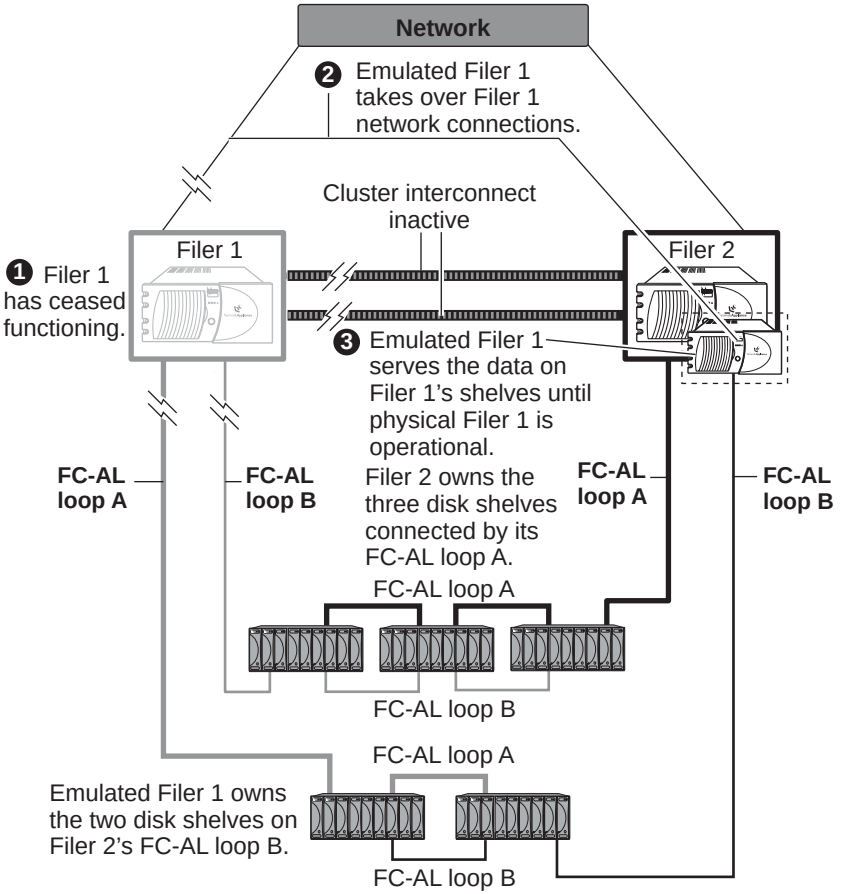
```
filer1> cf takeover
cf: takeover initiated by operator
filer1> Mon Aug 23 23:33:57 GMT [filer1:
cf.misc.operatorTakeover:warning]: Cluster monitor: takeover
initiated by operator
Mon Aug 23 23:33:57 GMT [filer1:
cf.fsm.nfo.acceptTakeoverReq:warning]: Negotiated failover:
accepting takeover request by partner, reason: operator initiated
cf takeover. Asking partner to shutdown gracefully; will takeover
in at most 180 seconds.
Mon Aug 23 23:34:01 GMT [filer1:
cf.fsm.nfo.partnerShutdown:warning]: Negotiated failover: partner
has shutdown
Mon Aug 23 23:34:01 GMT [filer1: cf.fsm.stateTransit:warning]:
Cluster monitor: UP --> TAKEOVER
Mon Aug 23 23:34:01 GMT [filer1: cf.fm.takeoverStarted:warning]:
Cluster monitor: takeover started
Replaying takeover WAFL log
Mon Aug 23 23:34:02 GMT [filer2/filer1: wafl.replay.done:info]:
WAFL log replay completed, 0 seconds
Mon Aug 23 23:34:03 GMT [filer1 (takeover):
cf.fm.takeoverComplete:warning]: Cluster monitor: takeover
completed
```

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.

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 starts 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)> cf giveback
filer1(takeover)> Mon Aug 23 23:37:54 GMT [filer1 (takeover):
cf.misc.operatorGiveback:info]: Cluster monitor: giveback initiated
by operator
Mon Aug 23 23:37:54 GMT [filer1: cf.fm.givebackStarted:warning]:
Cluster monitor: giveback started
Mon Aug 23 23:37:55 GMT [filer1: timed_restarter:error]: cf.timed:
couldn't start cluster time daemon; time will not be synchronized
between nodes in cluster
Mon Aug 23 23:37:55 GMT [filer1: cf.fm.givebackComplete:warning]:
Cluster monitor: giveback completed
Mon Aug 23 23:37:55 GMT [filer1: cf.fsm.stateTransit:warning]:
Cluster monitor: TAKEOVER --> UP
Mon Aug 23 23:37:55 GMT [filer1:
cf.fsm.takeoverByPartnerDisabled:notice]: Cluster monitor: takeover
of filer1 by filer2 disabled (unsynchronized log)
Mon Aug 23 23:38:03 GMT [filer1:
cf.fsm.partnerNotResponding:notice]: Cluster monitor: partner not
responding
Mon Aug 23 23:38:07 GMT [filer1: cf.fm.timeMasterStatus:info]:
Acting as cluster time master
Mon Aug 23 23:38:07 GMT [filer1: cf.fsm.partnerOk:notice]: Cluster
monitor: partner ok
Mon Aug 23 23:38:07 GMT [filer1:
cf.fsm.takeoverOfPartnerEnabled:notice]: Cluster monitor: takeover
of filer2 enabled
Mon Aug 23 23:38:09 GMT [filer1:
cf.fsm.takeoverByPartnerEnabled:notice]: Cluster monitor: takeover
of filer1 by filer2 enabled
Mon Aug 23 23:38:11 GMT [filer1:
cf.fsm.takeoverOfPartnerDisabled:notice]: Cluster monitor: takeover
of filer2 disabled (unsynchronized log)
Mon Aug 23 23:38:12 GMT [filer1:
cf.fsm.takeoverOfPartnerEnabled:notice]: Cluster monitor: takeover
of filer2 enabled
```

Understanding interface configurations for takeover

Requirements of network interfaces in a cluster

In clusters, network interfaces must 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 supports only the identity of the local filer.

Cluster network interface configurations

There are two kinds of interface configurations for clusters:

- Shared interfaces

  In this configuration, each interface can provide a network connection for both clustered filers. When a failover occurs, the interface on the takeover node takes over the network connection for the failed node.
- Dedicated and standby interfaces

  In this configuration, each filer has two interfaces, one dedicated and one standby. The dedicated interface provides the network connect for each filer during normal mode. When a failover occurs, the standby interface on the takeover node takes over the network connection for the failed node.

Note

Most cluster interfaces are configured as shared interfaces, because they do not require an extra network card. However, some Asynchronous Transfer Mode (ATM) configurations require interfaces to be configured as dedicated and standby.

Interface roles in normal and takeover modes

A cluster can be in one of two modes:

- 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
Shared	Supports the identity of the local filer	Supports the identity of both the local filer and the failed filer
Dedicated	Supports the identity of the local filer	Supports the identity of the local filer
Standby	Idle	Supports the identity of the failed filer

About configuring interfaces

For information about configuring shared, dedicated, and standby interfaces for a NetApp cluster, see the *NetApp Hardware Cluster Guide*.

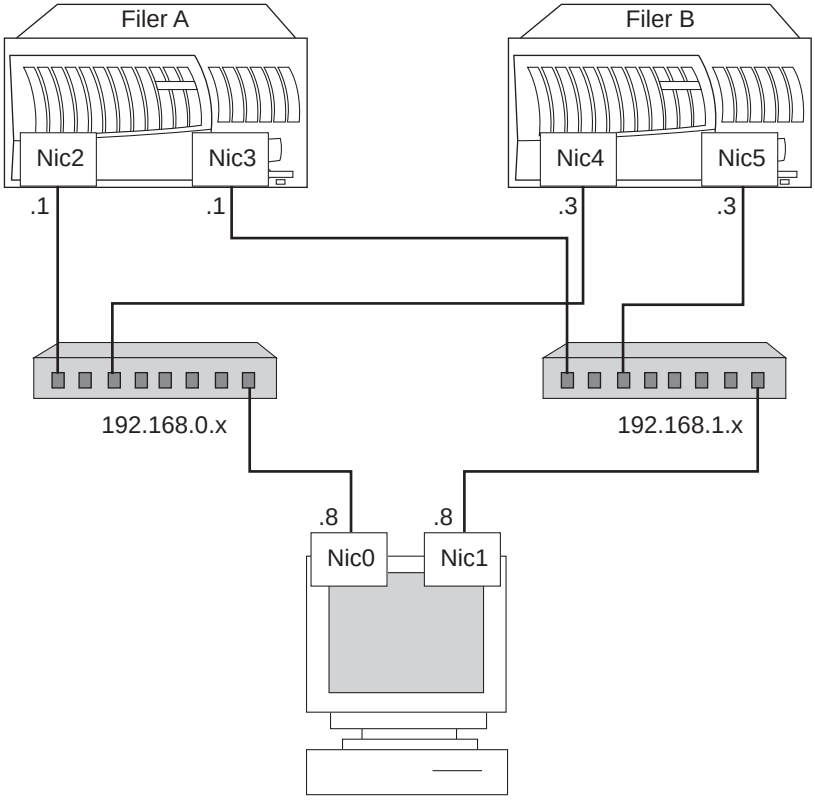
Enabling cluster failover in a DAFS configuration

About the failover path

When a filer in a cluster fails, the DAFS Database Accelerator client must find the failover paths over the virtual interface (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 204
- "Enabling failover paths on the filer" on page 207

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 202.

```
# 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 filerA has the following host names:

- `filerA-viip0`—The local NIC host name

- `filerB-failover-viip0`—A failover NIC host name that indicates that filerA is the takeover filer when filerB fails

How the client identifies the failover path

If filerA 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.
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`

Step	Action
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 and Protocols 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 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 210
- "Enabling negotiated failover between filers" on page 218
- "Managing a cluster in takeover mode" on page 220
- "Managing an emulated filer" on page 222
- "Backing up and restoring from the emulated filer" on page 227
- "Performing a giveback" on page 228
- "Enabling automatic giveback" on page 231
- "Configuring automatic giveback" on page 232

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.
`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.
`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. takeover 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.

Message	Meaning
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 interconnect links 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 `aggr 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 **`aggr status -r`**

Enabling and disabling takeover

You might want to disable takeover if you are doing maintenance that typically causes a takeover. You can 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 takes over immediately (within 60 seconds) without waiting until the core dump is completed.

Note

Until the core dump to the spare disk is complete, the spare disk is not 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 steps.

Step	Action
1	Ensure that you have enabled cluster takeover with the `cf enable` command.
2	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 apply. **Note** By default this option is turned off, except when the iSCSI and FCP licenses have been previously enabled or when the nodes are FAS270 systems. 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`

Unlicensing the cluster feature

You can unlicense the cluster feature only in normal mode with takeover disabled.

Note

Unlicensing the cluster feature will not disable the cluster. To disable the cluster, see the Cluster Guide for your filer model.

If you unlicense the cluster feature, you must reboot both nodes before relicensing clustering.

To unlicense the cluster feature, complete the following steps on both nodes.

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 and describes the commands you can use when initiating a takeover.

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 (if the `cf.takeover.on_panic` option is set to `false`). The takeover starts either after the partner halts successfully or after a timeout.
`cf takeover -f`	Initiates an immediate takeover of the partner of the local filer regardless of whether the other filer is dumping its core. The partner node is not allowed to halt gracefully.
`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 risk losing uncommitted data when you use this command.

Command	Description
`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 is 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 Online Backup and Recovery 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 or disable negotiated failover on a cluster, complete the following step.

Step	Action
1	Enter the following command: **`options cf.takeover.on_disk_shelf_miscompare {on|off}`** Use `on` to turn negotiated failover on, and use `off` 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 15.

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.

Note

If your cluster is failing over only because one of the filers is too busy to respond to its partner, increase the value of the `cf.takeover.detection.seconds` option on the partner.

Caution

If the specified time is less than 15 seconds, unnecessary takeovers can occur. Use caution when assigning a takeover time of less than 15 seconds.

Determining negotiated failover status

To determine negotiated failover status on a cluster, complete the following step.

Step	Action
1	Enter the following command: **`cf status`** **Results:** The `cf 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

Differences between takeover mode and 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. 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.

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.
`aggr status`	When it displays hardware information, the `aggr 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 aggr 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 are 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 Tape Backup and Recovery 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
- Volume creation (traditional volume or FlexVol)
- Aggregate creation
- Disk ownership assignment
- 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 229.

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 risk losing 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.

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. To determine whether your NetApp system boots from CompactFlash, see your Hardware and Service guide.

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. To determine whether your NetApp system boots from CompactFlash, see your Hardware and Service guide.

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`** This option is Off by default, except for platforms booted from CompactFlash, where it is On by default. **Note** To determine whether your NetApp system boots from CompactFlash, see your Hardware and Service guide.

Setting giveback 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 giveback 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 through 999. The default is 5 minutes.

Setting giveback 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 giveback to terminate long-running processes

To set giveback 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 utility

NetApp provides a Cluster Configuration Checker utility that you can download and run to protect against the most common configuration causes of failover problems between clustered filers.

Where you can run the utility

You can run the utility, `cf-config-check.cgi`, as a command from a UNIX shell, or you can install the Common Gateway Interface (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, `rsh` access to both clustered filers is required.

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.
- Differences in the cfmode settings between the clustered filers.

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.

Cluster LCD messages

About this section

This section lists some of the messages you could see on your LCD display (on the front of the filer) regarding your cluster, and steps you need to take.

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.

Message	Explanation	What you do
`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.
`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 or aggregates 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 or aggregates 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.

System Information and Performance 11

About this chapter

This chapter describes commands that give you information about the NetApp system and how it is performing.

Note

Some options for different commands can gather the same information. For example, the `aggr status -r` command and `sysconfig -r` command gather the same RAID information and present it in the same format.

For more information about any of the commands described in this chapter, see the corresponding man page.

Topics in this chapter

This chapter discusses the following topics:

- "NetApp system configuration information" on page 240
- "Storage information" on page 243
- "Aggregate information" on page 245
- "Volume information" on page 247
- "File statistics for volumes" on page 249
- "Environmental status information" on page 256
- "Fibre Channel information" on page 259
- "Getting NetApp system information using the stats command" on page 260
- "Improving NetApp system performance" on page 270

NetApp system configuration information

Version information

The `version` command displays the version of Data ONTAP currently running on a NetApp system.

Hardware configuration

The `sysconfig` command displays information about the NetApp system'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 NetApp system.

RAID and checksum information

The `sysconfig -r` command displays the status of plexes and aggregates; 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 NetApp system
- 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 an aggregate

You can obtain the information displayed by `sysconfig -r` from SNMP, using the NetApp 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.

Traditional volume and aggregate information

The `sysconfig -V` command displays RAID group and disk information about each traditional volume and aggregate.

Tape library information

The `sysconfig -m` command displays tape library information. Before you use this option, ensure that the NetApp system was booted with the autoload setting of the tape library off.

Overall NetApp system 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 NetApp system. 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 NetApp, provide the information displayed by `sysconfig -v`. This information is useful for diagnosing system problems.

Comprehensive NetApp system information

The `sysconfig -A` command displays NetApp system 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 aggregates.

Getting system information using the stats command

You can also get system information, either interactively or with a script, using the `stats` command. For more information about the `stats` command, see "Getting NetApp system information using the stats command" on page 260.

Storage information

Command for displaying storage information

You use the `storage show` command to display information about storage components for a NetApp system, including information about all disks and adapters.

For more information on the `storage` command, see the na_storage(1) man page and the chapter on disks in the *Storage Management Guide*.

Adapter information

The `storage show adapter` command displays information about all storage adapters currently in your NetApp system. 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 NetApp system. 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 NetApp system. 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 NetApp system. The `storage show mc` *name* command displays information for only the medium changer specified.

Switch port information

The `storage show port` command displays information about all ports on all switches attached to your NetApp system. 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 NetApp system. 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 NetApp system. 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.

Aggregate information

Command for displaying aggregate information

You use the `aggr status` command to display information about aggregate configurations. The types of information displayed depend on the command options. When you specify an aggregate, such as `aggr status aggr0`, the information for that aggregate is displayed; when you do not specify an aggregate, the status of all aggregates and traditional volumes in the NetApp system is displayed.

The aggregate status commands work for aggregates that were created explicitly, as well as for the aggregates created automatically when traditional volumes were created. Because a traditional volumes is tightly coupled with its containing aggregate, the aggregate status commands return information for both aggregates and traditional volumes. In both cases, it is the aggregate information that is returned.

For more information about the `aggr` command, see the na_aggr(1) man page.

Aggregate state information

With no options, the `aggr status` command displays a concise synopsis of aggregate states. This includes the name of the aggregate, whether it is an aggregate or traditional volume, whether it is online, offline, or restricted, whether its RAID type is RAID-4 or RAID-DP, and other states, such as partial or degraded, and any options that are enabled, either by default or through the `aggr options` or `vol options` command.

Disk information

The `aggr 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 `aggr status -r` command displays RAID, plex, and checksum information for an aggregate. The display is the same as the `sysconfig -r` display.

RAID information for each group

The `aggr status -v` command displays information about each RAID group within an aggregate or traditional volume, and the settings of the aggregate options.

Getting aggregate information using the stats command

You can also get aggregate information, either interactively or with a script, using the `stats` command. For more information about the `stats` command, see "Getting NetApp system information using the stats command" on page 260.

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 NetApp system 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 a FlexVol or a traditional volume; whether it is online, offline, or restricted; other status, such as partial and degraded; and what, if any, options are enabled for the volume or its containing aggregate (either by default or through the `aggr options` or `vol options` command). The `vol` command also displays RAID information for the volume's containing aggregate.

Disk information

The `vol status -d` command displays information about the volume's containing aggregate's disks. The information displayed is the same as for the `sysconfig -d` command.

RAID and checksum information

The `vol status -r` command displays RAID, plex, and checksum information for the volume's containing aggregate. The information displayed is the same as for the `sysconfig -r` command.

Volume options and RAID information

The `vol status -v` command displays the state of all per-volume options and information about each plex and RAID group within the volume's containing aggregate.

Volume language

The `vol status -l` command displays the language used by each volume.

Getting volume information using the stats command

You can also get volume information, either interactively or with a script, using the `stats` command. For more information about the `stats` command, see “Getting NetApp system information using the stats command” on page 260.

File statistics for volumes

About file statistics

The `filestats` command displays a summary of file statistics within a volume on a NetApp system by reading file information from a snapshot that you specify. 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.

Detailed information

This section contains the following topics:

- "About the filestats command" on page 250
- "Options for the filestats command" on page 253

About the filestats command

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] [async] [ages` *`ages`*`] [timetype {a,m,c,cr}] [sizes` *`sizes`*`] snapshot` *`snapshot_name`* `[volume` *`volume_name`*`] [style` *`style`*`] [file` *`output_file`*`]`

The `snapshot` argument is required; if the volume name is not specified, vol0 is assumed.

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 253.

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 NetApp system.

Example with ages option specified

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

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

- Use the daily age breakdown displayed in the Cumulative Total KB column of the Age output to determine the average change in data per day.
- Divide the amount of disk space you want to reserve for snapshots by the daily change average. For example, if you find that the average daily change rate is 3 GB and you have a 200-GB volume, 40 GB (or 20 percent) of which you want to reserve for snapshots, divide 40 by 3 to determine the number of daily snapshots you can have before exceeding your space limit. In this example, 13 daily snapshots is your limit.

Example of determining volume capacity

You can also use the `filestats` command to determine when the 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 the 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 Online Backup and Recovery 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`
- `file`
- `async`

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 user 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 251.

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 251.

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.

About the file option

The `file` option prints the results of the `filestats` command to the specified output file, rather than the console. The output file is created in the /etc/log directory.

About the async option

The `async` option causes the `filestats` command to run independently of the console. This option is designed for use with the `file` option.

Note

Running more than one asynchronous `filestats` command simultaneously can adversely affect system performance.

Environmental status information

Command for displaying environmental status

You use the `environment` command to display information about the NetApp system 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 NetApp system environment:

- Shelf status and temperature
- NetApp system component information
- NetApp system temperature
- Devices attached to the NetApp system

You can query information about the following items:

- Disk shelves
- The NetApp system power supply
- The NetApp system 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 NetApp system is crossed.
- When you enter the command from the command line.

  You run this command manually to monitor the NetApp system subsystems, especially when you suspect a problem and when reporting abnormal conditions to NetApp.

All environment information

The `environment status` command with no options displays all NetApp system 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.

Detailed chassis environment status

The `environment chassis list-sensors` command displays detailed information from all chassis sensors.

Command for monitoring UPS devices

You can add UPS devices to be monitored, enable or disable monitoring of UPS devices, or display the status of UPS devices. The `ups status` command displays the status of all UPS devices.

Specifying a UPS device to be monitored

To specifying a UPS device to be monitored by the NetApp system'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 cyclic redundancy check (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.

Getting Fibre Channel information using the stats command

You can also get Fiber Channel information, either interactively or with a script, using the fcp object for the `stats` command. For more information about the `stats` command, see “Getting NetApp system information using the stats command” on page 260.

Getting NetApp system information using the stats command

About the stats command

The `stats` command provides access, through the command line or scripts, to a set of pre-defined data collection tools in Data ONTAP called counters. These counters provide you with information about your NetApp system, either instantaneously or over a period of time.

About objects, instances, and counters

Stats counters are grouped by what object they provide data for. Stats objects can be physical entities such as `system`, `processor` or `disk`; logical entities such as `volume` or `aggregate`; protocols such as `iscsi` or `fcp`, or other modules on your NetApp system. To see a complete list of the stat objects, you can use the `stats list objects` command.

Each object can have zero or more instances on your NetApp system, depending on your system configuration. Each instance of an object has its own name. For example, for a system with two processors, the instance names are `processor0` and `processor1`.

Counters have an associated privilege mode; if you are not currently running with sufficient privilege for a particular counter, it is not recognized as a valid counter.

How you use the stats command

When you use the stats command to get information about your NetApp system, you need to make the following decisions:

- What counters do you want to collect information from, on what object instances?

  For more information, see "About objects, instances, and counters" on page 260.
- Do you want to specify the counters on the command line or do you want to use a predetermined set of counters called a preset file?

  NetApp provides some preset files. You can also create your own.

  For more information, see "About the stats preset files" on page 269.
- How do you want the information to be returned and formatted?

  You can control where the information is returned (to the console or to a file) and how it is formatted.

  For more information, see "Changing the output of a stats command" on page 267.

- How do you want to invoke the stats command?

  You can invoke the stats command using the following methods:

  - A single invocation

    This method retrieves information from the specified counters once and stops.

    For more information, see "Using the stats command interactively in singleton mode" on page 265.
  - A periodic invocation

    For this method, information is retrieved from the specified counters repeatedly, at a time interval of your choice. You can specify a number of iterations to be performed, or the stats command can run until you stop it explicitly.

    For more information, see "Using the stats command interactively in repeat mode" on page 266.
  - As a background process

    This method enables you to initiate a stats command process that runs in the background until you terminate it explicitly, when the average values for the specified counters are returned.

    For more information, see "Using the stats command in background mode" on page 267.

Viewing the list of available counters

To display the list of counters for a particular object on the command line, complete the following step.

Step	Action
1	Enter the following command: **`stats list counters`** ***`object_name`*** *object_name* is the name of the object you want to list the available counters for.

Example:

```
cli> stats list counters system
Counters for object name: system
        nfs_ops
        cifs_ops
        http_ops
        dafs_ops
```

```
fcp_ops
iscsi_ops
net_data_recv
net_data_sent
disk_data_read
disk_data_written
cpu_busy
avg_processor_busy
total_processor_busy
num_processors
```

Getting more details about a counter

If you are planning to process the information you get from a `stats` command, you might want to understand how a counter presents its information. Or, you might want to get more details about the information that counter collects. To get more details about a specific counter, complete the following step.

Step	Action
1	Enter the following command: **`stats explain counters`** ***`object_name`*** **`[`** ***`counter_name`*** **`]`** *object_name* is the name of the object the counter is associated with. *counter_name* is the name of the counter you want more details about. If *counter_name* is omitted, information about all counters on the specified object is returned. **Result:** The following fields are returned for every specified counter: ◆ Name ◆ Description ◆ Properties The Properties field describes the type of information that is returned by this counter. Properties include the following types: ❖ `percent` for values that are a percentage value, such as `cpu_busy` ❖ `rate` for values that describe a value per time, such as `disk_data_read` ❖ `average` for values that return an average, such as `write_latency` ❖ `raw` for simple values that have no type, such as `num_processors` ◆ Unit The Unit field describes how value returned by this counter can be interpreted. The Unit field can be in one of the following groups of values: ❖ `percent` for counters with a Properties of `percent` ❖ The unit per time period for counters with a Properties of `rate`, such as `kb_per_sec` or `per_sec`. ❖ The time unit for counters that return timing values, such as `write_latency`

Example:

```
cli> stats explain counters system cpu_busy
Counters for object name: system
Name: cpu_busy
Description: Percentage of time one or more processors is busy in
the system
Properties: percent
Unit: percent
```

Using the stats command interactively in singleton mode

When you want to see a set of information about the system's current state at the command line, you can use the `stats` command in singleton mode. To do so, complete the following step.

Step	Action
1	Enter the following command: **`stats show object_def [object_def...]`** *object_def* is one of the following values: ◆ An object name (`object_name`) **Example:** `stats show system` This returns statistics from all counters provided for all instances of the specified object. ◆ The name of a specific instance (`object_name:instance_name`) **Example:** `stats show processor:processor0` This returns statistics from all counters provided for the specified instance of the specified object. ◆ The name of a specific counter (`object_name:instance_name:counter_name`) **Example:** `stats show system:*:net_data_recv` **Note** To see the statistic for all instances of the object, use an asterisk (*) for the instance name. To specify an instance name that includes spaces, enclose the name in double quotes (`"name with spaces"`). To specify an instance name that contains a colon (:), repeat the colon (`disk:20::00::00::20::37::de::4a::8e`). ◆ An asterisk (*) This returns statistics for all instances of all objects.

Example: The following command shows all current statistics for a volume named myvol.

```
cli> stats show volume:myvol
volume:myvol:total_ops:132/s
```

```
volume:myvol:avg_latency:13ms
volume:myvol:read_ops:5/s
volume:myvol:read_data:1923b/s
volume:myvol:read_latency:23ms
volume:myvol:write_ops:186/s
volume:myvol:write_data:1876b/s
volume:myvol:write_latency:6ms
volume:myvol:other_ops:0/s
volume:myvol:other_latency:0ms
```

Using the stats command interactively in repeat mode

If you want to see a statistic every few seconds, you can invoke the `stats` command in repeat mode. To do so, complete the following step:

Step	Action
1	Enter the following command: **`stats show [ -n num ] [ -i interval ] object_def [object_def...]`** *num* specifies the number of times you want the command to be run. If this parameter is omitted, the command is repeated until you issue a break. *interval* specified the interval between the iterations of the `stats` command. The default value is one second. *object_def* is specified as described in "Using the stats command interactively in singleton mode" on page 265.

Example: You could use the following command to display how your processor usage is changing over time:

```
stats show -i 1 processor:*:processor_busy
Instance processor_busy
                      %
processor0            32
processor1             1
processor0            68
processor1            10
processor0            54
processor1            29
processor0            51
...
```

Using the stats command in background mode

You can collect information from a specified set of counters over time in the background using the `stats start` and `stats stop` commands. The information collected is averaged over the period and displayed when the `stats stop` command is issued. You can initiate multiple stats commands in background mode, giving each of them a name so you can control them individually.

To collect system information over time, complete the following steps:

Step	Action
1	To start collecting the information, enter the following command: **`stats start [ -I identifier ] object_def [ object_def...]`** *identifier* names this instance of the `stats` command so you can refer to it later to show results. If you are running only one background `stats` command, you can omit this parameter. *object_def* is specified as described in "Using the stats command interactively in singleton mode" on page 265.
2	If you want to display interim results without stopping the background `stats` command, enter the following command: **`stats show [ -I identifier ]`** *identifier* names the instance of the `stats` command you want to display interim results for. If you are running only one background `stats` command, you can omit this parameter.
3	To stop data collection and output the final results, enter the following command: **`stats stop [ -I identifier ]`** *identifier* names the instance of the `stats` command you want to stop and display results for. If you are running only one background `stats` command, you can omit this parameter.

Changing the output of a stats command

Data ONTAP enables you to control the format and destination of the output of the `stats` command. This could be useful if you are processing the information with another tool or script, or if you want to store the output in a file so you can process it at a later time.

There are three ways you can affect the format and destination of the output of a `stats` command:

- You can send the output to a file rather than the console.
- You can determine whether the output appears in columns or rows.
- You can specify a delimiter.

Sending stats output to a file: To send `stats` output to a file, you add `-o` *filename* to your `stats show` or `stats stop` command line, where *filename* is the pathname to the file you want to receive the `stats` output. The file does not need to exist, although any directory in the path must already exist.

Using rows and columns: To determine whether the output is formatted in rows or columns, you add the `-r` or `-c` option to your `stats show` or `stats stop` command line.

The `-r` option formats the output in rows, and is the default if the `-i` option is not specified.

Example output in rows:

```
cli> stats show qtree:*:nfs_ops
qtree:vol1/proj1:nfs_ops:186/s
qtree:vol3/proj2:nfs_ops:208/s
```

The `-c` option formats the output in columns, and is the default only if the `-i` option is specified.

Example output in columns:

```
cli> stats show -c qtree:*:nfs_ops
Instance nfs_ops
              /s
vol1/proj1     143
vol3/proj2     408
```

Note

The "/s" line shows the unit for the applicable column. In this example, there is one column, and it is number of operations per second.

Specifying a delimiter: You might need to specify a delimiter so that your output can be imported into a database or spreadsheet. To add a delimiter, you use the `-d` *delimiter* option to your `stats show` or `stats stop` command line. The `-d` option only has effect if your output is in column format.

Example output with specified delimiter: In the following example, the same counter is listed as for the column output example, except that it is comma-delimited.

```
cli> stats show -d , -c qtree:*:nfs_ops
Instance nfs_ops
              /s
vol1/proj1,265
vol3/proj2,12
```

About the stats preset files

NetApp provides some XML files that output a predetermined set of statistics that you can use without having to construct a script or type in a complicated command on the command line. The preset files are located in the /etc/stats/preset directory.

Using the preset files

To use a preset file, you add `-p` *filename* to your `stats show` or `stats stop` command line. You can also add counters on the command line. If any options you specify on the command line conflict with the preset file, your command line options take precedence.

You can also create your own preset files.

For more information about preset files, see the na_stats_preset(5) man page.

Improving NetApp system performance

About this section

This section describes configuration procedures that might improve your system's performance.

Balancing NFS traffic on network interfaces

Attach multiple interfaces on the NetApp system to the same physical network to balance network traffic among different interfaces. For example, if two Ethernet interfaces on the system 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 each of the clients generates about the same amount of traffic.

The NetApp system 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 NetApp system 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 NetApp system 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 NetApp system 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 NetApp system does very aggressive read-ahead.

Adding disks to a disk-bound aggregate

If you have a single traditional volume or single-aggregate NetApp system, use the `sysstat -u` command to determine the fraction of time that the busiest disk is active. If the fraction is greater than 80 percent, add disks to the traditional volume or aggregate using the `aggr add` command.

Sizing aggregates appropriately

When creating an aggregate or a traditional volume, be sure to provide enough data disks for its anticipated data access load. Performance problems due to insufficient data disks are especially noticeable for single-data-disk aggregates (either an aggregate/FlexVol combination or an aggregate bound to a traditional volume) (two disks for RAID-4 and three disks for RAID-DP).

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* on the NOW site at http://now.netapp.com/NOW/knowledge/docs/docs.shtml for recommended card placement.

Maintaining adequate free blocks and free inodes

If free blocks or free inodes make up less than 10 percent of the space 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 over 90 percent of blocks are used, increase blocks by adding disks to the volume's containing aggregate or by deleting snapshots. If fewer than 10 percent of inodes are free, increase inodes by deleting files or using the `maxfiles` command.

Optimizing file and volume layout

If read performance on a particular large file or LUN degrades over time, use the `reallocate` command to optimize its layout. If you add disks to an aggregate, use `reallocate` to redistribute the data equally across all of the disks in the aggregate.

Using oplocks for CIFS NetApp systems

Oplocks allow CIFS clients to cache more data locally, reducing traffic to the NetApp system. On a system running CIFS that is not in a database environment, set oplocks to On by using the `options cifs.oplocks.enable on` command. If the system 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 system 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

NetApp recommends that you call NetApp 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 NetApp system 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 NetApp system running CIFS. Use the `options nfs.tcp.recvwindowsize 64240` command to maximize the TCP window size on a system 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 NetApp system has multiple tape drives and a volume with 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.

Troubleshooting Tools 12

About this chapter

This chapter describes what to do when you experience serious problems with your NetApp system, and lists some tools available to you to help you understand and avoid problems.

Topics in this chapter

This chapter discusses the following topics:

- "NetApp system panics" on page 276
- "Error messages" on page 277
- "Other troubleshooting resources" on page 279

NetApp system panics

About NetApp system panics

If your NetApp system has a serious problem, such as a problem with the hardware or a bug in the system software, it panics. When a system panics, it performs the following actions:

- The system core is dumped into a core file, which is placed in /etc/crash.
- A panic message is output to the console and to /etc/messages.
- The NetApp system reboots.

The panic message contains important information that can help you and NetApp technical support determine what happened, and how you can prevent the panic from happening in the future.

What to do if your NetApp system panics

If your NetApp system panics, following these steps can help NetApp technical support troubleshoot the problem more quickly.

Step	Action
1	Access the panic message on the console messages or in the /etc/messages file.
2	On the NOW site, navigate to the Panic Message Analyzer tool at http://now.netapp.com/NOW/cgi-bin/pmsg/.
3	Copy the panic message and Data ONTAP version number into the Panic Message Analyzer tool to determine whether your panic was caused by a known software issue.
4	If the panic is due to a known issue that was fixed in a later release, and upgrading to that release is feasible, you can download the new release from the NOW site and upgrade to resolve the issue.
5	Otherwise, call NetApp technical support.

Note

If you have AutoSupport enabled, AutoSupport automatically alerts NetApp when your system panics.

Error messages

About error message logging

If a hardware, software, or configuration problem exists on your system that is not severe enough to cause a panic, the NetApp system logs a message to alert you to the problem. That message can be logged to the console, a file, or to a remote system, depending on how you have configured message logging.

For more information, see "Configuring message logging" on page 128.

Note

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 of important events.

Using the Syslog Translator to get more information about error messages

Error messages need to be relatively brief to avoid clogging the error logging system. Some messages have more information available through the Syslog Translator on the NOW site.

To get more information about an error message using the Syslog Translator, complete the following steps.

Step	Action
1	Go the NOW site and select Technical Assistance & Documentation > Syslog Translator.
2	In the Software field, select `Data ONTAP`.

Step	Action
3	Cut and paste the error message into the Search String field and click Translate. **Result:** If more information is available about the message you have received, it is displayed, including the following information: ◆ Severity ◆ Description ◆ Corrective action ◆ Related information ◆ Data ONTAP versions this message applies to ◆ Details about the syslog message ◆ Details about the SNMP trap initiated by this message

Accessing the Syslog Translator using FilerView

To access the Syslog Translator though FilerView, complete the following steps.

Step	Action
1	From FilerView, select Filer > Syslog Messages. **Result:** The /etc/messages file is displayed.
2	Click on any message displayed as a hot link to access the Syslog Translator for that message. If a message is not listed as a hot link, no further information is available from the Syslog Translator for that message.

Other troubleshooting resources

Using the NOW site for help with errors

The NOW site is a powerful resource to help you diagnose and solve problems with your NetApp system. The NOW site includes the following tools:

- Knowledgebase Solutions

  A database of technical tips and articles to help with specific errors and problems. To access this tool, select Service & Support to access the natural language search tool. Make sure that the Knowledgebase Solutions check box is selected.

  You can also browse the Knowledgebase by selecting Browse the Knowledgebase.

- Bugs Online

  NetApp provides information about known issues and any workarounds using this tool. To access Bugs Online, select Service & Support > Bugs Online & Release Tools.

  If you know the bug ID, you can view the information for that particular bug. Otherwise, you can use either the Bugs Online search capabilities or the natural language search as described for the Knowledgebase Solutions tool to search for a bug that matches your issue.

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 NetApp expansion slot. See expansion card.

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 system setup for managing the system. The setup program automatically configures the system 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 system to work with other clients in the same way it does with the administration host.

aggregate — A manageable unit of RAID-protected storage, consisting of one or two plexes, that can contain one traditional volume or multiple FlexVol volumes. For more information about aggregates, see the *Storage Management Guide*.

API — Application Programming Interface. A software toolkit designed to provide system access to external programs. Data ONTAP provides an API called ZAPI.

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 system's domain, or by the system itself, using its /etc/passwd file.

AutoSupport A system daemon that triggers e-mail messages from the customer site to NetApp or another specified e-mail recipient when there is a potential system problem. For more information about AutoSupport, see "Using AutoSupport" on page 145.

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.

CLI Command Line Interface. The Data ONTAP system prompt is an example of a Command Line Interface.

client A computer that shares files on a NetApp system.

cluster A pair of NetApp systems connected so that one system can detect when the other is not working and, if so, can serve the failed system data. When NetApp systems are clustered, each system is also referred to as a node.

cluster interconnect Cables and adapters with which the two NetApp systems in a cluster are connected and over which heartbeat and WAFL log information are transmitted when both systems are running.

cluster monitor Software that administers the relationship of NetApp systems in the cluster through the `cf` command.

community A name used as a password by the SNMP manager to communicate with the NetApp system agent.

console A terminal that is attached to a NetApp system's serial port and is used to monitor and manage NetApp system operation.

continuous media scrub A background process that continuously scans for and scrubs media errors on the NetApp system disks.

copy-on-write The technique for creating snapshots without consuming excess disk space.

degraded mode The operating mode of a NetApp system 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 NetApp system 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 NetApp system.

emulated NetApp system A software copy of a failed NetApp system that is hosted by its takeover NetApp system. The emulated NetApp system appears to users and administrators to be a functional version of the failed NetApp system. For example, it has the same name as the failed NetApp system.

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 NetApp system expansion slot. See adapter card.

expansion slot The slots on the NetApp system board into which you insert expansion cards.

failed NetApp system A physical NetApp system that has ceased operating. In a clustered configuration, it remains the failed NetApp system until a 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 NetApp system to the failed NetApp system, resulting in a return to normal operation; the reverse of takeover.

group A group of users defined in the NetApp system's /etc/group file.

heartbeat A repeating signal transmitted from one NetApp system to the other that indicates that the NetApp system is in operation. Heartbeat information is also stored on disk.

hot spare disk A disk installed in the NetApp system 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 NetApp system 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 NetApp system and in a UNIX file system.

interrupt switch A switch on some NetApp system 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 NetApp system The NetApp system 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 NetApp system that is used to store the cluster state information of a NetApp system. If that system stops operating, the takeover system uses the information in the mailbox disks in constructing a virtual NetApp system. Mailbox disks are also used as file system disks.

mail host The client host responsible for sending automatic e-mail to NetApp when certain NetApp system events occur.

maintenance mode An option when booting a NetApp system from a system boot disk. Maintenance mode provides special commands for troubleshooting hardware and configuration.

MultiStore™ An optional software product that enables you to partition the storage and network resources of a single NetApp system so that it appears as multiple NetApp systems on the network.

NDMP Network Data Management Protocol. A protocol that allows NetApp systems 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 card.

normal mode The state of a NetApp system when there is no takeover in the cluster.

NVRAM cache Nonvolatile RAM in a NetApp system, used for logging incoming write data and NFS requests. Improves system performance and prevents loss of data in case of a NetApp system or power failure.

NVRAM card An adapter card that contains the NetApp system's NVRAM cache.

NVRAM mirror A synchronously updated copy of the contents of the NetApp system NVRAM (nonvolatile random access memory) contents kept on the partner NetApp system.

panic A serious error condition causing the NetApp system 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 NetApp system, the other NetApp system in a cluster.

partner mode The method you use to communicate through the command-line interface with a virtual NetApp system during a takeover.

POST Power-on self-tests. The tests run by a NetApp system after the power is turned on.

qtree A special subdirectory of the root of a volume that acts as a virtual subvolume with special attributes. For more information about qtrees, see the *Storage Management Guide*.

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 systems 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 NetApp system models.

serial console An ASCII or ANSI terminal attached to a NetApp system's serial port. Used to monitor and manage NetApp system operations.

share A directory or directory structure on the NetApp system that has been made available to network users and can be mapped to a drive letter on a CIFS client.

SID Security identifier used by the Windows operating system.

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 NetApp system to tape while the NetApp system is in use.

system board A printed circuit board that contains a NetApp system's CPU, expansion bus slots, and system memory.

takeover The emulation of the failed NetApp system identity by the takeover NetApp system in a cluster; the opposite of *giveback*.

takeover NetApp system A NetApp system that remains in operation after the other NetApp system stops working and that hosts a virtual NetApp system that manages access to the failed NetApp system disk shelves and network connections. The takeover NetApp system maintains its own identity and the virtual NetApp system maintains the failed NetApp system identity.

takeover mode The method you use to interact with a NetApp system while it has taken over its partner. The console prompt indicates when the NetApp system 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 NetApp system.

UID User identification number.

Unicode A 16-bit character set standard. It was designed and is maintained by the nonprofit consortium Unicode Inc.

vFiler A virtual NetApp system you create using MultiStore, which enables you to partition the storage and network resources of a single NetApp system so that it appears as multiple NetApp systems on the network.

volume A file system. For more information about volumes, see the *Storage Management Guide*.

WAFL Write Anywhere File Layout. The WAFL file system was designed for the NetApp system to optimize write performance.

WINS Windows Internet Name Service.

workgroup A collection of computers running Windows NT or Windows for Workgroups that is grouped for browsing and sharing.

Index

Symbols

/etc directory 9, 51
/etc/hosts.equiv file 41, 51
/etc/log/auditlog file 132
/etc/messages file 52, 128
/etc/rc file 51, 135
/etc/syslog.conf file
    configuring message logging in 131
    file format and parameters of 129
/etc/usermap.cfg file, character coding of 53
/home file, contents of 50
/vol/vol0, root volume 49

Numerics

3DES, for SecureAdmin 165

A

access to filer *See* filer access
ACL 281
adapter card 281
address resolution 281
administration hosts
    adding 42
    definition of 40, 281
    deletion of 42
    reasons for using 40
    removing 42
    use of 40
    when you can add 40
    where they are specified 41
administrative level commands 21
administrator access, managing 85
administrator accounts
    changing the password of (passwd) 111
    reasons for creating 85
advanced level commands 21
aggregate snapshot management 116
aggregates
    aggr copy command 11
    aggr status command, description of 245
    aggregate state, displaying (aggr status) 245
    definition 281
    disk statistics, displaying (aggr status) 245
    performance improvements for disk-bound aggregates 271
    root option 59
API 281
ATM 281
auditlog file 85, 132
authentication
    definition of 282
    public key-based 175
    with SSH 165
    with SSL 168
AutoSupport
    about 146
    commands used 160
    configuring 152
    contents of email 159
    default transport protocol 146
    defined 146
    events that trigger e-mail 158
    mail host support for 147
    NetApp technical support and 146
    options 148
    options autosupport.option (configures AutoSupport) 152
    reboots and 127
    requirements for 147
    testing 152
    testing (options autosupport.doit) 152
    transport protocols 146
    troubleshooting 154
    when system reboots 127

B

banner message for Telnet sessions 30
big-endian 282
boot options
    FAS200 series appliances 69
    for F87 systems 67

booting
F825, F880, FAS series systems 74
F87 systems 67
FAS200 series systems 69
from command prompt 62
from diskette 64, 65, 66, 74, 76
from firmware prompt 62, 70
options for the F87 filer 67, 69
reasons to 64, 74
browsers, improving security through 168
Bugs Online 279

C

capabilities
assigning to roles 97
assignment to users 86
definition of 86
list of supported types 97
modifying others' 92
types of 97
cards, expansion, displaying information about 240
certificate-authority signed certificates 179
certificates
domain names and 182
generating 179
installing 179, 181, 183
regenerating 183
testing 181
types of 179
used by SSL protocol 179
cf giveback command 228
cf-config-check.cgi utility 234
change privileges, file ownership 46
character coding for configuration files 53
checksums, displaying information 245
CIFS
accessing /etc directory 54
accessing home directory 55
administrator accounts in 85
client, requirements to manage filer 41
definition of 282
editing configuration files using 57
CIFS share, definition of 288
CLI 282
client decryption 166
clients
CIFS, requirements 41
editing configuration file from 57
NFS, requirements 41
platforms supported by FilerView 35
SecureAdmin supported 165
cluster commands
cf (enables and disables takeover) 213
cf forcegiveback 229
cf forcetakeover (forces a takeover) 216
cf partner (displays partner's name) 212
cf status (determines status) 219
cf status (monitors status) 210
cf takeover (initiates takeover) 216
halt (halts filer without takeover) 214
partner (accesses emulated filer) 222
partner restore (performs restore on emulated filer) 227
sysconfig/aggr status (displays disk information) 213
takeover (description of all takeover commands) 216
cluster configuration checker 234
clusters
benefits of 186
characteristics of 188
checking configuration through a utility 234
command exceptions when accessing emulated filer 224
configuration information for 212
considerations 188
definition of 186
displaying disk information 213
displaying partner's name (cf partner) 212
emulated filer
accessing from takeover filer (partner) 222
accessing remotely 224
description of 222
dumps and restores 227
performing restores (partner restore) 227
enabling and disabling takeover (cf) 213
giveback
description of 191

executing 228
messages 199
mode 198
halting filer without takeover (halt) 214
hardware components described 187
LCD messages on takeover filer 235
licensing and unlicensing cluster system 215
mailbox disks 188
managing in normal mode 210
messages on failed filer 237
monitoring cluster status (cf status) 210
negotiated failover 218
network interfaces, role of 200
normal mode 193
operating modes 192
partner command 191, 222, 223
preparing to configure 189
status messages, descriptions of 211
takeover
messages indicating a 196
mode, description of 195
statistics, description of 220
when occurs 190
with SNMP 221
types of operation modes 192
command line editing
with editing mechanism 18
with history mechanism 18
commands
administrative level 21
advanced level 21
autosupport.option (sets autosupport options) 152
boot (boots the filer) 62
date (sets system date and time) 121
halt (halts the filer) 81, 82
license 119
options autosupport.doit (tests autosupport) 152
options wafl.root_only_chown (sets file ownerships changes) 46
passwd (changes administrative user password) 111
passwd (changes filer system password) 110
priv 21
privilege levels 21
reboot (reboots the filer) 80
rsh command list 33
savecore, what it does 126
stats 260
useradmin 85
CompactFlash cards
booting an F87 filer 67
booting F800 systems 74
booting FAS900 systems 74
booting the FAS200 series filers 69
checking the Data ONTAP version of 71, 77
description of 4
configuration
display, using sysconfig 240
message logging 128
of autosupport (options AutoSupport.option) 152
configuration files
/etc 50
about backup and cloning 139
accessing 16
backing up 139
cloning 140
comparing backups 141
descriptions 51
editing from CIFS client 57
editing from NFS client 56
editing from NFS client-setup 56
restoring 140
within /etc directory 51
configuring
HTTP services 77
TFTP services 78
core files 126
criticaltime (UPS option) 142

D

data access management 9, 12
data migration management 9
Data ONTAP, check version of 71
data organization management 8
data protection 9
data storage management 7, 8

DataFabric 2
date, setting filer time and 121
decryption, between client and filer 166
default directories 50
default root volume 49
deletion of administration hosts 42
device carrier 3
directories
    default permissions 50
disks, displaying statistical information for 240
domain names, changing filer 182
domainusers
    definition of 85
    deleting 107
    granting access to 90
    listing 102
DSA key pair, for SSH 175

E

editing mechanism, editing command line with 18
emulated filer
    accessing from takeover filer 222
    accessing remotely 224
    description of 222
    dumps and restores 227
    performing restores (partner restore) 227
encryption
    with SSH 166
    with SSL 168
encryption algoritims supported by SecureAdmin 165
environmental adapter (filer main unit) 5
error message logging, about 277

F

F800, booting with CompactFlash 74
F87 filer, booting 67
failover paths
    about enabling 203
    configuring addresses for 205
    enabling 207
    for DAFS cluster 202
    how the client identifies the 205
failover, negotiated
    determining status in a cluster 219
    enabling and disabling of 218
    in a cluster 218
FAS200 series systems
    booting 69
    preparing for Netboot startup 71
    starting through netboot option 72
    starting through the netboot option 72
FAS900, booting with CompactFlash 74
FAS900, booting with CompactFlash cards 74
file ownership change privileges 46
filer access 40
    /etc directory, accessing from CIFS client 54
    /etc directory, accessing from NFS client 54
    /home directory, accessing from CIFS client 55
    /home directory, accessing from NFS client 55
    from the console 26
    inaccessiblility, reasons for 138
    using FilerView 35
    using rsh command from a PC client 34
    using rsh command from a UNIX client 33
    with Remote Shell (rsh) 32
    with Telnet 28, 29
filer hardware
    components 4
    environmental adapter 5
    expansion slots 6
    LCD 5
    memory 4
    system board 4
filers
    booting 69
    booting the system (boot) 62
    changing domain name of 182
    decryption 166
    disabling licenses for 120
    displaying current licenses for (license) 120
    displaying overall statistics for 241
    displaying version of 240
    editing boot configuration file in 137
    enabling licenses for (license) 119
    halting (halt) 82
    improving performance in 270
    managing CIFS using FilerView 17
    performance improvements

aggregate sizing 271
avoiding inodes access time update 270
backup rate 273
balancing NFS traffic on interfaces 270
caching client data to reduce traffic 272
disk-bound volume 271
large transfer 272
maintain adequate free blocks and inodes 271
read-ahead 271
reallocate command 272
using TCP 270
rebooting the system (reboot) 80
rsh (Remote Shell) access to 32
savecore (saves the core file) 126
security
controlling file ownership changes (options wafl.root_only_chown) 46
controlling mount privileges (options nfs.mount_rootonly) 45
limiting Telnet access 43
options 47
setting date and time (date) 121
FilerView
accessing filer through 17, 35
accessing Syslog Translator through 278
description 35
Help system defined 39
supported by client platforms 35
using the interface 37
Web browser requirements 35
files, configuration 50
filestats command
about 250
options for 253
forcing a takeover 216
forcing giveback 229

G

generating certificates 179
giveback
forcing 229
mode, description of 191, 198
processes that prevent giveback 228
givebacks for clusters 191
groups
assigning roles to 93
assigning users to 88
definition of 86
deleting 107
listing 102
naming requirements 86
predefined 93
renaming 95
Windows special 87

H

history mechanism, editing command line with 18
HMAC, for SecureAdmin 165
host keys
changing the size of 172
determining sizes 170
setting 170
uses of 165
using with SSH 170
where stored 170
host support for AutoSupport, mail 147
hosts
definition of 40
deletion of administration 42
use of administration 40
when you can add 40
HTTP access to log files 55
HTTP services, configuring 77

I

installing certificates 179, 181
interface, use of FilerView 37

K

keys
host. *See* host keys
private and public. *See* RSA key pair
public-based authentication. *See* RSA key pair.
server. *See* server keys
session 165

used by SSH protocol 165
Knowledgebase Solutions 279

L

LCD
cluster messages, description of 235
on filer chassis 5
level commands
administrative 21
advanced 21
licenses 119, 120
log files, accessing using HTTP 55

M

mail host support for AutoSupport 147
mailbox disks 188
Manage ONTAP Developer SDK software 15, 16, 17
man-in-the-middle warning 173
manual (man) pages 20
message files, accessing using HTTP 55
message logging, configuring 128
mount privileges, controlling of (options nfs.mount_rootonly) 45
multiprotocol file and block sharing 7

N

naming requirements for useradmin command 86
NDMP 12
negotiated failover
description of 218
determining status (cf status) 219
NetApp clusters *See* clusters
NetApp storage system, defined 2
Netboot startup
preparing FAS200 series filers 71
server requirements 77
starting FAS200 series systems 72
using filers as Netboot servers 77
Network file service 7
network interfaces
role in a cluster system 200
NFS client
access to /etc directory 54
access to /etc/directory 55
editing configuration files from 56
editing configuration files using 56
requirements to manage filer 41
non-local users, granting access to 90
normal mode for clusters 193
NVFAIL 12
NVRAM
description of 4
halt command to save data to disk 81

O

obsolete domain names, and SSL 182
OpenSSH
for SecureAdmin 165
generating key pairs in 175
operating modes for clusters 192
operation modes, in a cluster 192
ownership change privileges, file 46

P

panics 276
password requirements, changing 111
passwords, changing (passwd) 110, 111
permissions of default directories (/etc, /home) 50
plexes, displaying information about 245
private and public keys. *See* RSA key pair 175
privilege levels for Data ONTAP commands 21
privileges, file ownership change 46
protocols. *See* SSH (Secure Shell) protocol and SSL (Secure Sockets Layer) protocol
public-key encryption 165, 166
PuTTY for SecureAdmin 165

Q

qtrees 287
quota file, character coding for 53

R

RAID 287
displaying statistics (aggr status) 245

displaying statistics (sysconfig -r) 240
displaying statistics (vol status) 247
rebooting the system
from the console 80
regeneration of certificates 183
reinitialization
of SSH 172
of SSL 183
Remote Management Controller 5
Remote Shell access to filer (rsh) 32
requirements
for Web browser 35
to manage filer on NFS clients 41
RMC 5
roles
assigning capabilities to 97
assigning to groups 93
creating 99
definition of 86
deleting 107
listing 102
naming requirements 86
predefined 96
root option for aggregates 59
root password, changing 110
root volume
changing 58
default name 49
directories contained within 50
RSA key pair
definition of 175
generating for SSH 1.x 175
generating for SSH 2.0 176
where stored 175, 176
RSA/DSA, for SecureAdmin 165
rsh (Remote Shell)
access to filer 32
using with Windows 34
rsh commands
accessing filer from a PC client 34
accessing filer from a UNIX client 33
format used with user name and password 32
list of 33
privilege levels 21
use with user names and passwords 32

S

secure connection, testing 181
Secure FilerView, improving security using 168
secure session, creating with SecureAdmin 166
SecureAdmin
authentication supported 165
creating a secure session with 166
displaying status of 184
encryption algorithims supported 165
improving security with SSH 165
improving security with SSL 168
managing SSH portion 169
managing SSL portion 179
security
controlling file ownership changes (options wafl.root_only_chown) 46
controlling mount privileges (options nfs.mount_rootonly) 45
improving using Secure FilerView 168
improving using SecureAdmin 164
improving using SSH 165
limiting Remote Shell access 44
limiting Telnet access 43
options 47
password options 111
self-signed certificates 179
serial ports (filer hardware) 6
server keys
changing the size of 172
setting 170
size guidelines for 170
uses of 165
using with SSH 170
server requirements for netboot 77
session keys, uses of 166
share, definition of 288
single partitions 62
slots, expansion (filer hardware) 6
SnapMirror 10
SnapRestore 10
snapshot 288
Snapshot software 9
snapshots, aggregate 116
SnapVault 10
SNMP protocol, in clustered takeover mode 221

SSH (Secure Shell) commands
secureadmin disable all 184
secureadmin disable ssh 172, 174
secureadmin enable all 184
secureadmin enable ssh 170, 172, 174
secureadmin setup -f ssh 172
secureadmin setup ssh 170
secureadmin status 184
SSH (Secure Shell) protocol
authentication with 165
creating a secure session with 166
determining host and server key size using 170
disabling or enabling 170, 174
encryption with 165
host keys. *See* host keys
improving security with 165
keys used by 165
managing 169
reinitializing 172
server keys. *See* server keys
session keys 166
setting up and starting 170
SSH Communications Security client, for SecureAdmin 165
SSL (Secure Sockets Layer) commands
secureadmin addcert ssl 179, 183
secureadmin disable all 184
secureadmin disable ssl 183
secureadmin enable all 184
secureadmin enable ssl 179
secureadmin enable ssl (enables SSL) 183
secureadmin setup -f ssl 183
secureadmin setup ssl 179
secureadmin status 184
SSL (Secure Sockets Layer) protocol
authentication with 168
certificates used with 179
disabling or enabling 183
improving security with 168
managing 179
reinitializing 183
setting up and starting 179
startup commands in /etc/rc 135
statistics commands
aggr status command, description of 245
checking expansion cards 240
displaying adapter information 243
displaying aggregate state statistics 245
displaying chassis environment status 257
displaying Data ONTAP version 240
displaying disk information
aggr status 245
storage show 243
sysconfig -d 240
vol status 247
displaying Fibre Channel driver statistics 259
displaying link statistics 259
displaying medium changer information 243
displaying overall filer information 241
displaying RAID and checksum information 240, 245, 247
displaying relative environment information 256
displaying relative physical drive position 259
displaying shelf environment status 257
displaying tape drive information 241
displaying tape information 244
displaying tape library information 241
displaying volume
information (sysconfig -v) 241
language (vol status) 247
displaying volume state statistics 247
environment, description of 256
Fibre Channel statistics, description of 259
storage command, description of 243
sysconfig command, description of 240
vol status command, description of 247
stats command
about 260
background mode 267
controlling output 267
counters 260
instances 260
objects 260
preset files 269
repeat mode 266
singleton mode 265
status commands
aggr status (displays aggregate state) 245
aggr status -d (displays disk statistics) 245

aggr status -r (displays RAID statistics) 245
environment chassis (displays shelf environment information) 257
environment command, description of 256
environment status (displays all filer environment information) 256
environment status shelf (displays shelf environment information) 257
fcstat device_map (displays relative physical drive position) 259
fcstat fcal_stats (displays fibre channel driver statistics) 259
fcstat link_stats (displays link statistics) 259
filestats command, description of 249
storage command, description of 243
storage hub (displays hub information) 243
storage show adapter (displays adapter information) 243
storage show disk (displays disk information) 243
storage show mc (displays medium changer information) 243
storage show port (displays switch port information) 243
storage show switch (displays switch information) 243
storage show tape (displays tape information) 244
storage stats tape (displays tape statistics) 244
sysconfig -c (checks expansion cards) 240
sysconfig -d (displays disk information) 240
sysconfig -m (displays tape drive information) 241
sysconfig -m (displays tape library statistics) 241
sysconfig -r (displays RAID information) 240
sysconfig -v (displays overall filr statistics) 241
sysconfig -V (displays volume statistics) 241
ups (displays UPS environment information) 257
vol status (displays volume state) 247
vol status -d (displays disk statistics) 247
vol status -l (displays volume language) 247
vol status -r (displays RAID statistics) 247
status, displaying SecureAdmin 184
storage system 2
support for AutoSupport, mail host 147
SyncMirror 11
syncronizing system time 123
Syslog Translator
  accessing through FilerView 278
  using 277
system
  board (filer hardware) 4
  date and time, setting 121
  management 12
  memory (filer main unit) 4
  panics 276
  password, changing 110
  rebooting, from the console 80

T

takeover
  in clusters 190
  mode with SNMP 221
  mode, description of 195
  statistics in 220
tape backup and restore 10, 11
tape drives, displaying statistics 241
tapes, displaying tape library statistics 241
Telnet
  access to filer 28
  configuring a timeout period 30
  console and Telnet sessions
    console and Telnet sessions 25
  limiting access to 43
  sharing console and Telnet sessions 25
  termination of session with the filer 29
TFTP
  enabling console logging of accessed files 79
  root directory, specifying 78
  services, configuring 78
time
  setting filer date and 121
  synchronizing 123
time servers, about 123
timed, using 123
troubleshooting

LCD messages on clustered takeover filer 235
messages on failed clustered filer 237
reasons to boot from diskette 64, 74

U

Uninterruptible Power Supply (UPS)
adding a device to be monitored 257
enabling or disabling monitoring of 257
UPS
management 142
shutdown process 142
use of administration hosts 40
use of FilerView interface 37
user account, changing password for 111
useradmin
examples 108
naming requirements 86
usermap file, character coding for 53
users
assigning to groups 88
capabilities and 86
changing passwords 111
creation example 108
creation examples 108
definition of 85
deleting 107
examples of creating 108
listing 102
modifying capabilities of 92
naming requirement 86

V

Vandyke SecureCRT, for SecureAdmin 165
version checking, Data ONTAP 71
vol copy 11
volumes
disk statistics, displaying (vol status) 247
vol status command, description of 247
volume language, displaying (vol status) 247
volume state, displaying (vol status) 247
volume statistics, displaying 241

W

WAFL 2
warnings
man-in-the-middle 173
obsolete domain names 182
warningtime (UPS option) 142
Windows
administrator accounts in 85
domain users, granting access to 90
network commands 16
special groups 87
using rsh with 34